Isometries of Hermitian Lattices over Discrete Valuation Rings

Simon Larose

Thesis submitted to the University of Ottawa in partial fulfillment of the
requirements for the degree of
Master of Science Mathematics and Statistics*

Department of Mathematics and Statistics
Faculty of Science
University of Ottawa

*The M.Sc. program is a joint program with Carleton University, administered by the Ottawa-Carleton Institute of Mathematics and Statistics

Abstract

In this thesis, we study the generation of the unitary group of a Hermitian lattice over a separable extension of discrete valuation rings by symmetries of the lattice. After introducing basic algebraic definitions, we give an exposition of the theory of Dedekind domains, with the goal of describing discrete valuation rings as their local rings. We then extend results about Hermitian lattices over valuation rings of non-Archimedean local fields to our setting. Finally, we extend a known result about the generation of the unitary group of certain 2-adic Hermitian lattices to lattices over arbitrary separable extensions of discrete valuation rings, and we count the number of symmetries that appear in the factorization given by our theorem in the case where the isometry is strictly diagonalizable.

Acknowledgements

The author was generously supported by a Natural Sciences and Engineering Research Council (NSERC) Canada Graduate Scholarship-Master's, the Government of Ontario's Ontario Graduate Scholarship Program, and Dr. Kirill Zaynullin's NSERC Discovery Grant (RGPIN-2022-03060) through the writing of this thesis.

I would like to thank Dr. Monica Nevins for her encouragement throughout my time as a student at the University of Ottawa. Her magical ability to find time and energy to help those who need it has left a profound impression on me. I would also like to thank Dr. Nevins and Dr. Antonio Lei for taking the time to read and review my thesis. Finally, I would like to thank my supervisor, Dr. Kirill Zaynullin, for introducing me to the world of research, and for his unwavering support of my academic projects.

Contents

Chapter 1

Introduction

Finding generators for groups by using the data of spaces on which they act is a fascinating subject at the intersection of algebra and group theory. In Euclidean space, orthogonal transformations, the linear maps that preserve distance, are products of reflections — distinguished isometries which are defined with respect to a single vector. This is true because of a beautiful connection between hyperplane reflections and the orthogonal structure of $\mathbb{R}^n$. In this thesis, we will explore the delightful interplay between isometries and the geometric fabric they inhabit. Our starting point is the following celebrated result of Élie Cartan and Jean Dieudonné.

Theorem 1.0.1 (Cartan-Dieudonné)**.** *[Lam05] Let K be a field of characteristic not equal to 2, and let (V, q) be a quadratic space over K. Then any isometry τ of the orthogonal group $O(V, q)$ of V can be written as a product of reflections.*

In [Sc50], Scherk extended this theorem by showing that a minimal factorization of an element $\tau \in O(V, q)$ in terms of reflections has length $n - \dim(\mathrm{Fix}(\tau))$, where $n = \dim(V)$ and $\mathrm{Fix}(\tau)$ is the 1-eigenspace of τ.

For Hermitian spaces over a field, such as $\mathbb{C}^n$ with the inner product induced by complex conjugation, a strikingly similar generation theorem in terms of rank 1 isometries holds. As one often does in algebra, we can investigate what happens when we move from fields to rings, and, so long as the base ring is equipped with an anti-automorphism of order at most 2, an involution, every object mentioned above has an analogue. In particular, we obtain a group of isometries — the unitary group of a

Hermitian lattice. Studying the geometry of the unitary group of certain Hermitian lattices is therefore our task.

Recently, interest in the generation theory of the unitary group of a Hermitian lattice over the valuation ring of a non-Archimedean local field was renewed. In [BHM22], the authors give algorithms to decompose isometries of a Hermitian lattice over such a ring. An algorithm not dissimilar to that of the classical Cartan-Dieudonné Theorem works in the case that the corresponding extension of fraction fields is unramified. When there is ramification, the question of generation is much more subtle and serves as the focus of their paper. The ramified case of the generation theorem of [BHM22] enabled the computation of the determinant group of some Hermitian lattices over global fields and formed a key step in their classification of certain automorphism groups of K3 surfaces in [BH23]. Our goal is to find the correct setting to generalize the simpler algorithm of the unramified case to a broader class of Hermitian lattices. We have found that this is possible over separable extensions of discrete valuation rings, and it will be the purpose of this thesis to explain how this occurs.

We now give an outline of the material. Chapter 2 is dedicated to developing the algebraic background necessary to define discrete valuation rings and separable extensions thereof. Sections 2.1 to 2.3 aim to remind the reader of some standard results in commutative algebra. We give a treatment of field traces in Sections 2.4 to prove that the integral closure of an integrally closed integral domain in a separable extension is finitely generated. In Section 2.5, we present separable algebras over a commutative ring, the correct generalization of separable field extensions to algebras over a ring. In Section 2.6, we define the algebra trace and prove a surjectivity result under separability hypotheses.

Chapter 3 introduces the Dedekind domains, central in algebraic number theory. Our goal is to explore their localizations at primes, which are known to specialists as discrete valuation rings. This chapter culminates in the presentation of arithmetic and geometric examples. In Section 3.1, we introduce the language of fractional ideals, define Dedekind domains, and prove that their ideals factor uniquely into products of primes. We also prove that Dedekind domains behave well under extension. In Section 3.2, we give some algebraic background on localization, and in Section 3.3, we describe how ideals factor, or ramify, under extension. We then describe the

local theory of Dedekind domains in Section 3.4. In particular, we introduce discrete valuation rings and their extensions. To finish the chapter, we provide examples. In Section 3.5, we introduce the p-adic numbers, and in Section 3.6, we give a geometric example.

Chapter 4 presents the main result of the thesis. We introduce rings with involution in Section 4.1, which, as mentioned above, broadly generalize the complex numbers with conjugation. In Sections 4.2 and 4.3, we define Hermitian forms over rings with involution and Hermitian lattices over these rings. In Section 4.4, we then specialize the theory of Hermitian lattices to those over quadratic separable extensions of discrete valuation rings. In Sections 4.5 and 4.6, we describe orthogonality relations in these Hermitian lattices, verifying a generalized version of the Gram-Schmidt algorithm and giving an invariant of the lattice, the Jordan decomposition. In Section 4.7, we show that the Jordan decomposition is very nearly a complete isometry invariant when our ring extension is separable — up to knowledge of norms of the involution. In the case that the base ring is the valuation ring of a non-Archimedean local field, we recover, with the help of a classic result from [Se79], that the Jordan decomposition does indeed determine the lattice. In Section 4.8, we define the unitary group of a Hermitian lattice and some distinguished elements thereof — symmetries and Eichler isometries. In Section 4.9, we prove the main result.

Theorem 1.0.2. *Let $\mathcal{O}_E/\mathcal{O}_K$ be a quadratic separable extension of discrete valuation rings with non-trivial standard involution σ. Let L be a Hermitian lattice over $\mathcal{O}_E$ of rank n. The unitary group of L is generated by symmetries. Any minimal factorization of a strictly diagonalizable isometry τ of L in terms of symmetries has length at most $2\operatorname{rk}(L_\tau) - \dim_a(\tau, [L])$, where $[L]$ is the Jordan decomposition afforded by the eigenvectors of τ.*

The theorem not only extends the algorithm of [BHM22] for unitary groups of separable p-adic Hermitian lattices to the more general setting of those over discrete valuation rings, but describes the length of the resulting factorization in terms of a new invariant — the anisotropic dimension $\dim_a(\tau)$ of an isometry — which we will define in this thesis. In particular, we give a new upper bound for the length of a minimal factorization of a strictly diagonalizable isometry, which we define as an element of the unitary group having an orthogonal basis of eigenvectors for the lattice.

We can summarize the contributions of this thesis as follows. The material of Chapters 2 and 3 is standard and is intuitively understood by experts in the field. These chapters have been designed to be expository, and most proofs are either elementary or reproductions of results from texts. Chapter 4 serves to solidify the written theory of Hermitian lattices over separable extensions of discrete valuation rings. Indeed, until Section 4.7, the bulk of the material of Chapter 4 has the purpose of verifying, or in some cases, proving anew, that certain foundational results that hold in the setting of non-Archimedean local fields of characteristic zero also hold for discrete valuation rings. Happily, much of the theory contained therein is easily extended from the theory of Hermitian lattices over non-Archimedean local fields, and, in Chapter 4, many results and their proofs are reproductions or very slight generalizations of results from [Jac62] and [OMe73]. These sources deal with valuation rings of non-Archimedean local fields of characteristic zero. An advantage of our treatment is that it verifies that the results mentioned hold independently of characteristic, when separability hypotheses are applied. While much of the verification was routine, certain proofs, such as those of Lemmas 4.6.5, 4.6.6, and 4.6.7 concerning the Jordan decomposition, are based on sketches of proofs given for similar results concerning orthogonal lattices over non-Archimedean local fields and required substantial effort. In Section 4.8, we then define our counting invariant, the anisotropic dimension of an isometry. This novel invariant, inspired by the treatment of unitary groups over fields in [Wa59], will allow us to predict the number of elementary isometries — symmetries — which appear when using the algorithm of Theorem 1.0.2, at least in the case of strictly diagonalizable isometries. Then, in Section 4.9, we finally put the theoretical framework we have constructed to use, verifying without difficulty that the algorithm of [BHM22] effectively factors isometries when we relax its original hypotheses. We also use our new invariant to count the number of symmetries which appear in the factorization of Theorem 4.9.1, when the isometry is strictly diagonalizable. We then describe an effective factorization method coming from the proof of the Theorem and provide examples of Cartan-Dieudonné decompositions over both a 2-adic ring and a geometric ring. To finish, we propose some future research directions.

This thesis provides fruitful directions for future work. Indeed, the study of the factorization length of an isometry into a product of symmetries is inspired by the study of geodesics in classical geometry, and one might seek an algorithm which yields a factorization of minimal length. This point of view motivated our definition of the

anisotropic dimension of an isometry. We are especially hopeful that the algorithm of Theorem 4.9.1 provides such a minimal decomposition, since it already does in the case that a strictly diagonalizable isometry has full anisotropic dimension. One can also look at certain non-commutative algebras over rings with involution of the second kind, the so-called Azumaya algebras. The group of unitary similitudes of these algebras is closely related to the unitary group of a Hermitian lattice. It would be natural to attempt to extend the work of this thesis to an Azumaya algebra having a separable quadratic extension of discrete valuation rings with involution as its center.

Chapter 2

Commutative algebra

We begin the thesis by giving an outline of some useful results in commutative algebra. In particular, Sections 2.1, 2.2, and 2.3 are intended as a reference, and their contents will be familiar to graduate students. Section 2.4 gives us some facts about field traces, introduces the trace bilinear form, and proves a key fact about integral closures of integrally closed domains. Section 2.5 introduces a new concept, separability of an algebra, which will constitute a key hypothesis that we will need to extend the Cartan-Dieudonné theorem. To finish the chapter, we define the trace of an algebra in Section 2.6 and prove a surjectivity result when the algebra is separable. In this chapter, many results are part of the mathematical lore, but the more difficult proofs are reproductions of those in the literature.

We now set our conventions regarding the basic theory of algebra for the remainder of this work, and remind the reader of foundational facts. Our reference for ring theory is [DF04, §7-§9]. In this thesis, all rings are unital, and unless stated otherwise, all rings are commutative. Consequently, ring homomorphisms map the identity to the identity. The non-commutative rings we will transiently meet all arise as R-algebras, where R is a commutative ring. A subring contains the multiplicative identity of the parent ring.

2.1 Rings and modules

This section is dedicated to outlining basic facts of the theory of rings and modules, which will no doubt be familiar to the reader. This section is intended as a reference for future chapters.

Definition 2.1.1. Let R be a ring. The spectrum of R, denoted $\mathrm{Spec}\,(R)$, is the set of prime ideals of R.

Definition 2.1.2. Let I and J be ideals in the ring R. The sum of the ideals I and J is the ideal

$$I + J = \{a + b \mid a \in I, b \in J\}.$$

The product of the ideals I and J is the ideal

$$IJ = \{\sum_{k=1}^{n} a_k b_k \mid a_k \in I, b_k \in J\}.$$

We can iterate these constructions. In particular, the k-fold product of the ideal I with itself is denoted I^k.

Lemma 2.1.3 (Mod $\mathfrak{p}$ irreducibility)**.** Let R be an integral domain, and let $f(X) \in R[X]$ be a non-constant monic polynomial. If there exists a prime ideal $\mathfrak{p}$ of R such that the image of $f(X)$ is irreducible in $(R/\mathfrak{p})\,[X]$, then $f(X)$ is irreducible in $R[X]$.

The following is the famous Corollary to Gauss' Lemma.

Proposition 2.1.4. Let R be a unique factorization domain with field of fractions K, and let $f(X) \in R[X]$ be a polynomial. If $f(X)$ is irreducible in $K[X]$, then $f(X)$ is irreducible in $R[X]$.

Two ideals I, J of a ring R are called comaximal if $I + J = R$.

Lemma 2.1.5 (Chinese Remainder Theorem)**.** Let $\mathfrak{p}_1, \ldots, \mathfrak{p}_k$ be pairwise comaximal ideals of the ring R. Then, for any choice of positive integers $e_1, \ldots, e_k$, we have

$$R/\mathfrak{p}_1^{e_1} \cdot \mathfrak{p}_k^{e_k} \cong R/\mathfrak{p}_1^{e_1} \times \cdots \times R/\mathfrak{p}_k^{e_k}.$$

Definition 2.1.6. Let R be a ring, and let A be a ring. We say that A is an R-algebra with structure map θ if there exists a unital ring homomorphism $\theta\colon R \to A$, When the choice of structure map is clear, we say that A is an R-algebra. A homomorphism in the category of R-algebras is a ring homomorphism that commutes with the structure map.

We now begin outlining some facts about modules over a commutative ring. Many proofs are omitted, and a good reference for the material of the rest of the section is [Fo17, §1].

Definition 2.1.7. Let A, B, C be objects in the same category of rings, R-modules, or R-algebras, and let $f\colon A \to B$ and $g\colon B \to C$ be morphisms. We say the diagram

$$0 \longrightarrow A \xrightarrow{f} B \xrightarrow{g} C \longrightarrow 0$$

is a short exact sequence if the image of each arrow is the kernel of the next, where the arrows to and from 0 are the canonical maps. We say the sequence splits if there exists a morphism $\varphi\colon C \to B$ such that $\varphi \circ g = id_C$, or equivalently, if there exists a morphism $\psi\colon A \to B$ such that $\psi \circ f = id_A$.

Remark 2.1.8. In the categories of rings, R-modules, or R-algebras, the short exact sequence above can be thought of as the isomorphism $B/A \cong C$ from the first isomorphism theorem, and the notions are (almost) interchangeable. The main advantage of representing the isomorphism induced by the quotient map as an exact sequence is that the sequence cleanly describes the morphisms involved.

We define finitely generated, finitely presented, and projective modules over a ring.

Definition 2.1.9. Let M be an R-module. We say that M is a finitely generated R-module if M is the homomorphic image of R^n for some $n \geq 0$. We say that M is finitely presented if M is a finitely generated R-module and the kernel of the generating homomorphism is finitely generated.

Definition 2.1.10. Let R be a ring, and P an R-module. We say P is projective as an R-module if one and thus all of the following equivalent conditions hold.

1. P is isomorphic to a direct summand of a free module.

2. Every short exact sequence of R-modules

$$0 \longrightarrow N \longrightarrow M \xrightarrow{g} P \longrightarrow 0$$

splits.

3. The covariant functor $\mathrm{Hom}_R(P, \_)$ is right exact.

4. For every homomorphism $\psi\colon P \to M$ and surjective homomorphism $f\colon N \to M$, there exists $\varphi\colon P \to N$ such that

$$\begin{array}{ccccc} & & & & P \\ & & \overset{\exists\varphi}{\swarrow} & & \downarrow{\scriptstyle \psi} \\ N & \xrightarrow[f]{} & M & \longrightarrow & 0 \end{array}$$

commutes.

We immediately deduce the following.

Lemma 2.1.11. Let R be a commutative ring, and let M be a finitely generated projective R-module. Then, M is a finitely presented R-module.

Proof. If $\varphi\colon R^n \to M$ is a surjection, then φ splits and $\ker(\varphi)$ is a direct summand of R^n. Therefore, $\ker(\varphi)$ is a homomorphic image of R^n, hence finitely generated. ∎

We state the following from [Fo17, p.25].

Lemma 2.1.12. Let R, S be commutative rings, and let $\phi\colon R \to S$ be a homomorphism. If the R-module M is projective, then the S-module $S \otimes_R M$ is projective.

We remind the reader of some key facts involving the localization of rings and modules that we will use in this thesis.

Lemma 2.1.13. Let R be a commutative ring, $S \subseteq R$ be a multiplicative set, and M an R-module. There is a natural isomorphism

$$S^{-1}R \otimes_R M \to S^{-1}M.$$

Corollary 2.1.14. Let M be a finitely generated projective R-module, and let S be a multiplicative set in R. The localization $S^{-1}M$ is a finitely generated projective $S^{-1}R$-module.

Proof. By Lemma 2.1.13, localization is a tensor product. We appeal to Lemma 2.1.12 and to the right exactness of the tensor functor. ∎

We introduce the property of flatness, and state that the localization functor is flat.

Definition 2.1.15. Let R be a commutative ring and let M be an R-module. We say that M is a flat R-module if the tensor functor $M \otimes_R {}^-$ is exact. We say an R-algebra S is flat if it is flat as an R-module.

Lemma 2.1.16. Let R be a commutative ring and $S \subseteq R$ be a multiplicative set. The localization $S^{-1}R$ is a flat R-module.

We define faithfulness of a module, then solve an exercise from [Fo17, p.54] which will be used to prove Lemma 2.1.23.

Definition 2.1.17. We say an R-module M is faithful if for all non-zero $r \in R$, there exists $m \in M$ such that $rm \neq 0$.

Lemma 2.1.18. Let R be a commutative ring, $S \subseteq R$ a multiplicative set in R, and M a finitely generated R-module. Then $S^{-1}M = 0$ if and only if there exists $s \in S$ such that $sM = 0$.

Proof. If there exists $s \in S$ such that $sM = 0$, then it is immediate from the definition of the localized module that $S^{-1}M$ is 0. Conversely, suppose that for all $x \in M$, we have that $sx = 0$ for some $s \in S$. Let $x_1, \ldots, x_m$ be a generating set for M, with elements $s_1, \ldots, s_m \in S$ such that $s_i x_i = 0$ for all $1 \leq i \leq m$. Since R is commutative, $\prod_{i=1}^{m} s_i$ annihilates M. ∎

Localization allows us to verify many properties of modules by verifying the equivalent property at each localization away from a prime ideal independently.

Lemma 2.1.19. Let R be a commutative ring and let M be a finitely generated R-module. The following are equivalent.

1. The module M is the zero module.
2. The localization $M_{\mathfrak{p}}$ is the zero module for every prime ideal $\mathfrak{p}$ of R.
3. The localization $M_{\mathfrak{m}}$ is the zero module for every maximal ideal $\mathfrak{m}$ of R.

Corollary 2.1.20. Let M and N be R-modules, and let $f\colon M \to N$ be an R-module homomorphism. We have that f is surjective (resp. injective) if and only if the localized map $f_{\mathfrak{m}}\colon M_{\mathfrak{m}} \to N_{\mathfrak{m}}$ is surjective (resp. injective) for every maximal ideal $\mathfrak{m}$ of R.

Proof. We can check that $\ker(f)_{\mathfrak{m}} = \ker(f_{\mathfrak{m}})$. Then, if $f_{\mathfrak{m}}$ is injective, we have that $\ker(f)_{\mathfrak{m}} = 0$ for every maximal ideal $\mathfrak{m}$ of R. By the Lemma, $\ker(f) = 0$. Similarly, $\operatorname{coker}(f) = 0$ if $\operatorname{coker}(f_{\mathfrak{m}}) = 0$ for all maximal ideals $\mathfrak{m}$. Since localization is exact, we have that f is surjective (resp. injective) implies that the corresponding sequence localized at $\mathfrak{p} \in \operatorname{Spec}(R)$ is exact, i.e., that $f_{\mathfrak{p}}$ is surjective (resp. injective). ∎

We need the following result on localization from [Fo17, p.73] to prove Lemma 2.1.23.

Proposition 2.1.21. Let S be a flat commutative R-algebra. Let M be a finitely presented R-module and let N be an R-module. The natural map

$$S \otimes_R \operatorname{Hom}_R(M, N) \to \operatorname{Hom}_S(S \otimes_R M, S \otimes_R N)$$

induced by the universal property of the tensor product is an isomorphism.

Corollary 2.1.22. Let S be a multiplicative set in R. Localization induces a natural isomorphism

$$S^{-1} \operatorname{Hom}_R(M, N) \to \operatorname{Hom}_{S^{-1}R}\left(S^{-1}M, S^{-1}N\right).$$

We prove the following technical Lemma, as in [Knu91, p.4]. This Lemma will eventually allow us to describe involutions on quadratic algebras over a ring.

Lemma 2.1.23. Let R be a commutative ring, and let A be an R-algebra. If A is finitely generated, projective, and faithful as an R-module, then $R \cdot 1_A$ is an R-module direct summand of A.

Proof. We want to show the structure map $\theta\colon R \to A$ admits a section as a map of R-modules. It suffices to show that the map $\mathrm{Hom}_R(A, R) \to \mathrm{Hom}_R(R, R)$ induced by θ is surjective. By the assumption that A is finitely generated and projective as an R-module, Corollary 2.1.22 says that localization at a maximal ideal $\mathfrak{m}$ of R commutes with $\mathrm{Hom}_R(\_, R)$ and we can consider the induced map $\mathrm{Hom}_{R_\mathfrak{m}}(A_\mathfrak{m}, R_\mathfrak{m}) \to \mathrm{Hom}_{R_\mathfrak{m}}(R_\mathfrak{m}, R_\mathfrak{m})$. Since A is faithful as an R-module, Lemma 2.1.18 says that $A_\mathfrak{m}$ is non-zero and therefore that $A_\mathfrak{m} \cong R_\mathfrak{m}^n$ as R-modules, for some $n \geq 1$. But $1_{A_\mathfrak{m}} = (1, \ldots, 1)$ and $A_\mathfrak{m}$ is free over $R_\mathfrak{m}$ so that the $R_\mathfrak{m}$-algebra map $\theta_\mathfrak{m}\colon R_\mathfrak{m} \to A_\mathfrak{m}$ admits a section as an $R_\mathfrak{m}$-module homomorphism. Hence, the induced map is surjective at an arbitrary maximal ideal $\mathfrak{m}$ of R. ∎

To finish the section, we state one of the most useful results in commutative algebra and three immediate corollaries, as in [AM69, p.21]. These results are each known, interchangeably, as Nakayama's Lemma. The last formulation gives us an idea of why the use of Nakayama's Lemma is ubiquitous: if R is a local ring with maximal ideal $\mathfrak{m}$, and M is a finitely generated R-module, then any basis for the $R/\mathfrak{m}$-vector space $M/\mathfrak{m}M$ lifts to a minimal generating set for M. This essentially allows us to treat finitely generated modules over local rings like vector spaces over a field by checking that a property holds over a suitably defined quotient $R/\mathfrak{m}$-module.

Theorem 2.1.24. *Let R be a commutative ring, I an ideal of R contained in every maximal ideal of R, and M a finitely generated R-module. If $IM = M$, then there exists $r \in R$ such that $r \equiv 1 \pmod{I}$ and $rM = 0$.*

Corollary 2.1.25. Let M be a finitely generated R-module and an ideal I contained in every maximal ideal of R. Then, $IM = M$ implies that $M = 0$.

Proof. Let $r \in R$ be such that $r \equiv 1 \pmod{I}$ and $rM = 0$. As I is contained in every maximal ideal, the maximal ideals of R/I are the images of the maximal ideals of R. The condition that r is 1 modulo I therefore means that r must be contained in no maximal ideal of R. So, r is a unit of R, and $rM = 0$ implies that $M = 0$. ∎

Corollary 2.1.26. Let M be a finitely generated R-module, N a submodule of M, and I an ideal contained in every maximal ideal of R. Then $M = IM + N$ implies that $M = N$.

Proof. We have that $IM + N$ is a submodule of M, so we can consider the quotient $(IM + N)/N$. Then, as $M = IM + N$, we have that $M = IM$ modulo N. By Corollary 2.1.25, M/N is equal to N/N, so M is contained in N. This implies $M = N$. ∎

Corollary 2.1.27. Let R be a local ring with maximal ideal $\mathfrak{m}$, and let M be a finitely generated R-module. If the images of $x_1, \ldots, x_k$ generate $M/\mathfrak{m}M$ as an $R/\mathfrak{m}$-vector space, then the $x_1, \ldots, x_k$ generate M as an R-module.

Proof. Let N be the submodule of M generated by $x_1, \ldots, x_k$. Then, $N + \mathfrak{m}M$ generates every equivalence class of $M/\mathfrak{m}M$, so that $M = N + \mathfrak{m}M$. By Corollary 2.1.26, $M = N$. ∎

2.2 Noetherian rings and modules

We define Noetherian rings and modules and present some basic facts, following [Ash02, §7.5]. Noetherian rings and modules satisfy very strong finite generation properties.

Definition 2.2.1. Let M be an R-module, and let $M_1, M_2, \ldots$ be an increasing sequence of R-submodules of M, so that

$$M_1 \subseteq M_2 \subseteq \ldots.$$

If there exists $i \in \mathbb{N}$ such that $M_i = M_j$ for all $j \geq i$, we say that the sequence $M_1, M_2, \ldots$ stabilizes. We say that the R-module M satisfies the ascending chain condition if every increasing sequence of R-submodules of M stabilizes.

We define the Noetherian property for modules.

Definition 2.2.2. Let R be a ring. We say that an R-module M is Noetherian if M satisfies the ascending chain condition. A ring R is Noetherian if it is Noetherian as a module over itself.

The rest of the section is dedicated to stating fundamental properties of Noetherian modules for future reference.

Proposition 2.2.3. Let M be an R-module. The following conditions are equivalent.

1. M is Noetherian.
2. Every non-empty collection of R-submodules of M has a maximal element with respect to inclusion.
3. Every R-submodule of M is finitely generated.

Proposition 2.2.4. If N is a submodule of the R-module M, then M is Noetherian if and only if N and M/N are Noetherian R-modules.

Corollary 2.2.5. If $M_1, \ldots, M_n$ are Noetherian R-modules, then so is $M_1 \oplus \cdots \oplus M_n$.

Proof. Put $M = M_1 \oplus \cdots \oplus M_n$ and $N = M_1 \oplus \cdots \oplus M_{n-1}$. We have that $M/N \cong M_n$ is Noetherian, and by induction, so is N. ∎

Corollary 2.2.6. If M is a finitely generated module over the Noetherian ring R, then M is Noetherian.

Proof. The finitely generated R-module M is a quotient of the Noetherian R-module R^n. ∎

Finally, we state the Hilbert Basis Theorem, which says that polynomial rings over Noetherian rings are Noetherian.

Theorem 2.2.7 (Hilbert Basis Theorem)**.** *If the ring R is Noetherian, then so is $R[X]$. In particular, $R[X_1, \ldots, X_r]$ is Noetherian.*

Proof. An elementary argument is given in [Eis95, p.27]. ∎

2.3 Integral extensions

The material for this section introduces basic results regarding integral extensions of commutative rings, and is taken from and follows the exposition in [AM69, §5].

Longer proofs are omitted. Integral extensions of $\mathbb{Z}$ are of particular interest in algebraic number theory, because these rings are those which contain the solutions to integer polynomial equations in a fixed algebraic closure of $\mathbb{Q}$.

Definition 2.3.1. Let A be a subring of the ring R. We say that $\alpha \in R$ is integral over A if there exists a monic polynomial $f(X) \in A[X]$ such that $f(\alpha) = 0$. In this case, we say that $f(X) = 0$ is an equation of integral dependence for α over A.

Proposition 2.3.2. Let A be a subring of R, and let $\alpha \in R$. Then, the following are equivalent.

1. α is integral over A.
2. $A[\alpha]$ is a finitely generated A-submodule of R.
3. α is contained in a subring B of R containing A such that B is a finitely generated A-module.

Lemma 2.3.3. Let A be a subring of R, with $\alpha_1, \ldots, \alpha_n \in R$. If α_1 is integral over A, and α_i is integral over $A[\alpha_1, \ldots, \alpha_{i-1}]$, for all $2 \leq i \leq n$, then $A[\alpha_1, \ldots, \alpha_n]$ is a finitely generated A-module.

Proposition 2.3.4. Suppose A, B, C are subrings of R. If C is integral over B and B is integral over A, then C is integral over A.

Proof. Let $\alpha \in C$, and consider a monic $f(X) \in B[X]$ with $f(\alpha) = 0$. Write $f(X) = X^n + \sum_{i=0}^{n-1} b_i X^i$ for coefficients $b_i \in B$. Each b_i is integral over A by assumption, and is thus integral over $A[b_0, \ldots, b_{i-1}]$. Since $f(\alpha) = 0$, we have that α is integral over $A[b_0, \ldots, b_{n-1}]$. By Lemma 2.3.3, the subring $A[b_0, \ldots, b_{n-1}, \alpha]$ of C containing α must be a finitely generated A-module. This is one of the equivalent conditions of Proposition 2.3.2 for α to be integral over A. ∎

Definition 2.3.5. Let A be a subring of the ring R. The integral closure of A in R is the set of elements of R integral over A. We say that A is integrally closed in R if A is equal to its integral closure in R. We say that the integral domain A is integrally closed if A is integrally closed in its field of fractions.

Proposition 2.3.6. Let A be a subring of the ring R. The integral closure B of A in R is a subring of R.

Proof. Let $x, y \in R$ be integral over A. Since x is integral over A and y is integral over $A[x]$, $A[x, y]$ is a finitely generated A-module by Proposition 2.3.3. We apply Proposition 2.3.2 to conclude that every element of $A[x, y]$ is integral over A. In particular, $x - y$ and xy are integral over A. ∎

Proposition 2.3.7. The integral closure B of a subring $A \subseteq R$ is integrally closed in R.

Proof. The integral closure C of B in R is integral over B. We apply Proposition 2.3.4 to get that C is integral over A, so that $C \subseteq B$. ∎

Lemma 2.3.8. Unique factorization domains are integrally closed.

Proof. We adapt the proof of the rational roots theorem to show the unique factorization domain A is integrally closed in its field of fractions K. Let $x = \frac{r}{s} \in K$ be integral over A, where $r, s \in A$ are coprime. If we have $a_0, \ldots, a_{n-1} \in A$ such that

$$(\frac{r}{s})^n + a_{n-1}(\frac{r}{s})^{n-1} + \cdots + a_0 = 0,$$

we can multiply through by s^n to obtain

$$a_{n-1}s(r^{n-1}) + \cdots + a_0 s^n = -r^n.$$

In this case, s divides r^n. Since r and s are coprime in the UFD A, this means that s must be a unit of A, so that x is contained in A. ∎

Example 2.3.9. Let $K = \mathbb{Q}(\alpha)$ be any finite extension of $\mathbb{Q}$, and let B be the integral closure of $\mathbb{Z}$ in K. The integral domain $\mathcal{O}_K := B$ is known as the ring of integers of K. By Proposition 2.3.7, $\mathcal{O}_K$ is integrally closed. Not every ring of integers satisfies unique factorization of elements. For example, it can be shown that $2 \cdot 3$ and $(1 - \sqrt{-5}) \cdot (1 + \sqrt{-5})$ are two distinct prime factorizations in $\mathbb{Z}\left(\sqrt{-5}\right)$, the ring of integers of $\mathbb{Q}\left(\sqrt{-5}\right)$. However, all is not lost. As we will see later, all rings of this type satisfy unique factorization of ideals. We also note that, more generally, the notation $\mathcal{O}_K$ is used to denote an integral structure in a domain.

2.4 The trace form and the structure of the ring of integers

We recall that a finite field extension E/K of degree n is endowed with the structure of a K-vector space $V \cong K^n$, and that left multiplication $\ell_\alpha \colon E \to E$ by every $\alpha \in E$ is a K-linear map on $E = K^n$. In this section, we provide details on the norm and trace of a field extension, with the goal of defining the trace form on E. The facts and proofs given in this section can be found in [Ash02, §7.3].

Definition 2.4.1. Let E/K be a finite field extension, and let $\alpha \in E$. We define the norm of α relative to the extension E/K to be the determinant of the K-linear map ℓ_α given by left multiplication by α on E. We define the trace of α to be the trace of the linear map ℓ_α. We denote the norm of α relative to E/K by $\mathrm{N_{E/K}}(\alpha)$ and the trace of α relative to E/K by $\mathrm{Tr_{E/K}}(\alpha)$.

Definition 2.4.2. Let E/K be a finite field extension, and let $\alpha \in E$. We define the characteristic polynomial of α to be the characteristic polynomial of a transformation matrix for $\ell_\alpha \colon E \to E$. The characteristic polynomial of $\alpha \in E$ (relative to E/K) is denoted by $\chi_{E/K,\alpha}(X) \in K[X]$.

We informally recall that the characteristic polynomial, norm, and trace of α above do not depend on the choice of basis for E/K. In addition, the coefficients of the characteristic polynomial, the norm, and the trace all belong to the base field K, as these quantities are obtained by arithmetic operations on matrix entries in K.

Proposition 2.4.3. Let E/K be a finite extension of degree n, and let $\alpha \in E$ with minimal polynomial $m_\alpha(X) \in K[X]$. Then

$$\chi_{E/K,\alpha}(X) = m_\alpha(X)^r,$$

where $r = [E : K(\alpha)]$.

We give a formula for the norm and trace in terms of K-monomorphisms. Note that in the case that E/K is Galois, these monomorphisms are the elements of the Galois group of E/K.

Proposition 2.4.4. Let E/K be a separable extension of degree n, and let $\sigma_1, \ldots, \sigma_n$ be the distinct K-monomorphisms of E into an algebraic closure $\overline{E}$. Then

$$\mathrm{Tr}_{\mathrm{E/K}}(\alpha) = \sum_{i=1}^{n} \sigma_i(\alpha) \quad \text{and} \quad \mathrm{N}_{\mathrm{E/K}}(\alpha) = \prod_{i=1}^{n} \sigma_i(\alpha)$$

for all $\alpha \in E$.

Proposition 2.4.5. Let E be any finite extension of K, the field of fractions of the integral domain A. Suppose that $\beta \in B$, the integral closure of A in E. The coefficients of the minimal polynomial $m_\beta(X)$ of β are integral over A.

Proof. We have that $m_\beta(X) = \Pi_{i=1}^{m}(X - \beta_i)$ splits into linear factors over an algebraic closure C of K. Set $\beta_1 = \beta$. For $2 \leq i \leq m$, there exist K-isomorphisms $\sigma_i \colon K(\beta) \to K(\beta_i)$, so that if $\sum_{j=0}^{\ell} a_j \beta^j = 0$ is an equation of integral dependence for β, we have that

$$0 = \sum_{j=0}^{\ell} a_j \sigma_i(\beta^j) = \sum_{j=0}^{\ell} a_j \beta_i^j$$

is an equation of integral dependence for each $2 \leq i \leq n$. Therefore, the coefficients of

$$m_\beta(X) = \prod_{i=1}^{m} (X - \beta_i)$$

are sums and products of elements integral over A, so that by Proposition 2.3.6 these coefficients are themselves integral over A. ∎

Corollary 2.4.6. The trace, norm, and more generally, the coefficients of the characteristic polynomial $\chi_\beta(X)$ of β with respect to E/K are integral over A. In particular, if A is integrally closed, these quantities belong to A.

Proof. The above belong to the field of fractions K of A. We then combine Proposition 2.3.6 with the fact that $\chi_\beta(X) = (m_\beta(X))^d$, where $d = [E : K(\beta)]$. The result follows from the fact that the norm and the trace are obtained as coefficients of $\chi_\beta(X)$. ∎

We would now like to describe the relationship between an integrally closed integral domain A and its integral closure B in a finite separable extension E of the field of fractions K of A. As we saw in Proposition 2.3.2, we have that B contains

every subring of E which is also finitely generated as an A-module. It turns out that B is itself finitely generated as an A-module, when E/K is finite and separable. The goal of this section is to explain why this happens. This exposition in this section and the proofs within are as in [Mil20, §2, 2.23-2.28], except when stated otherwise. To begin, we will need to define the trace form.

Definition 2.4.7. Let E/K be a field extension of degree n. The map

$$\mathrm{B_T}\colon E \times E \to K$$

defined by

$$(\alpha, \beta) \mapsto \mathrm{Tr_{E/K}}(\alpha\beta)$$

is called the trace form of the extension E/K. This map is a symmetric K-bilinear form. We define the discriminant of the extension E/K to be the discriminant of the trace form $\mathrm{B_T}$, so that

$$\mathrm{disc}\,(E/K) := \det\left(\left[\mathrm{Tr_{E/K}}\,(x_i x_j)\right]\right)$$

with respect to any K-basis $\{x_1, \ldots, x_n\}$ of E/K, considered as an element of

$$K^\times/(K^\times)^2 \cup \{0\}.$$

Definition 2.4.8. Let $A \subseteq B$ be an inclusion of rings. Suppose that B is a free A-module of rank m, and let $\beta_1, \ldots, \beta_m$ be elements of B. Their discriminant is the quantity

$$\mathrm{disc}(\beta_1, \ldots, \beta_m) := \det([\mathrm{Tr_{B/A}}(\beta_i\beta_j)]).$$

The discriminant of B over A is the ideal of A generated by all $\mathrm{disc}(\beta_1, \ldots, \beta_m)$, where the generators run over all m-tuples of elements of B.

Lemma 2.4.9. Suppose that $A \subseteq B$ are rings, with B a free A-module of rank m. If $\gamma_1, \ldots, \gamma_m \in B$ can be written as $\gamma_i = \sum_{j=1}^m a_{ij}\beta_j$, then

$$\mathrm{disc}(\gamma_1, \ldots, \gamma_m) = \det(a_{ij})^2 \cdot \mathrm{disc}(\beta_1, \ldots, \beta_m).$$

If the β_i and γ_j each form an A-basis of B and $\gamma_i = T\beta_i$, then their discriminants must vary by a unit, as the matrix T must be invertible. The previous lemma then shows that their discriminants must vary up to a square, so that the discriminant of the free A-module B is well defined in $A^\times/(A^\times)^2$.

Proposition 2.4.10. Let E be a separable extension of K of degree m and let $\sigma_1, \ldots, \sigma_m$ be the distinct K-monomorphisms of E into $\overline{E}$. Then for every basis $\beta_1, \ldots, \beta_m$ of E/K, we have

$$\operatorname{disc}(\beta_1, \ldots, \beta_m) = \det(\sigma_i(\beta_j))^2.$$

In particular, $\operatorname{disc}(\beta_1, \ldots, \beta_m)$ is non-zero.

Corollary 2.4.11. If E/K is a finite separable extension, the trace form $\mathrm{B_T}\colon E \times E \to K$ is non-degenerate.

Lemma 2.4.12. Let A be an integral domain with field of fractions K, and let E be a finite separable extension of K. If B is the integral closure of A in E, there exists a basis for E/K consisting of elements of B.

Proof. Let $\beta_1, \ldots, \beta_n$ be a basis for E/K. We show that $k_i\beta_i$ is in B for appropriate $k_i \in A$. Since a set consisting of non-zero scalar multiples of a basis for E/K remains a basis for E/K, this will demonstrate the claim.

We have that β_i is algebraic over K, so that for $1 \leq j < m$, there exist $\frac{p_{i,j}}{q_{i,j}}$ with $p_{i,j}, q_{i,j} \in A$, such that

$$0 = \sum_{j=0}^{m-1} \frac{p_{i,j}}{q_{i,j}} \beta_i^j + \beta_i^m.$$

Setting $k_i = \Pi_{j=1}^m q_{i,j}$ and multiplying the above linear dependence relation by k_i^m, we observe that

$$0 = \sum_{j=0}^{m-1} k_i^m \frac{p_{i,j}}{q_{i,j}} \beta_i^j + k_i^m \beta_i^m = \sum_{j=0}^{m-1} k_i^{m-j} \frac{p_{i,j}}{q_{i,j}} (k_i\beta_i)^j + (k_i\beta_i)^m.$$

But since k_i clears denominators, the coefficients $k_i^{m-j}\frac{p_{i,j}}{q_{i,j}}$ are in A, so that $k_i\beta_i$ is integral over A. Since B is the integral closure of A in E, we have $k_i\beta_i \in B$, so that $\{k_1\beta_1, \ldots, k_n\beta_n\}$ is a basis in B for E/K. ∎

Proposition 2.4.13. Let A be an integrally closed integral domain with field of fractions K, and let B be the integral closure of A in a separable extension E/K of degree m. There exist free A-submodules M and M' of E such that $M \subseteq B \subseteq M'$.

Proof. By Lemma 2.4.12, E/K has a basis $\{\beta_1, \ldots, \beta_m\}$ with all β_i in B. Since E/K is separable, the trace form is non-degenerate by Corollary 2.4.11, so there exists a dual basis $\{\gamma_1, \ldots, \gamma_m\}$ for $\{\beta_1, \ldots, \beta_m\}$. We set

$$M = A\beta_1 \oplus \cdots \oplus A\beta_m$$

and

$$M' = A\gamma_1 \oplus \cdots \oplus A\gamma_m.$$

We have that $M \subseteq B$ by construction, so it remains to show that $B \subseteq M'$. Let $\beta \in B$. We can write $\beta = \sum_{j=1}^{m} b_j\gamma_j$. Since A is integrally closed and $\beta\beta_i \in B$ for all i, Corollary 2.4.6 tells us that $\mathrm{Tr_{E/K}}(\beta\beta_i)$ is in A. Since the γ_j are dual to the β_i with respect to the trace form, we have that

$$\mathrm{Tr_{E/K}}(\beta\beta_i) = \mathrm{Tr_{E/K}}((\sum_{j=1}^{m} b_j\gamma_j)\beta_i) = \sum_{j=1}^{m} b_j\delta_{i,j} = b_i,$$

so that b_i is in A. This shows that $B \subseteq M'$. ∎

Corollary 2.4.14. If A is an integrally closed Noetherian domain, then B is a finitely generated A-module.

Proof. We have that B is contained in the Noetherian A-module M', so B is a finitely generated A-module. ∎

Corollary 2.4.15. If A is a principal ideal domain, then B is a free A-module of rank m.

Proof. Since A is a PID, A is Noetherian. Applying the previous corollary, it remains to show that B is free of rank m. We have that B is a submodule of the free A-module M', where A is a PID, so that B is free with rank at most m. Since $M \subseteq B$ is a free A-module of rank m, B has rank at least m, which forces equality. ∎

Corollary 2.4.16. The integral closure B of A in E is the largest subring of E finitely generated as an A-module.

Proof. We have just shown that B is finitely generated over A. Apply Proposition 2.3.2 and the definition of the integral closure of A in E to see that every other subring of E which is finitely generated as an A-module is contained in B. ∎

As in [Ash02, Prop. 7.5.13], we show that integral closures are Noetherian.

Corollary 2.4.17. Let A be integrally closed with field of fractions K, and let E/K be a separable extension of K. If A is a Noetherian ring, then the integral closure B of A in E is a Noetherian ring.

Proof. By the Proposition, B is a submodule of a free A-module M' of finite rank. But A is Noetherian, so M' is Noetherian. By Proposition 2.2.3, B is a Noetherian A-module. If N is any B-submodule of B, then N is by restriction an A-submodule of the Noetherian A-module B. This means N is finitely generated as an A-module, and therefore as a B-module. By Proposition 2.2.3, B is a Noetherian ring. ∎

2.5 Separable algebras

Since we want to understand extensions of rings, we generalize the concept of a separable field extension to R-algebras. The connection between these two concepts is not immediately obvious, and the reader can treat the results presented in this section as a black box for the purposes of the thesis. We let R be a commutative ring, but the R-algebra A is not necessarily commutative in this section. Unless stated otherwise, the material and proofs are as in [Fo17].

Definition 2.5.1. Let A be an R-algebra. The R-algebra

$$A^e = A \otimes A^{op}$$

is the enveloping algebra of A.

We endow A with the structure of a left A^e module by setting

$$\sum_{i=1}^{m} a_i \otimes b_i \cdot x = \sum_{i=1}^{m} a_i x b_i$$

for all $a = \sum_{i=1}^{m} a_i \otimes b_i \in A^e$ and $x \in A$. Notice that if we set $x = 1_A$, we recover the multiplication of A and obtain a map

$$\mu \colon A^e \to A$$

defined by

$$\mu\left(\sum_{i=1}^{m} a_i \otimes b_i\right) = \sum_{i=1}^{m} a_i b_i.$$

Since A is unital, μ is onto. So, we get a short exact sequence

$$0 \longrightarrow J \longrightarrow A^e \xrightarrow{\mu} A \longrightarrow 0 \tag{2.5.1}$$

of A^e-modules, where J is the kernel of μ. This construction will allow us to define separability. To this end, we begin by solving an exercise from [Fo17, p.117].

Proposition 2.5.2. Let R be a commutative ring, and let A be an R-algebra. The following are equivalent.

1. A is projective as a left A^e-module.
2. The short exact sequence (2.5.1) splits.
3. There is an element $e \in A^e$ such that $\mu(e) = 1$ and $Je = 0$.
4. There is an idempotent $e \in A^e$ such that J is the principal left ideal generated by $(1 - e)$.

Proof. The first two points are equivalent by construction of the short exact sequence and the definition of a projective A^e-module. If (2.5.1) splits via $\nu\colon A \to A^e$, then set $e = \nu(1)$. For $j \in J$, we then have

$$je = j\nu(1) = \nu(j \cdot 1) = \nu(\mu(j)) = 0.$$

Now, such an element e satisfies $\mu(e) = 1 = \mu(1 \otimes 1)$, so that $e = 1 + j$, for some $j \in J$. This tells us that e is idempotent, since

$$e^2 = e + je = e + 0 = e.$$

Moreover, if $j \in J$, then $j(1 - e) = j - je = j$, by the assumption that $Je = 0$. So, $J = R(1 - e)$. If $e \in A^e$ is an idempotent such that J is generated as an ideal by $(1 - e)$, then $(1 - e) \in J$, so that $\mu(1 - e) = 0$. Then, $1 = \mu(1) = \mu(e)$. Moreover,

$$(1 - e)e = e - e^2 = e - e = 0,$$

so that right multiplication by e annihilates the ideal J generated by $(1 - e)$.

To show A is projective, we define the A^e-module homomorphism $\nu\colon A \to A^e$ by $\nu(1) = e$. We have just shown that ν has a left inverse, so ν is injective. Therefore, we can identify A with $im(\nu) = A^e e$. Now, we can write arbitrary $a \in A^e$ as $a = a(1-e)+ae$. If $a \in \ker(\mu)\cap A^e e$, then there exist $r, s \in A^e$ such that $re = a = s(1-e)$. Then, $re^2 = s(1-e)e = 0$ and $re^2 = re = a$. So, $a = 0$. This shows that

$$A^e = J \oplus \mathrm{im}(\nu) \cong J \oplus A.$$

So, A is projective as an A^e-module. ∎

Corollary 2.5.3. The R-algebra A is separable if and only if there exists an idempotent $e \in A^e$ such that $\mu(e) = 1$, $J = A^e(1-e)$, and $Je = 0$.

Proof. The proof of the Proposition shows that the same e satisfies the last two points of the statement. ∎

Definition 2.5.4. If the R-algebra A satisfies any of the equivalent conditions of Proposition 2.5.2, we say that A is separable over R. The element $e \in A^e$ of the Proposition is called a separability idempotent for A.

We solve the following important exercise from [Fo17, p.119].

Lemma 2.5.5. Let R be a commutative ring, A an R-algebra, and J the kernel of the multiplication map $\mu\colon A^e \to A$. Then $J = \ker(\mu)$ is generated as a left ideal by the set

$$\{a \otimes 1 - 1 \otimes a \mid a \in A\}.$$

Proof. Let $y = \sum_{k=1}^m a_k \otimes b_k \in J$. Then, $0 = \mu(\sum_{k=1}^m a_k \otimes b_k) = \sum_{k=1}^m a_k b_k$. For each $1 \le j \le m$, we have that $-a_j b_j = \sum_{k\neq j}^m a_k b_k$. Notice that $(a_j b_j \otimes 1)$ is a preimage of $a_j b_j$ and $(1 \otimes -a_j b_j)$ is a preimage of $-a_j b_j$ under μ. So,

$$\sum_{k=1}^m a_k \otimes b_k = a_j b_j \otimes 1 - 1 \otimes a_j b_j + x_j$$

for some $x_j \in J$. For fixed j, this equation is equivalent to

$$-a_j b_j \otimes 1 + a_j \otimes b_j = -\sum_{k\neq j}^m a_k \otimes b_k + a_j b_j \otimes 1 + x_j,$$

and therefore

$$a_j \otimes 1(-b_j \otimes 1 + 1 \otimes b_j) = -\sum_{k \neq j}^{m} a_k \otimes b_k + a_j b_j \otimes 1 + x_j.$$

This shows the right hand side is contained in $I_j := \langle b_j \otimes 1 - 1 \otimes b_j \rangle$. Summing these equations over all j, we have that

$$\sum_{k=1}^{m} a_k \otimes b_k - \sum_{k=1}^{m} a_k b_k \otimes 1 = \sum_{k=1}^{m} i_k$$

with each $i_k \in I_k$. But

$$\sum_{k=1}^{m} a_k b_k \otimes 1 = (\sum_{k=1}^{m} a_k b_k) \otimes 1 = 0,$$

by the assumption that $y \in \ker(\mu)$. So, $y = \sum_{i=1}^{k} i_k$. This is what we set out to prove. ∎

Example 2.5.6. The $\mathbb{R}$-algebra $\mathbb{C}$ is separable and $e = \frac{1}{2}(1 \otimes 1 - i \otimes i)$ is a separability idempotent for $\mathbb{C}/\mathbb{R}$. We can easily check that $e^2 = e$ and $\mu(e) = 1$. We show that $Je = 0$. By Lemma 2.5.5, it suffices to check that $(z \otimes 1 - 1 \otimes z)e = 0$ for all $z \in \mathbb{C}$. Let $z = a + bi$, $a, b \in \mathbb{R}$, be arbitrary in $\mathbb{C}$. Seeing as the tensor product is taken over $\mathbb{R}$, we have that $a \otimes 1 - 1 \otimes a = 0$. Therefore,

$$\begin{aligned}(a + bi \otimes 1 - 1 \otimes a + bi)e &= ((a \otimes 1 - 1 \otimes a) + (bi \otimes 1 - 1 \otimes bi))e \\ &= \frac{b}{2}(i \otimes 1 - 1 \otimes i) - (i^2 \otimes i + i \otimes i^2) \\ &= \frac{b}{2}(i \otimes 1 - i \otimes 1 - 1 \otimes i + 1 \otimes i) = 0.\end{aligned}$$

As we will see by the end of the section, a field extension is separable as an algebra over the base field if and only if it is a separable field extension.

Definition 2.5.7. Let A be an R-algebra, and let M be an A-bimodule. We say that M is a two-sided A/R module if the action by R through A satisfies

$$rx = (r \cdot 1)x = x(r \cdot 1) = xr,$$

i.e., that the R-action is equal on both sides.

Remark 2.5.8. If A is an R-algebra, a left A^e-module is the same as a two-sided A/R-module. Indeed, if M is a left A^e-module, then

$$ax := (a \otimes_R 1)x$$

and

$$xa := (1 \otimes_R a)x$$

endows M with the structure of a two-sided A/R-module because the tensor product is taken over R. The same equations endow a two-sided A/R-module with the structure of a left A^e-module.

There is also a functorial definition of separability, which will be of use. We introduce this point of view.

Definition 2.5.9. Let M be a two-sided A/R-module. We call the left R-module

$$M^A := \{x \in M \mid ax = xa \text{ for all } a \in A\}$$

the centralizer of A in M.

Lemma 2.5.10. Viewing left A^e-modules as two-sided A/R modules, the assignment

$$M \mapsto M^A$$

together with the canonical restriction map on module homomorphisms define a covariant functor from left A^e-modules to left R-modules.

Lemma 2.5.11. Let M be a left A^e-module. The map

$$\mathrm{Hom}_{A^e}(A, M) \to M^A$$

$$f \mapsto f(1)$$

induces a natural isomorphism of functors

$$\mathrm{Hom}_{A^e}(A, \_) \cong (\_)^A.$$

From the Lemma, we get a useful criterion for separability.

Corollary 2.5.12. We have that A is a separable R-algebra if and only if $(\_)^A$ is right-exact.

Proof. The algebra A is a separable over R if and only if A is projective as an A^e module if and only if $\mathrm{Hom}_{A^e}(A, \_)$ is right exact. ∎

We give some further technical lemmas from [Fo17].

Lemma 2.5.13. Let A be an R-algebra with structure map $\theta\colon R \to Z(A)$, and let I be an ideal in R with $I \subseteq \ker(\theta)$. Then A is separable over R if and only if A is separable over R/I.

Proof. By the assumptions on I, A is an R/I algebra. If e is a separability idempotent over R, then e is also a separability idempotent over R/I. Conversely, let $e \in A^e$ satisfy $Je \subseteq I$ and $\mu(e) = 1 + I$. Since $I \subseteq \ker(\theta)$, then $Je = 0$ and $\mu(e) = 1$. By Proposition 2.5.2, A is a separable R-algebra. ∎

Lemma 2.5.14. If A is a separable R-algebra and I is a two-sided ideal of A, then A/I is a separable R-algebra.

Proof. By Corollary 2.5.12, it is sufficient to check that the functor $(\_)^{A/I}$ is exact. Let M be a two-sided $(A/I)/R$-module. Then, we can consider M as a two-sided A/R-module through the projection $A \to A/I$. With this module structure, $M^A = M^{A/I}$. Since $(\_)^A$ is exact, then so is $(\_)^{A/I}$. ∎

We state the following result from [Fo17, $p.141$], and prove an important Corollary, which, as we will see, generalizes the phenomenon of non-ramification of ideals to arbitrary R-algebras. We will define ramification in Chapter 3.

Theorem 2.5.15. *Let K be a field, and let S be a commutative K-algebra. Then, S is separable over K if and only if S is isomorphic to a finite direct product of finite separable field extensions of K.*

Corollary 2.5.16. Let R and S be local rings with respective maximal ideals $\mathfrak{p}$ and $\mathfrak{P}$. If S is a separable R-algebra and the structure map $\theta\colon R \to S$ satisfies $\theta(\mathfrak{p}) \subseteq \mathfrak{P}$, then

$$\mathfrak{p}S = \mathfrak{P}$$

and $S/\mathfrak{P}$ is a finite separable field extension of $R/\mathfrak{p}$.

Proof. Since S is a commutative separable R-algebra, $\mathfrak{p}$ is a two-sided ideal, so $S/\mathfrak{p}S$ is separable over R by Lemma 2.5.14. Since the ideal $\mathfrak{p}$ of R is 0 in $S/\mathfrak{p}S$ and S is a commutative separable R-algebra, Lemma 2.5.13 tells us that $S/\mathfrak{p}S$ is separable over $R/\mathfrak{p}$. But $R/\mathfrak{p}$ is a field, so by Theorem 2.5.15, $S/\mathfrak{p}S$ is a finite product of finite separable field extensions of $R/\mathfrak{p}$. Suppose that this product of fields has n non-zero factors. A product of n non-zero rings has at least n maximal ideals, but S is local, which means that $S/\mathfrak{p}S$ is local. Therefore, we must have that $n = 1$. This means that $S/\mathfrak{p}S$ is a field, so that $\mathfrak{p}S$ is a maximal ideal of S, i.e., that $\mathfrak{p}S = \mathfrak{P}$. Moreover, $S/\mathfrak{P} = S/\mathfrak{p}S$ is a finite separable field extension of $R/\mathfrak{p}$. ∎

To finish this section, we state that tensor products of separable algebras are separable as in [Fo17, p.124], and as a corollary, that separability is preserved under base change with a commutative ring.

Theorem 2.5.17. *Let R be a commutative ring and S_1 and S_2 commutative R-algebras. Let A_1 be a separable S_1-algebra and A_2 be a separable S_2-algebra. Then $A_1 \otimes_R A_2$ is separable over $S_1 \otimes_R S_2$.*

Proof. The proof consists of showing that $(\_)^{A_1 \otimes_R A_2}$ is an exact functor by showing that certain diagrams commute. Please see [Fo17, p.124] for details. ∎

From the Theorem, we obtain that the base change of a separable R-algebra by a commutative R-algebra is separable.

Corollary 2.5.18. If A is a separable R-algebra and S is a commutative R-algebra, then

$$S \otimes_R A$$

is a separable $S \otimes_R R$-algebra.

Proof. Use $S_1 = A_1 = S$, $S_2 = R$, and $A_2 = A$ in the Theorem. ∎

2.6 The trace ideal of a separable algebra

The goal of this section is to define the trace of an algebra, and to prove that the trace map for a finitely generated, projective, and separable R-algebra S is surjective.

The results and proofs in this section are presented as in [Fo17, $p.148-152$].

Definition 2.6.1. Let M be an R-module. We call the two sided ideal

$$\mathfrak{T}_R(M) := \{\sum_{i=1}^{n} f_i(m_i) \mid f_i \in \mathrm{Hom}_R(M,R), m_i \in M\}$$

the trace ideal of M in R. We say that M is a generator of R if $\mathfrak{T}_R(M)$ is equal to R.

Definition 2.6.2. Let M be an R-module and I an indexing set. We call a collection of pairs

$$\{(m_i, f_i)\}_{i\in I}$$

with $m_i \in M$ and $f_i \in \mathrm{Hom}_R(M,R)\}$ a dual basis for M if for every $m \in M$:

1. $f_i(m) = 0$ for all but finitely many $i \in I$,
2. $m = \sum_{i\in I} f_i(m)m_i$.

Lemma 2.6.3. Let M be an R-module. Then, M is projective as an R-module if and only if M has a dual basis over R.

Proof. Suppose M has a dual basis $\{(m_i, f_i)\}_{i\in I}$, $m_i \in M$, $f_i \in \mathrm{Hom}_R(M,R)$ for all $i \in I$. Then, M is the homomorphic image of the free module $\mathcal{F}$ on the m_i. We claim the natural projection map $\varphi\colon \mathcal{F} \to M$ splits. Indeed, define $\psi\colon M \to \mathcal{F}$ by $\psi(m) = \sum_{i=1}^{m} f_i(m)m_i$. This map is R-linear because of the linearity of f_i and satisfies $\varphi \circ \psi = id_M$. Therefore, $\mathcal{F}$ is the direct sum of M and $\ker(\varphi)$. Conversely, suppose M is projective as an R-module, so that there exists a free R-module $\mathcal{F}$ and a surjection $\varphi\colon \mathcal{F} \to M$ such that φ is split by $\psi\colon M \to \mathcal{F}$. Let $\{m_i\colon i \in I\}$ be an R-basis for $\mathcal{F}$ and define $f_i\colon M \to R$ by the identity $\psi(m) = \sum_{i=1} f_i(m)m_i$ in $\mathcal{F}$. Since the m_i form a basis of $\mathcal{F}$, the expression f_i is well defined, R-linear, and $f_i(m) = 0$ for all but finitely many i. Since φ surjects, we can write every $m \in M$ as $m = \varphi(\psi(m)) = \sum_{i\in I} f_i(m)\varphi(m_i)$, so that $\{(\varphi(m_i), f_i)\}_{i\in I}$ is a dual basis for M over R. ∎

Example 2.6.4. A finitely generated free module R^n has a dual basis given by the projections onto each factor, together with the corresponding basis vector.

We need an equivalent version of Nakayama's Lemma 2.1.24.

Lemma 2.6.5. Let R be a commutative ring, I an ideal of R, and M a finitely generated R-module. We have that $IM = M$ if and only if $I + Ann_R(M) = R$.

Proof. Let $s \in R$, and let us show that s is an element of $I + Ann_R(M)$. By Theorem 2.1.24, there exists $r \in R$ such that $r \equiv 1 \pmod{I}$ and $rM = 0$. Then, there exists $i \in I$ such that $r - i = 1$ and therefore that $rs - is = s$, i.e., that $s \in I + Ann_R(M)$. Conversely, if $I + Ann_R(M) = R$, then $IM = (I + Ann_R)M = RM$. ∎

Lemma 2.6.6. Let R be a commutative ring and I, J be ideals of R that $I + J = R$. Then $I \cap J = IJ$.

Proof. We always have $IJ \subseteq I \cap J$. If $1 = i + j$, $i \in I, j \in J$, then $a \in I \cap J$ can be written as $a = a \cdot 1 = a(i + j) = ai + aj = ia + aj \in IJ$. ∎

Nakayama's Lemma then helps us show that R is essentially generated by the trace ideal, up to annihilators, as in [Fo17, p.10].

Proposition 2.6.7. Let R be a commutative ring and M a finitely generated projective R-module. Then

$$R = \mathfrak{T}_R(M) \oplus Ann_R(M)$$

Proof. Since M is finitely generated projective, the construction of Lemma 2.6.3 affords us a finite dual basis $(f_i, m_i)_{1 \leq i \leq n}$. So, for any $m \in M$,

$$m = \sum_{i=1}^{n} f_i(m) m_i.$$

This shows that $\mathfrak{T}_R(M)M = M$, so that by Nakayama's Lemma 2.6.5,

$$R = \mathfrak{T}_R(M) + Ann_R(M).$$

By Lemma 2.6.6, it suffices to show $\mathfrak{T}_R(M) Ann_R(M) = 0$. But for $f \in \mathrm{Hom}_R(M, R)$, $r \in Ann_R(m)$, $m \in M$, we have that $rf(m) = f(rm) = 0$. Since,$\mathfrak{T}_R(M) Ann_R(M)$ is generated by elements of the form $rf(m)$, $\mathfrak{T}_R(M) Ann_R(M) = 0$. ∎

We immediately get the following from the Proposition and the definition of a faithful module.

Corollary 2.6.8. If M is finitely generated, projective, and faithful as an R-module, then $R = \mathfrak{T}_R(M)$.

We want to generalize the trace map of a separable field extension to finitely generated, projective R-algebras.

Definition 2.6.9. Let A be an R-algebra, and let M be an A-module that is finitely generated and projective as an R-module. By Lemma 2.6.3, we have a dual basis $f_1, \ldots, f_k \in \mathrm{Hom}_R(M, R)$ and $m_1, \ldots, m_k$ for M. The trace map from A to R with respect to M is the map

$$T_R^{A,M} \colon A \to R$$
$$x \mapsto \sum_{i=1}^{k} f_i(x \cdot m_i).$$

When $A = M$, we simply write T_R^A.

We show that the trace map is well-defined, independently of the choice of dual basis for M over R.

Lemma 2.6.10. Let A be an R-algebra and let M be an A-module that is finitely generated projective as an R-module. Suppose that $\{(f_i, m_i)\}_{1 \le i \le k}$ and $\{(g_j, n_j)\}_{1 \le i \le \ell}$ are dual bases for M over R. Then,

$$\sum_{i=1}^{m} f_i(x \cdot x_i) = \sum_{j=1}^{\ell} g_j(x \cdot n_j)$$

for all $x \in A$.

Proof. Since the f_i's and m_i's are a dual basis for M over R, we can write every n_j, $1 \le j \le \ell$, as $\sum_{i=1}^{k} f_i(n_j) m_i$. Since the f_i and g_j are R-linear, we have

$$\begin{aligned}
\sum_{j=1}^{\ell} g_j(x \cdot n_j) &= \sum_{j=1}^{\ell} g_j(x \cdot \sum_{i=1}^{k} f_i(n_j) m_i) \\
&= \sum_{i,j} f_i(n_j) g_j(x \cdot m_i) \\
&= \sum_{i,j} f_i(g_j(x \cdot m_i) n_j)
\end{aligned}$$

$$= \sum_{i=1}^{k} f_i(\sum_{j=1}^{\ell} g_j(x \cdot m_i)n_j)$$
$$= \sum_{i=1}^{k} f_i(x \cdot m_i)$$

for all $x \in A$. So, the trace map is independent of the choice of dual basis. ∎

Example 2.6.11. Let E/K be a field extension of degree n. Then, the usual field trace $\mathrm{Tr_{E/K}}$ is equal to the algebra trace map T_K^E of Definition 2.6.9. Let $\{e_1, \ldots, e_n\}$ be a basis for E/K. Then, if $f_i \in \mathrm{Hom}_K(E, K)$ is the projection onto e_i, we have that the f_i and e_i form a dual basis for E over K. Now, the field trace of $\beta \in E$ is defined as the trace of the representing matrix for β in the left regular representation of E with respect to our chosen basis. This means that $\mathrm{Tr_{E/K}}(\beta) = \sum_{i=1}^{n} a_{ii}$, where a_{ii} is the projection onto the i^{th} basis vector e_i of the field element βe_i. So,

$$\mathrm{Tr_{E/K}}(\beta) = \sum_{i=1}^{n} f_i(\beta e_i).$$

By Lemma 2.6.10, this shows the algebra trace map is equal to the field trace map. In particular, if E/K is Galois with group G, the algebra trace map satisfies

$$T_K^E(\beta) = \sum_{\sigma \in G} \sigma(\beta).$$

The following exercise from [Fo17] asserts that the trace map behaves nicely with respect to direct sums of modules.

Lemma 2.6.12. Let A be an R-algebra, M a left A-module which is finitely generated, projective over R. If $M = M_1 \oplus M_2$ as A-modules, we have

$$T_R^{A,M} = T_R^{A,M_1} + T_R^{A,M_2}.$$

Proof. Let $m_1, \ldots, m_k \in M$, $f_1, \ldots, f_k \in \mathrm{Hom}_R(M, R)$ be a dual basis for M over R. If $m_i = m_i^1 + m_i^2$ for $m_i^j \in M_j$, $1 \leq j \leq 2$, we have that the $\{f_i|_{M_j}\}_{1 \leq i \leq k}$ and $M_1^j, \ldots, M_k^j$ form a dual basis for each M_j by basic properties of the direct sum of modules. By Lemma 2.6.10, the trace relative to M is defined by any choice of dual basis for M. The identity of the Lemma then follows from a quick computation using the linearity of the f_i. ∎

We obtain our main surjectivity theorem for the trace map of an algebra, in the case that this algebra is separable, finitely generated, and projective over the base ring. The following result and its corollaries are a reproduction of those in [Fo17, p.150].

Theorem 2.6.13. *Let R be a commutative ring, and let S be a commutative R-algebra such that R is a subring of S. Then S is finitely generated as an R-module, projective, and separable over R if and only if there exists an element $T \in \mathrm{Hom}_R(S, R)$ and elements $x_1, \ldots, x_n, y_1, \ldots, y_n \in S$ satisfying*

1. $\sum_{j=1}^{n} x_j y_j = 1$ *and*
2. $\sum_{j=1}^{n} x_j T(y_j x) = x$ *for all* $x \in S$.

The map T is always equal to the trace map T_R^S.

Proof. We will only give the proof of the forward direction. The proof of the converse may be found in [Fo17, p.150]. Let $a_1, \ldots, a_m \in S$ and $f_1, \ldots, f_m \in \mathrm{Hom}_R(S, R)$ be a dual basis for S over R. Note that since S is finitely genereated and projective over R, $S \otimes S$ is finitely generated and projective over $S \otimes R$ by Lemma 2.1.12. In particular, a dual basis for $S \otimes S$ is given by the $1 \otimes a_i$ and $1 \otimes f_i$, so that the trace map $T_{S\otimes R}^{S\otimes S}$ is equal to $1 \otimes T_R^S$. Since S is separable over R, Corollary 2.5.18 tells us that $S \otimes S$ is separable over $S \otimes R$.

Let e be a separability idempotent for $S \otimes R$ over R. Since the multiplication map μ splits the sequence of Proposition 2.5.2, we have that μ induces an isomorphism $(S \otimes R)\, e \cong S \otimes R$, and have a direct sum $S \otimes S \cong J \oplus (S \otimes R)\, e \cong J \oplus S \otimes R$. Lemma 2.6.12 then says that

$$T_{S\otimes R}^{S\otimes S} = T_{S\otimes R}^{S\otimes S, J} + T_{S\otimes R}^{S\otimes S, S\otimes R} = T_{S\otimes R}^{J} \oplus T_{S\otimes R}^{S\otimes R},$$

where $T_{S\otimes R}^{J}$ is the restriction of the trace on all $S \otimes S$ to J and $T_{S\otimes R}^{S\otimes R}$ is the restriction of the trace to $(S \otimes R)\, e$. In particular,

$$T_{S\otimes R}^{S\otimes S}\left((x \otimes 1)e\right) = T_{S\otimes R}^{S\otimes R}\left((x \otimes 1)e\right).$$

Let $f\colon (S \otimes R)e \to S \otimes R$ be the isomorphism induced by the multiplication map μ. We have that $e \in S \otimes R$ and the isomorphism f form a dual basis for $S \otimes R$ as an

$S \otimes R$ module, so that

$$T_{S\otimes S}^{S\otimes S}((x\otimes 1) = T_{S\otimes R}^{S\otimes R}(x\otimes 1) = f((x\otimes 1)e) = x\otimes 1.$$

Let $x_1, \ldots, x_n, y_1, \ldots, y_n \in S$ be such that $e = \sum_j x_j \otimes y_j$. Then, $\sum_j x_j y_j = \mu(e) = 1$, which proves point 1. For any $x \in S$, we have by our computation of the trace map that

$$\begin{aligned}\sum_j x_j \otimes T_R^S(y_j x) &= 1 \otimes T_R^S\left(\sum_j x_j \otimes y_j x\right)\\ &= 1 \otimes T_R^S\left((1\otimes x)e\right)\\ &= T_{S\otimes R}^{S\otimes S}\left((1\otimes x)\, e\right)\\ &= x\otimes 1.\end{aligned}$$

Applying the multiplication homomorphism μ allows us to conclude that the second point of the statement holds. Noting that $\{x_1, \ldots, x_n\}$, $\{T(y_1\cdot), \ldots, T(y_n\cdot)\}$ is a dual basis for S over R, an easy computation and use of Lemma 2.6.10 shows that $T = T_R^S$ if T satisfies point 2. ∎

Since S/R is an extension of commutative rings, the set $\mathrm{Hom}_R(S,R)$ has the structure of a right S-module, with the S-action defined by $(f\cdot\alpha)(x) = f(\alpha x)$ for all $x, \alpha \in S$, In other words, $f\cdot\alpha = f\circ\ell_\alpha$, where ℓ_α is the image of α in the left regular representation $S \to \mathrm{Hom}_R(S,S)$.

Corollary 2.6.14. Let S/R be a separable extension of commutative rings such that S is a finitely generated projective R-module. The singleton $\{T_R^S\}$ is an S-basis for $\mathrm{Hom}_R(S,R)$.

Proof. Let $x_1, \ldots, x_n, y_1, \ldots, y_n \in S$ be as in the Theorem. For any $f \in \mathrm{Hom}_R(S,R)$, we have

$$\begin{aligned}f(x) &= \sum_{j=1}^n f\left(x_j T_R^S(y_j x)\right)\\ &= \sum_{j=1}^n T_R^S(y_j x)\, f(x_j)\\ &= T_R^S\left(\sum_{j=1}^n y_j x f(x_j)\right)\end{aligned}$$

$$= T_R^S(x) \cdot \left(\sum_{j=1}^{n} y_j f(x_j) \right),$$

so that the trace map generates $\mathrm{Hom}_R(S,R)$ as a right S-module. To show that $\{T_R^S\}$ forms a basis of $\mathrm{Hom}_R(S,R)$, we show that the annihilator of T_R^S in S is $\{0\}$. If $0 = \left(T_R^S \cdot \alpha\right)(x)$ for all $x \in S$, we have by the second point of the Theorem that

$$\alpha = \sum_{j=1}^{n} x_n T_R^S(y_j \alpha)) = \sum_{j=1}^{n} x_j \left(T_R^S \cdot \alpha\right)(y_j) = 0.$$

So, the annihilator of the trace map in S is 0. ∎

Corollary 2.6.15. If S is a finitely generated, projective, and separable extension of R, then there is an element $c \in S$ such that $T_R^S(c) = 1$. Moreover, $R \cdot c$ is an R-module direct summand of S.

Proof. Since R is projective as an R-module, it suffices to show that there exists $c \in S$ such that $T_R^S(c) = 1$ to obtain that c is an R-direct summand of S. Since S is finitely generated and projective as an R-module, we have a dual basis $f_1, \ldots, f_n \in \mathrm{Hom}_R(S,R), y_1, \ldots, y_n \in S$, with which we can write $1 = \sum_{i=1}^{n} f_j(x_j)$. By the Corollary, there exists $\alpha_j \in S$ such that $f_j(x_j) = T_R^S(\alpha_j x_j)$ for each $j \in \{1, \ldots, n\}$. Therefore,

$$1 = T_R^S \left(\sum_{j=1}^{n} \alpha_j x_j \right),$$

and $c = \sum_{j=1}^{n} \alpha_j x_j$ is the required element of S. ∎

Chapter 3

Dedekind Domains

In Chapter 3, we introduce Dedekind domains. We are mainly interested in their local rings, the discrete valuation rings, which will be of interest when we define Hermitian lattices in Chapter 4. As we will see, Dedekind domains satisfy many nice properties, and for this reason, it is worth giving an exposition of their theory. In Section 3.1, we define Dedekind domains and prove that they satisfy unique factorization of ideals. Section 3.2 gives some technical results on localization in the context of Dedekind domains. Section 3.3 serves to present a result outlining the factorization of a prime ideal of a Dedekind domain in maximal integral extensions. Section 3.4 introduces non-Archimedean valuations and defines discrete valuation rings. In Sections 3.5 and 3.6, we give interesting examples of discrete valuation rings. From arithmetic, we present the p-adic numbers, which arise as the rings of integers of the completion of number fields at primes, and from geometry, we describe functions on algebraic curves.

3.1 Fractional ideals and Dedekind domains

We present the basic theory of Dedekind domains. In this section, we follow [Ash02, §7.6-7.8], and unless stated otherwise, the proofs we give are minor modifications of those presented in the reference. The product of two ideals I, J of a commutative

ring is defined to be

$$IJ = \{\sum_{i=1}^{m} a_i b_i \mid a_i \in I, b_i \in J\}.$$

The purpose of this section will be to define an inverse ideal with respect to the product operation, the fractional ideal. We will then introduce a special class of rings, the so-called Dedekind domains, which admit inversion of ideals, and equivalently, unique factorization of ideals. We fix R to be an integral domain with fraction field K.

Lemma 3.1.1. If $\mathfrak{p}$ is a prime ideal of R that contains a product $I_1 \dots I_r$ of ideals, then $\mathfrak{p}$ contains some I_j.

Proof. Suppose that every ideal I_j in the product is not entirely contained in $\mathfrak{p}$. Then, there exist a_j in each I_j such that a_j is not an element of $\mathfrak{p}$. But $\prod_{j=1}^{r} a_j$ is contained in $\mathfrak{p}$. Since $\mathfrak{p}$ is prime, one of the a_j is contained in $\mathfrak{p}$. ∎

Lemma 3.1.2. If I is a nonzero ideal of the Noetherian integral domain R, then I contains a product of nonzero prime ideals.

Proof. Let I be a maximal counterexample, which exists since R is Noetherian. In particular, I is not prime. So, there exist $a, b \in R$ such that $ab \in I$ but neither a nor b are in I. By the maximality of I, the ideals $I+aR$ and $I+bR$ each contain a product of non-zero primes $\mathfrak{p}_1 \dots \mathfrak{p}_r$ and $\mathfrak{q}_1 \dots \mathfrak{q}_s$. Since R is an integral domain, $\mathfrak{p}_1 \dots \mathfrak{p}_r \mathfrak{q}_1 \dots \mathfrak{q}_s$ is a non-zero product of primes contained in the ideal $I + abR \subseteq I$. ∎

Definition 3.1.3. Let R be an integral domain with field of fractions K. A fractional ideal is an R-submodule I of K such that there exists a non-zero $r \in R$ with $rI \subseteq R$. We say r is a denominator of I. The ideals of R are fractional ideals and are called integral ideals.

Lemma 3.1.4. If I is a finitely generated R-submodule of K, then I is a fractional ideal. If R is Noetherian and I is a fractional ideal of R, then I is finitely generated.

Proof. Let $\frac{a_i}{b_i}$, $1 \leq i \leq r$, be finitely many generators of $I \subseteq K$. Then, $\prod_{i=1}^{r} b_i$ is a denominator for I. Suppose R is Noetherian and that $rI \subseteq R$. Then, $I \subseteq r^{-1}R$. Seeing as $r^{-1}R \cong R$ as R-modules and R is Noetherian, I is a submodule of the Noetherian R-module $r^{-1}R$. By Proposition 2.2.3, I is finitely generated as an R-module. ∎

Definition 3.1.5. Let I and J be fractional ideals of R. We define the product of fractional ideals as

$$IJ := \{\sum_{i=1}^{r} a_i b_i \mid a_i \in I, b_i \in J\}.$$

The sum of fractional ideals $I + J$ is their sum as R-submodules of K.

Lemma 3.1.6. If I and J are fractional ideals of R, then $I \cap J$, IJ, and $I + J$ are fractional ideals of R.

Proof. If r_i is a denominator for I and r_j is a denominator for J, then r_i is a denominator for $I \cap J$ and IJ, and $r_i r_j$ is a denominator for $I + J$. ∎

The Lemma shows that the set of fractional ideals of an integral domain is closed under intersection, product, and sum. Moreover, the product of fractional ideals is an associative operation, and taking the product of a fractional ideal I with the ring R yields I. Therefore, we can define a commutative monoid structure on the fractional ideals of an integral domain with operation the product of ideals and with neutral element R. We will show that, in a certain class of rings, every non-zero element of this monoid is invertible under multiplication.

Definition 3.1.7. Let R be a ring, and let $\mathfrak{p}$ be a prime ideal of R. The height of $\mathfrak{p}$ is the supremum

$$ht(\mathfrak{p}) := \sup\{n \mid \mathfrak{p}_0 \subsetneq \mathfrak{p}_1 \subsetneq \cdots \subsetneq \mathfrak{p}_{n-1} \subsetneq \mathfrak{p}_n = \mathfrak{p}\}$$

of the length of chains of prime ideals contained in $\mathfrak{p}$. The Krull dimension of R is the supremum

$$\sup_{\mathfrak{p} \in \mathrm{Spec}(R)} ht(\mathfrak{p}).$$

An integral domain of dimension 0 is a field as 0 must be the only prime ideal. In an integral domain of dimension 1, the zero ideal is prime and there are non-trivial prime ideals. Since $0 \subsetneq \mathfrak{p}$ is a maximal chain, every non-zero prime is maximal. Since the Krull dimension is defined as a supremum of length of chains of prime ideals, it is not sufficient to assume that R is Noetherian to conclude that R has finite Krull dimension. Nagata provided an example of a Noetherian ring of infinite Krull dimension in [Nag62]. We now define Dedekind domains.

Definition 3.1.8. A Dedekind domain is an integrally closed Noetherian integral domain of Krull dimension 1.

Example 3.1.9. Examples of Dedekind domains abound. All principal ideal domains are Dedekind domains, as their prime ideals are maximal, they are Noetherian, and, by Lemma 2.3.8, the unique factorization property implies integral closure. In particular, the integers are a Dedekind domain. However, there are Dedekind domains which are not unique factorization domains. For example, one might show in an undergraduate ring theory course that $\mathbb{Z}[\sqrt{5}]$ is not a unique factorization domain by showing that $2 \cdot 3$ and $(1+\sqrt{-5}) \cdot \left(1-\sqrt{-5}\right)$ are two prime factorizations of 6, as in [Neu99, p.17]. A theorem from classical number theory ([Ash02, §7.2]) says that the ring of integers of $\mathbb{Q}(\sqrt{-5})$ is equal to $\mathbb{Z}[\sqrt{-5}]$. As we will see, this means that $\mathbb{Z}[\sqrt{-5}]$ is a Dedekind domain.

So, it is not true in general that integral extensions of $\mathbb{Z}$, i.e., rings whose elements are solutions to integer polynomial equations, satisfy the unique factorization property. Historically, the discovery that not all rings of integers satisfy unique factorization caused an uproar. Algebra was in its infancy: the French mathematician Gabriel Lamé had even given a proof of Fermat's last theorem, but his proof used in an essential way that rings of integers corresponding to roots of unity satisfy unique factorization. Kummer had done the same. He believed that algebraic integers satisfied a property like unique factorization, and theorized that unique factorization holds over some universal integral extension, the so-called ideal numbers. Given two factorizations for an element, say $2 \cdot 3 = 6 = (1+\sqrt{-5}) \cdot (1-\sqrt{-5})$, then there should be uniquely defined prime "elements", or ideal numbers, $\mathfrak{p}_1, \mathfrak{p}_2, \mathfrak{p}_3, \mathfrak{p}_4$ living in some extension of $\mathbb{Z}[\sqrt{-5}]$ such that $6 = \mathfrak{p}_1\mathfrak{p}_2\mathfrak{p}_3\mathfrak{p}_4$. As we now know, this is not the case, but the divisibility properties of these theorized ideal numbers implied that one should look at what we call ideals. To this end, Richard Dedekind was the one who eventually gave the correct definition of an ideal in an integral extension of $\mathbb{Z}$, and these ideals gave the best sense of unique factorization: the ideal generated by an element could be uniquely factorized. The historical background is further developed in [Neu99, p.17,p.38]. So, in what follows, we will prove that the rings of integers of number fields, and more generally, Dedekind domains, satisfy unique factorization of ideals. First, we reproduce the following technical result of [Ash02, §7.6].

Lemma 3.1.10. Let $\mathfrak{p}$ be a prime ideal of the Dedekind domain A. Set $J = \{x \in K \mid x\mathfrak{p} \subseteq A\}$. Then A is properly contained in J.

Proof. If $\mathfrak{p} = 0$, then $J = K$. Let us assume that $\mathfrak{p} \neq 0$. Since $\mathfrak{p}$ is an integral ideal, we must have that $A \subseteq J$. Let $a \in \mathfrak{p}$ be non-zero, and consider the ideal $aA \subseteq \mathfrak{p}$. Since A is Noetherian, we apply Lemma 3.1.2 and obtain a non-zero product of prime ideals $\mathfrak{p}_1 \dots \mathfrak{p}_r$ contained in $aA \subseteq \mathfrak{p}$. We can assume that r is minimal. By Lemma 3.1.1, we have that some $\mathfrak{p}_i$ is contained in $\mathfrak{p}$. Without loss of generality, $i = 1$. Since A has dimension 1 and $\mathfrak{p}_1$ is non-zero, $\mathfrak{p}_1 = \mathfrak{p}$. Suppose first that $r \geq 2$. By minimality of r, aA is not contained in $\mathfrak{p}_2 \dots \mathfrak{p}_r$. We can therefore choose $b \in \mathfrak{p}_2 \dots \mathfrak{p}_r$ such that $b \in A \backslash aA$. Now, $b\mathfrak{p} \subseteq aA$, so that $\mathfrak{p}ba^{-1}\mathfrak{p} \subseteq A$. Therefore, $ba^{-1} \in J$. On the other hand, if ba^{-1} were an element of A, then b would be an element of aA. So, $ba^{-1} \in J \backslash A$. If the minimal r above is 1, we get that $\mathfrak{p}$ is the principal ideal aA by the containment $\mathfrak{p} \subseteq aA \subseteq \mathfrak{p}$ and maximality of $\mathfrak{p}$. Then, we can choose $b \in A \backslash aA$ and obtain that ba^{-1} is not in A, as otherwise $b \in aA$. Then, $ba^{-1}\mathfrak{p} = ba^{-1}aA = bA \subseteq A$, so that $ba^{-1} \in J$. ∎

The next two results show that in a Dedekind domain, prime ideals have fractional inverses under the product operation. We follow the proof given in [Ash02, §7.6].

Lemma 3.1.11. Let I be a non-zero ideal of the integrally closed Noetherian integral domain R, and suppose J is a fractional ideal of R such that $JI = I$. Then $J \subseteq R$.

Proof. We show that J is integral over R. Let $x \in J$. Then, $x^n I \subseteq I$ for any $n \in N$, and thus $rx^n I \subseteq I$ for any $r \in R$. So, for any $a \in I$, we have that $rx^n a \in R$, for all $r \in R$ and all $n \in \mathbb{N}$. So, $aR[x] \subseteq R$. This shows that $R[x]$ is a fractional ideal of the Noetherian integral domain R, so that by Lemma 3.1.4, $R[x]$ is a finitely generated R-module. By Lemma 2.3.2, x is integral over R. Since R is integrally closed, $x \in R$. So, $J \subseteq R$. ∎

We give an explicit construction for the inverse of an ideal, before moving on to our main theorem.

Proposition 3.1.12. Let $\mathfrak{p}$ be a nonzero prime ideal of the Dedekind domain A, and let $J = \{x \in K \mid x\mathfrak{p} \subseteq A\}$. Then J is a fractional ideal of A and $\mathfrak{p}J = A$.

Proof. By construction, J is an A-submodule of K. Now, any $r \in \mathfrak{p}$ satisfies $rJ \subseteq A$, i.e., r is a denominator for J. Therefore, J is a fractional ideal of A. To show $\mathfrak{p}J = A$,

suppose that $\mathfrak{p}J$ is properly contained in A. By Lemma 3.1.10, we have that $A \subsetneq J$, so $\mathfrak{p} = \mathfrak{p}A \subseteq \mathfrak{p}J$. By assumption, $\mathfrak{p}J \subseteq A$, so that by maximality of the prime $\mathfrak{p}$, $\mathfrak{p} = \mathfrak{p}J$ or $\mathfrak{p}J = A$. Suppose then that $\mathfrak{p} = \mathfrak{p}J$. By Lemma 3.1.11 we have that that $J \subseteq A$. But we know that $A \subsetneq J$, by Lemma 3.1.10. So, $\mathfrak{p}$ is properly contained in $\mathfrak{p}J$, hence $\mathfrak{p}J = A$. ∎

We show that fractional ideals of a Dedekind domain factor uniquely into a product of primes, following the proof of [Ash02, §7.6]

Theorem 3.1.13. *If I is a nonzero fractional ideal of the Dedekind domain A, then I can be factored uniquely as*

$$\mathfrak{p}_1^{n_1} \dots \mathfrak{p}_r^{n_r}$$

where the $\mathfrak{p}_i$ are prime ideals and the n_i are integers. In other words, the nonzero fractional ideals form a group under multiplication.

Proof. We will first look at integral ideals. Since A is Noetherian, we can consider a maximal counterexample I_0. Then, since I_0 is not prime, I_0 is a properly contained in a maximal ideal I. By maximality of I_0, I admits a prime factorization. By Proposition 3.1.12, there exists a fractional ideal J that is the inverse of I, and since I is prime, Lemma 3.1.10 says that J properly contains A. Now, $II_0 \subseteq I_0$, so $I_0 = IJI_0 \subseteq JI_0$. But Lemma 3.1.11 says that since J properly contains A, J cannot satisfy $JI_0 = I_0$. So, $I_0 \subsetneq JI_0$. Moreover, we have that $JI_0 \subseteq JI = A$. So, JI_0 is an integral ideal of A that properly contains I_0, so that JI_0 has a prime factorization. Since I has a prime factorization, so does $IJI_0 = I_0$. If J is an arbitrary fractional ideal and r is a denominator for J, then $J = (rA)^{-1}(rA)J$. We know that rAJ is integral and $(rA)^{-1}$ has prime factorization obtained by inverting the prime factors of rA, so J has a prime factorization. To show uniqueness, we proceed by induction on the number of prime factors of the ideal I with multiplicity. Let $I = \mathfrak{p}_1^{n_1} \dots \mathfrak{p}_k^{n_k} = \mathfrak{q}_1^{m_1} \dots \mathfrak{q}_\ell^{m_\ell}$ for some primes $\mathfrak{p}_i, \mathfrak{q}_j$ of A. Without loss of generality, we can assume that all exponents are non-negative, as associativity of the ideal product allows us to clear denominators. The right-hand side is contained in $\mathfrak{p}_1$, so that by Lemma 3.1.1, $\mathfrak{p}_1$ contains $\mathfrak{q}_k$ for some k. By maximality, $\mathfrak{q}_k = \mathfrak{p}_1$, and we can cancel these prime ideals to apply induction. ∎

We can make some immediate observations following the theorem. By unique

factorization, the integral ideals are those with only non-negative exponents in their prime factorization. An ideal with prime factorization $\mathfrak{p}_1^{e_1} \dots \mathfrak{p}_k^{e_k}$ is contained in $\mathfrak{p}_1^{f_1} \dots \mathfrak{p}_k^{f_k}$ if and only if $e_i \geq f_k$ for every $1 \leq i \leq k$. We then say that I divides J if J is contained in I, and this definition allows us to make sense of greatest common divisors and least common multiples in the group of ideals of a Dedekind domain. It is not hard to show that the following definitions of GCD and LCM agree with the one coming from the divisibility relation.

Definition 3.1.14. Let I and J be fractional ideals of the Dedekind domain A. The greatest common divisor of I and J is the sum $I + J$. The least common multiple of I and J is the intersection $I \cap J$.

Remark 3.1.15. In passing, it is also nice to note that the existence and uniqueness of prime factorizations is a sufficient condition for an integral domain A to be Dedekind. Writing an ideal in its (finite, unique) prime factorization shows that chains of ideals are bounded above. Since to contain means to divide, prime ideals, with a prime factorization of length 1, are necessarily maximal. A brief argument utilizing the prime factorization of principal fractional ideals then shows that A is integrally closed.

We need one last result on integral extensions from [AM69, p.61], before proving that Dedekind domains behave well under extension of their fields of fractions.

Lemma 3.1.16. Let A and B be integral domains such that B is integral over A. We have that A is a field if and only if B is a field.

Proof. Suppose A is a field and let $b \in B$ be non-zero. Since b is integral over A, Lemma 2.3.2 says that the ring $A[b]$ is a finitely generated A-vector space. Consider the left multiplication map $\ell_b \colon B \to B$ induced by b. It suffices to show ℓ_b is an invertible map of vector spaces $A[b] \to A[b]$. Seeing as B is an integral domain, injectivity is immediate. Now, if $x \in A[b]$, we have that $x = p(b)$, for $p(X) \in A[X]$. So, $bx = bp(b)$ is in $A[b]$. Since the linear map $\ell_b \colon A[b] \to A[b]$ is injective, we must have that ℓ_b is surjective. This shows b^{-1} exists. Conversely, let B be a field. If $x \in A$, then $x^{-1} \in B$. Now, x^{-1} is integral over a, and therefore satisfies a minimal equation of integral dependence

$$x^{-n} + a_1 x^{-(n-1)} + \dots + a_n = 0.$$

Multiplying through by x^{n-1} and isolating, we have $x^{-1} = -(a_1 + a_2 x + \cdots + a_n x^{n-1})$, so x^{-1} is an element of A. ∎

Corollary 3.1.17. Let A and B be integral domains such that B is integral over A, let $\mathfrak{P}$ be a prime ideal of B, and set $\mathfrak{p} = \mathfrak{P} \cap A$. Then $\mathfrak{P}$ is maximal if and only if $\mathfrak{p}$ is maximal.

Proof. Since $\mathfrak{p} \subseteq \mathfrak{P}$ and $\mathfrak{p}$ is a prime ideal, the inclusion $A \subseteq B$ descends to an inclusion $A/\mathfrak{p} \subseteq B/\mathfrak{P}$ of integral domains. Since we can reduce any equation of integral dependence over A to obtain one over $A/\mathfrak{p}$, we get that $B/\mathfrak{P}$ is integral over $A/\mathfrak{p}$. By the Lemma, $B/\mathfrak{P}$ is a field if and only if $A/\mathfrak{p}$ is a field. ∎

By specializing the following result to finite extensions of a number field, we will obtain that the ring of integers of a number field is a Dedekind domain.

Theorem 3.1.18. *Let A be a Dedekind domain with field of fractions K, and let L be a finite separable extension of K. Let B be the integral closure of A in L. Then, B is a Dedekind domain with field of fractions L.*

Proof. We first check that B is an integrally closed and Noetherian ring of Krull dimension 1. By Corollary 2.4.17, we have that B is Noetherian. By the transitivity of integral extensions of Proposition 2.3.4, the integral closure of B in B is integral over A, and is therefore contained in B. This means that B is integrally closed. It remains to show that primes of B are maximal. If $\mathfrak{P}$ is a prime of B, we can consider the prime $\mathfrak{p} = \mathfrak{P} \cap A$ of A. It suffices to show that $\mathfrak{p}$ is non-zero. Indeed, if $\mathfrak{p}$ is non-zero, then $\mathfrak{p}$ is maximal, as a non-zero prime ideal in a ring of dimension 1. We can then apply Corollary 3.1.17 to show $\mathfrak{P}$ is also maximal. To show $\mathfrak{P} \cap A$ is not zero, let $x \in \mathfrak{P}$. We have that x satisfies a minimal equation of integral dependence

$$x^n + a_1 x^{n-1} + \cdots + a_n = 0.$$

In particular, $a_n \neq 0$. Isolating a_n, we see that x divides a_n in B. This means that $a_n \in \mathfrak{P} \cap A = \mathfrak{p}$, so that $\mathfrak{p}$ is non-zero. Since B is the integral closure of A in L, Lemma 2.4.12 says that L is the field of fractions of L. ∎

Corollary 3.1.19. Let A, K, L, B be as in the Theorem and let $\mathfrak{P}$ be a non-zero prime ideal of B. Then $A \cap \mathfrak{P}$ is a non-zero prime ideal of A.

Proof. Let $\mathfrak{P}$ be a non-zero prime ideal of B. Since B is a Dedekind domain, $\mathfrak{P}$ of B must be maximal. So, by Corollary 3.1.17, $\mathfrak{P} \cap A$ is a maximal ideal of A. Since A is an integral domain of Krull dimension 1, $\mathfrak{p}$ properly contains the zero ideal. ∎

Following [Ash02, §7.8], we state some basic properties of ideals in Dedekind domains. The Lemma is technical, and the omitted proof can be found in the reference.

Lemma 3.1.20. Let R be a Dedekind domain, and let I be a non-zero ideal of R. There exists a non-zero ideal I' such that II' is a principal ideal. Moreover, if J is any ideal not equal to I, I' can be chosen coprime to J.

Corollary 3.1.21. A Dedekind domain with only finitely many prime ideals is a principal ideal domain.

Proof. Write $J = \mathfrak{p}_1 \dots \mathfrak{p}_n$, the product of all the prime ideals of R, and let I be an arbitrary ideal. By Lemma 3.1.20, there exists a nonzero ideal I' such that II' is principal and that I' is coprime to J. Then I' is coprime to every prime ideal of R, which implies that $I' = R$. This means that I was already principal. ∎

Corollary 3.1.22. Let I be an ideal of the Dedekind domain R. If $a \in I$, then I is generated by at most 2 elements, one of which is a.

Proof. Let $a \in I$. Since the ideal generated by a is contained in I, we have $IJ = \langle a \rangle$, for some integral ideal J of R. By Lemma 3.1.20, there exists an ideal I' coprime to J such that $II' = \langle b \rangle$, for some $b \in I$. But the ideal generated by a and b is $IJ + II' = gcd(IJ, II') = I$. ∎

3.2 Localization of Dedekind domains

This section serves to introduce technical facts. We first remind the reader of the following fact about localizations of rings, which describes the ideals of a localization $S^{-1}A$ as those which are generated by those ideals of A disjoint from S. This correspondence will then enable us to describe the properties of localized Dedekind domains.

Theorem 3.2.1. *Let A be a ring and let S be a multiplicative set in A. The assignments*

$$I \mapsto S^{-1}I$$

induces a one-to-one inclusion-preserving correspondence between ideals of $S^{-1}A$ and ideals of A disjoint from S. Under this correspondence, prime ideals correspond to prime ideals. In particular, every ideal of $S^{-1}A$ is generated by the image of some ideal I of A in $S^{-1}A$.

Proof. See, for instance, [AM69]. ∎

The rest of the section consists of a collection of technical results describing the properties of localizations of certain rings.

Corollary 3.2.2. Let R be a ring, and let S be a multiplicative set in A. Then

$$\dim(S^{-1}R) \leq \dim(R).$$

Proof. By the Theorem, the length of a chain of prime ideals in $S^{-1}R$ is bounded above by the length of a chain of prime ideals in R. ∎

Lemma 3.2.3. Let A be a Noetherian ring, and let S be a multiplicative set in A. Then $S^{-1}A$ is Noetherian.

Proof. By Theorem 3.2.1, ideals of $S^{-1}A$ are in one-to-one correspondence with ideals of A disjoint from S, so that ascending chains of ideals in $S^{-1}A$ stabilize. ∎

Lemma 3.2.4. Let $A \subseteq B$ be rings, with B integral over A. If S is a multiplicative set in A, then $S^{-1}B$ is integral over $S^{-1}A$.

Proof. Let $\frac{\beta}{s} \in S^{-1}B$, for some $\beta \in B$ and $s \in S$. Since B is integral over A, there exists a monic $A(X) \in A[X]$ such that

$$A(\beta) = \sum_{k=0}^{n} a_k \beta^k = 0.$$

Then,

$$0 = \frac{1}{s^n} A(\beta) = (\frac{\beta}{s})^n + \sum_{k=0}^{n-1} \frac{a_k}{s^{n-i}} (\frac{\beta}{s^i})^k$$

is an equation of integral dependence over $S^{-1}A$. ∎

Lemma 3.2.5. Let B be an integral domain, integrally closed in its field of fractions L. Let S be a multiplicative set in B. Then $S^{-1}B$ is integrally closed in L.

Proof. Suppose $\alpha \in L$ satisfies an equation of the form

$$A(\alpha) = \sum_{k=0}^{n} \frac{b_k}{s_k}\alpha^k = 0,$$

$b_k \in B, s_k \in S$. Writing $s := \prod_{k=0}^{n} s_k$, we have that

$$0 = s^n A(\beta) = s^n\alpha^n + \sum_{k=0}^{n-1} \frac{s^n a_k}{s_k}\alpha^k = (s\alpha)^n + \sum_{k=0}^{n-1} \frac{s^{n-k} a_k}{s_k}(s\alpha)^k$$

is an equation of integral dependence for $s\alpha$ over B, seeing as s_k divides s^{n-k} for $k \leq n-1$. Therefore, $s\alpha$ is in B, so that $\alpha = \frac{s\alpha}{s}$ is in $S^{-1}B$. ∎

Lemma 3.2.6. Let A be an integral domain, L a finite extension of the field of fractions of A, and B be the integral closure of A in L. If S is a multiplicative set in A, then $S^{-1}B$ is the integral closure of $S^{-1}A$ in L.

Proof. By Lemma 3.2.4, $S^{-1}B$ is integral over $S^{-1}A$, so $S^{-1}B$ is contained in the integral closure of $S^{-1}A$. By Lemma 3.2.5, $S^{-1}B$ is integrally closed in L, and in particular, contains the integral closure of $S^{-1}A$ in L. ∎

Corollary 3.2.7. Let A be an integrally closed integral domain and S a multiplicative set in A. Then $S^{-1}A$ is integrally closed.

Proof. By definition, A is its own integral closure in K. By the Lemma, $S^{-1}A$ is integrally closed. ∎

Proposition 3.2.8. Let A be a Dedekind domain, and S a multiplicative set in A not containing 0. Then $S^{-1}A$ is a Dedekind domain or $S^{-1}A$ is a field.

Proof. Since 0 is not in S, we have that $S^{-1}A$ is an integral domain. We need to check that $S^{-1}A$ is Noetherian and integrally closed of dimension at most 1, as in this case, A is a field if A has dimension 0 or a Dedekind domain if A has dimension 1. By Lemma 3.2.3, $S^{-1}A$ is Noetherian, and by Corollary 3.2.7, $S^{-1}A$ is integrally closed. By Corollary 3.2.2, $\dim(S^{-1}A) \leq \dim(A) = 1$. ∎

Lemma 3.2.9. Let A be an integral domain. In the field of fractions of A, we have

$$A = \bigcap_{\mathfrak{p} \in \mathrm{Spec}(A)} A_{\mathfrak{p}}.$$

Proof. Clearly, we have that $A \subseteq A_{\mathfrak{p}}$ for every prime ideal $\mathfrak{p}$ of A. Suppose z is contained in $A_{\mathfrak{p}}$ for all $\mathfrak{p} \in \mathrm{Spec}(A)$. We form the ideal I generated by the set

$$\{s_{\mathfrak{p}} \in A \mid z = \frac{r_{\mathfrak{p}}}{s_{\mathfrak{p}}} \text{ in } A_{\mathfrak{p}}, \text{ some } r_{\mathfrak{p}} \in A\}.$$

If I were proper, then I would be contained in a maximal ideal $\mathfrak{m}$, but since the element $s_{\mathfrak{m}}$ is an element of $A \backslash \mathfrak{m}$, we cannot have $I \subseteq \mathfrak{m}$. Therefore, $I = A$. So, there exists a finite set of prime ideals $\{\mathfrak{p}_i\}_{1 \leq i \leq n}$ of A and coefficients $a_i \in A$ such that

$$1 = \sum_{i=1}^{n} a_i s_{\mathfrak{p}_i}.$$

If $z = \frac{r_i}{s_i}$ in each localization $A_{\mathfrak{p}_i}$, $1 \leq i \leq n$, we have that

$$z = z \left(\sum_{i=1}^{n} a_i s_{\mathfrak{p}_i} \right) = \sum_{i=1}^{n} \frac{r_i}{s_i} a_i s_i = \sum_{i=1}^{n} r_i a_i,$$

so that $z \in A$. ∎

3.3 Ramification

We now examine the factorization of primes in an integral extension as in [Mil20, pp.59-60]. Let A be a Dedekind domain with field of fractions K and let B be the integral closure of A in a finite separable extension L of K. Consider a prime ideal $\mathfrak{p}$ of A. The ideal generated by $\mathfrak{p}$ in B is often no longer a prime ideal of B, but since we have shown in Theorem 3.1.18 that B is a Dedekind domain, it makes sense to examine the factorization of the ideal generated by $\mathfrak{p}$ in B into a product of prime ideals of B. We will see that this factorization behaves in a predictable way, and that the degree of the corresponding extension of residue fields is determined by how $\mathfrak{p}$ ramifies in $\mathfrak{P}$. We start by introducing some definitions.

Definition 3.3.1. Take A and B as in the previous paragraph, and let $\mathfrak{p}$ be a prime ideal of A. Let $\mathfrak{p}B = \mathfrak{P}_1^{e_1} \dots \mathfrak{P}_g^{e_g}$ be the prime factorization of $\mathfrak{p}B$ in B.

1. If either g or any of the e_i are greater than 1, we say that $\mathfrak{p}$ is ramified in B.
2. If $\mathfrak{P}$ is a prime ideal of B appearing in the prime factorization of I, we say that $\mathfrak{P}$ divides $\mathfrak{p}$.
3. The number e_i is called the ramification index of $\mathfrak{p}$ in $\mathfrak{P}_i$.
4. The degree of the field extension $f_i = [B/\mathfrak{P}_i : A/\mathfrak{p}]$ is called the residue class degree of $\mathfrak{p}$ in $\mathfrak{P}_i$.

This section's main theorem will make use of different vector space structures induced by the action of A, B, and their quotients on various substructures of B. Let $\mathfrak{p}$ and $\mathfrak{P}$ be prime ideals of A and B respectively, and suppose that I is an ideal of B containing $\mathfrak{p}$. We first consider the quotient $B/\mathfrak{p}B$. We have that $\mathfrak{p}$ annihilates $B/\mathfrak{p}B$, so the action of A is well defined up to elements in $\mathfrak{p}$. More generally, $\mathfrak{p}B$ annihilates any ideal of B containing $\mathfrak{p}$. This gives us the following Lemma.

Lemma 3.3.2. Let A and B be as above, let $\mathfrak{p}$ be a prime ideal of A, and let I be an ideal of B containing $\mathfrak{p}B$. The action of A on B induces an $A/\mathfrak{p}$ vector space structure on B/I.

Let us now look at $M = \mathfrak{P}^k/\mathfrak{P}^{k+1}$, for some $k \geq 0$. Since $\mathfrak{P}^{k+1} = \mathfrak{P}\mathfrak{P}^k$, we have that $\mathfrak{P}$ annihilates M, so that

$$(b + \mathfrak{P}, m) \mapsto bm$$

endows M with the structure of a $B/\mathfrak{P}B$ module. From this module structure, we deduce that any $B/\mathfrak{P}B$-submodule of M must lift to a B-submodule of M. But we know that M is simple as a B-module, so M is a one dimensional vector space over $B/\mathfrak{P}B$ and we have the following.

Lemma 3.3.3. Let $\mathfrak{P}$ be an ideal of the Dedekind domain B. Then for any $k \geq 0$, the action of B on the ideal $\mathfrak{P}^k$ induces a $B/\mathfrak{P}$ vector space structure on $\mathfrak{P}^k/\mathfrak{P}^{k+1}$. Further, $\mathfrak{P}^k/\mathfrak{P}^{k+1}$ has dimension 1 over $B/\mathfrak{P}B$.

We are now ready to present the main result concerning factorization of prime ideals of the Dedekind domain A. This classic result gives us a formula relating the ramification index, the residue field extension degree, and the degree of the extension of the field of fractions. The proof is as in [Mil20].

Theorem 3.3.4. *Let $\mathfrak{p}$ be a prime ideal of the Dedekind domain A. Suppose that K is the field of fractions of A and that L is a separable extension of K with $[L:K]=n$. Write B for the integral closure of A in L. If $\mathfrak{P}_1,\dots,\mathfrak{P}_g$ are the prime ideals of B which divide $\mathfrak{p}$ so that $\mathfrak{p}=\mathfrak{P}_1^{e_1}\dots\mathfrak{P}_g^{e_g}$, then*

$$\sum_{i=1}^{g} e_i f_i = n.$$

Proof. We show that both sides of the equation are equal to the dimension m of $B/\mathfrak{p}B$ as an $A/\mathfrak{p}$ vector space. By the Chinese Remainder Theorem 2.1.5, we have

$$B/\mathfrak{p}B \cong \prod_{i=1}^{g} B/\mathfrak{P}_i^{e_i}$$

as a ring and also as an $A/\mathfrak{p}$ vector space. Since dimension is additive, it suffices to show that $\dim_{A/\mathfrak{p}}(B/\mathfrak{P}^{e_i}) = f_i e_i$. Since B is a Dedekind domain, the only ideals of B containing $\mathfrak{P}_i^{e_i}$ are the ideals $\mathfrak{P}_i^k$ for $0 \le k \le e_i$. So, $B/\mathfrak{P}_i^{e_i}$ has a filtration by ideals

$$B = \mathfrak{P}_i^0 \supset \mathfrak{P}_i^1 \supset \cdots \supset \mathfrak{P}_i^{e_i}$$

where $M_k = \mathfrak{P}_i^k/\mathfrak{P}_i^{k+1}$ is a one dimensional $B/\mathfrak{P}B$ vector space, by Lemma 3.3.2. But $f_i = \deg_{A/\mathfrak{p}}(B/\mathfrak{P}B)$, which means that each M_k is an f_i-dimensional $A/\mathfrak{p}$ vector space with respect to the induced action. By Lemma 3.3.2, $B/\mathfrak{P}_i^{e_i}$ is a $A/\mathfrak{p}$ vector space with respect to the same action, so the filtration by the $\mathfrak{P}_i^k$ induces a complete flag in $A/\mathfrak{P}_i^{e_i}$ of length $e_i f_i$. This shows that $\dim_{A/\mathfrak{p}}(B/\mathfrak{P}^{e_i}) = f_i e_i$.

To show the second equality, note that the natural map $A/\mathfrak{p} \to A_\mathfrak{p}/\mathfrak{p}A_\mathfrak{p}$ is an isomorphism. Put $k = A/\mathfrak{p}$ and $k_p = A_\mathfrak{p}/\mathfrak{p}A_\mathfrak{p}$. Since localization commutes with quotients, we have $(B/\mathfrak{p}B)_\mathfrak{p} \cong B_\mathfrak{p}/\mathfrak{p}B_\mathfrak{p}$. But $(B/\mathfrak{p}B)_\mathfrak{p} \cong B/\mathfrak{p}B \otimes A_\mathfrak{p}$, so that at the level of $k \cong k_\mathfrak{p}$ vector spaces

$$m = \dim_k(B/\mathfrak{p}B) = \dim_{k_\mathfrak{p}}(B/\mathfrak{p}B \otimes A_\mathfrak{p}) = \dim_{k_\mathfrak{p}}(B_\mathfrak{p}/\mathfrak{p}B_\mathfrak{p}).$$

It suffices to show that $\dim_{k_\mathfrak{p}}(B_\mathfrak{p}/\mathfrak{p}B_\mathfrak{p})$ is equal to n. By Proposition 3.2.8, $A_\mathfrak{p}$ is a local Dedekind domain and therefore has finitely many prime ideals. By Corollary 3.1.21, $A_\mathfrak{p}$ is a principal ideal domain. By Lemma 3.2.6, $B_\mathfrak{p}$ is the integral closure of $A_\mathfrak{p}$ in L, so that $B_\mathfrak{p}$ is finitely generated over $A_\mathfrak{p}$ and therefore free. Applying Corollary 2.4.15, $[L:K]=n$ is equal to the rank of $B_\mathfrak{p}$ over $A_\mathfrak{p}$. Base change of $B_\mathfrak{p}$ to $k_\mathfrak{p}$ then shows that $n = \mathrm{rk}_{A_\mathfrak{p}}(B_\mathfrak{p}) = \dim_{k_\mathfrak{p}}(B_\mathfrak{p}/\mathfrak{p}B_\mathfrak{p})$. ∎

Example 3.3.5. Consider $\mathbb{Q}(i)$ as an extension of $\mathbb{Q}$, so that $B = \mathbb{Z}(i)$ is the integral closure of $A = \mathbb{Z}$ in $\mathbb{Q}(i)$. Denote by $\mathfrak{p}$ and I the ideals generated by 2 in $\mathbb{Z}$ and $\mathbb{Z}(i)$ respectively. Since 2 is a prime number, $\mathfrak{p}$ is a prime ideal of $\mathbb{Z}$. As an element of $\mathbb{Z}(i)$, 2 admits the prime factorization $(1+i)(1-i)$, which, in particular, shows that 2 is not a prime element of $\mathbb{Z}(i)$. But $\mathbb{Z}(i)$ is a Euclidean domain, so the prime elements of $\mathbb{Z}(i)$ are precisely the generators of prime ideals, hence I cannot be a prime ideal of $\mathbb{Z}(i)$.

Let us compute the prime factorization of the ideal $I = (2)$ in $\mathbb{Z}(i)$. As we saw above, $2 = (1+i)(1-i)$ in $\mathbb{Z}(i)$, so we must have that $\mathfrak{P}$ is contained in the prime ideals $(1+i)$ and $(1-i)$. Since $\mathbb{Q}(i)$ is a quadratic extension of $\mathbb{Q}$, Theorem 3.3.4 says that we must have that $I = (1+i)(1-i)$, where the right-hand side is a product of ideals. By a brief argument as in [Neu99, p.3], we have that the units of $\mathbb{Z}(i)$ are the fourth roots of unity $\{\pm 1, \pm i\}$. We have that $1+i = i(1-i)$ so that the ideal generated by $(1+i)$ is equal to that generated by $(1-i)$. Hence, the prime factorization of I is $(1+i)^2$, and the ramification index e of $\mathfrak{p}$ in $(1+i)$ is 2. By the degree formula from Theorem 3.3.4, we must therefore have that $B/(1+i) \cong A/\mathfrak{p} = \mathbb{Z}/2\mathbb{Z}$.

3.4 Discrete valuation rings

In the proof of Theorem 3.3.4, we localized the integral closure B of the Dedekind domain A at a prime ideal $\mathfrak{p}$ of A to work with the nice properties afforded to finitely generated modules over a principal ideal domain. In fact, the localization $A_{\mathfrak{p}}$ was a local Dedekind domain. We would now like to understand rings of this type, the so-called discrete valuation rings. We present the following standard definitions and results concerning discrete valuation rings, which are referenced, for instance, in [Fo17].

A totally ordered group is an Abelian group equipped with a total order. Important examples of totally ordered groups include the positive real numbers under multiplication, with the usual order relation, and the integers with the usual order relation.

Definition 3.4.1. Let K be a field and G be a totally ordered group, with binary operation " $+$ " and identity element 0_G. A function

$$\nu \colon K^\times \to G$$

is called a (non-Archimedean) valuation if

$$\nu(xy) = \nu(x) + \nu(y)$$

for all $x, y \in K^\times$ and if ν respects the ultrametric inequality

$$\min\{\nu(x), \nu(y)\} \leq \nu(x + y)$$

for all $x, y \in K^\times$ such that $x + y \neq 0$. If K is equipped with a valuation, we say that K is valued field. Let

$$R := \{x \in K \mid \nu(x) \geq 0_G\}.$$

We call R the valuation ring of K.

An immediate consequence of the previous definition is that G must be Abelian. The following properties of the valuation ring of K follow immediately from the definition.

Lemma 3.4.2. Let K be a field with valuation $\nu\colon K^\times \to G$. The valuation ring R of K is a local ring with maximal ideal

$$\mathfrak{m} := \{0\} \cup \{x \in K^\times \mid \nu(x) > 0_G\},$$

K is the field of fractions of R, and R is integrally closed in K.

Lemma 3.4.3. Let G be a totally ordered additive group, K a valued field with valuation $\nu\colon K^\times \to G$, and let R be the associated valuation ring. Let $0_G \in G$ be identity element of G. Then

$$\nu(1) = 0_G$$

and

$$\nu(x^{-1}) = -\nu(x).$$

Suppose further that G has no elements of order 2. If $\nu(x) \neq \nu(y)$ and $x + y \neq 0$, then

$$\nu(x + y) = \min\{\nu(x), \nu(y)\}.$$

Proof. We have that $\nu(1) = \nu(1 \cdot 1) = \nu(1) + \nu(1)$, so that by the cancellation law in G, $\nu(1) = 0_G$. Then, $e = \nu(x^{-1}x) = \nu(x^{-1}) + \nu(x)$ so that $\nu(x^{-1}) = -\nu(x)$.

Let $x, y \in K$ satisfy $\nu(x) \neq \nu(y)$ with $x + y \neq 0$, and without loss of generality, $\nu(x) < \nu(y)$. Then, $\nu(x) = \nu(x + y - y) \geq \min\{\nu(x + y), \nu(-y)\}$. But $\nu(-y) = \nu(-1 \cdot y) = \nu(-1) + \nu(y)$. We have that $0_G = \nu((-1)^2) = \nu(-1) + \nu(-1)$, so that by the assumption that G has no non-trivial self inverse elements, $\nu(-1) = 0_G$. So, $\nu(-y) = \nu(y)$, and therefore

$$\nu(x) \geq \min\{\nu(x + y), \nu(y)\} = \nu(x + y).$$

Since

$$\nu(x + y) \geq \min\{\nu(x), \nu(y)\},$$

we get $\nu(x) = \nu(x + y)$. ∎

Corollary 3.4.4. Let R be a valuation ring of K, and let $x \in K^\times$. Then $x \in R$ or $x^{-1} \in R$.

Proof. Since $\nu(x) = -\nu(x^{-1})$, at least one of $\nu(x)$ and $\nu(x^{-1})$ is non-negative. ∎

Definition 3.4.5. Let $\nu\colon K^\times \to G$ be a valuation on the field K. We say ν is a discrete valuation if G is isomorphic to the group $\mathbb{Z}$ and if the valuation is surjective. In this case, we say that K is a discretely valued field. An integral domain R is called a discrete valuation ring, or DVR, if there exists a discrete valuation on the field of fractions of R such that R is the associated valuation ring.

Note that when the valuation is discrete, we denote the identity element 0_G of G by 0, considering that the valuation is a group homomorphism from the multiplicative group of the discretely valued field to the additive group $\mathbb{Z}$. Having defined a discrete valuation, we can state the following result, which follows at once from Lemma 3.4.3.

Corollary 3.4.6. The valuation $\nu\colon R\backslash\{0\} \to \mathbb{Z}$ on a discrete valuation ring R satisfies $\nu(1) = 0$ and

$$\min\{\nu(x), \nu(y)\} = \nu(x + y)$$

for all $x, y \in R$ such that $\nu(x) \neq \nu(y)$ and $x + y \neq 0$.

Proof. We note that the additive group $\mathbb{Z}$ has no non-trivial self inverse elements, so we apply the first and third points of Lemma 3.4.3. ∎

We want to highlight that we have not formally defined the value of ν at 0. However, we do like to think of 0 as having infinite valuation. As we will see, elements of positive valuation are precisely those which are not invertible in the valuation ring.

Corollary 3.4.7. The units of a valuation ring R are the elements of valuation equal to the group identity 0_G.

Proof. If $x \in R$ has positive valuation, then its inverse has negative valuation, so x^{-1} is not in R. If $x \in R$ satisfies $\nu(x) = 0_G$, then x^{-1} also satisfies $\nu(x^{-1}) = 0_G$. So $x^{-1} \in R$. ∎

We list some properties of discrete valuation rings.

Lemma 3.4.8. Let R be a discrete valuation ring with maximal ideal $\mathfrak{m}$ and associated discrete valuation ν. Write 1 for the generator of the valuation group of K. Then,

1. R is a principal ideal domain. In particular, R is a Dedekind domain.
2. The maximal ideal $\mathfrak{m}$ is generated by any element $\varpi \in R$ such that $\nu(\varpi) = 1$.
3. Fix $\varpi \in R$ with $\nu(\varpi) = 1$. The non-zero ideals of R are the principal ideals $\varpi^k R$, $k \in \mathbb{N}$.
4. R is a maximal subring of its field of fractions K.

Proof. Let I be a proper, non-zero ideal of R. Since $\nu(R)$ is a subset of the non-negative integers and the units of R are the elements of valuation 0, $\nu(I)$ is a subset of the positive integers. Therefore, $\nu(I)$ attains a minimum $k > 0$, so there exists $x \in I$ such that $\nu(x) = k$. Now, let $y \in I$ with $\nu(y) = \ell \geq k$. We have that $x^{-1} \in K$ and $\nu(x^{-1}y) = -\nu(x) + \nu(y) = -k + \ell \geq 0$, so that $x^{-1}y \in R$. Hence, $y = xx^{-1}y \in xI$. So, I is the principal ideal xI. This shows that R is a principal ideal domain, and as we have seen, principal ideal domains are Dedekind domains.

Taking $I = \mathfrak{m}$, we have by surjectivity of ν that there exists $\varpi \in \mathfrak{m}$ such that $\nu(\varpi) = 1$, so that specializing the argument above, $\mathfrak{m} = \varpi R$. Since $\nu(\varpi^k) = \sum_{i=1}^{k} \nu(\varpi) = k$ for arbitrary $k \in \mathbb{Z}$, and $\nu(R) = \mathbb{Z}$ by assumption, we can take

ϖ^k in the argument above as the element x of minimum valuation k of some non-zero ideal I of R. Moreover, each $k \in \mathbb{Z}$ defines a distinct non-zero ideal $\varpi^k R$ of R. So, the non-zero ideals of R are the $\varpi^k R$, $k \in \mathbb{Z}$. Lastly, suppose R is properly contained in some subring S of K, and let $x \in S$. Since x is not contained in R, we have by Corollary 3.4.4 that x^{-1} is contained in R, and, for the same reason, we have that x^{-1} is not a unit of R. Therefore x^{-1} is an element of the maximal ideal $\mathfrak{m}$. Since the multiplicative set generated by x^{-1} meets $\mathfrak{m}$, Theorem 3.2.1 implies that the ring obtained by inverting x^{-1} in K is all of K. So, $S = K$. ∎

Corollary 3.4.9. Let $\varpi \in R$ be such that $\nu(\varpi) = 1$. Up to multiplication by a unit, ϖ is the only prime element of R.

Proof. Since R is a principal ideal domain, this follows from our characterization of proper ideals. ∎

Definition 3.4.10. Any element $\varpi \in R$ satisfying $\nu(\varpi) = 1$ as in the Lemma is called a uniformizer.

Lemma 3.4.11. Let A be an integral domain. Then A is a discrete valuation ring if and only if A is a local Dedekind domain.

Proof. The forward direction holds by Lemma 3.4.8. Suppose that A is a local Dedekind domain. Then, A has a unique non-zero prime ideal $\mathfrak{p}$, so that by Corollary 3.1.21, A is a principal ideal domain. Let ϖ generate the maximal ideal $\mathfrak{p}$. Then $\nu(u\varpi^k) = k$ defines a discrete valuation on A. Since A is local, every element of the field of fractions of A is either invertible in A, is contained in $\mathfrak{p}$, or is the inverse of an element of $\mathfrak{p}$, i.e., is uniquely of the form $u\varpi^k$ for some $u \in A^\times$ and $k \in \mathbb{Z}$. So, ν extends uniquely to K and defines a discrete valuation on K. Since the elements of K which are not already in A are those of the third type, we recover A as the valuation ring of K. ∎

Corollary 3.4.12. Let $\mathfrak{p}$ be a non-zero prime ideal of the Dedekind domain A. The localization $A_\mathfrak{p}$ is a discrete valuation ring.

Proof. By Theorem 3.2.1 and Proposition 3.2.8, $A_\mathfrak{p}$ is a Dedekind domain with unique non-trivial maximal ideal $\mathfrak{p}A_\mathfrak{p}$. ∎

Example 3.4.13. It is reasonable to investigate where the construction of the proof of Lemma 3.4.11 fails to give us a discrete valuation ring when we relax the assumption that the PID A is local. Indeed, we could still try to define a valuation (with respect to a chosen prime of A) on the field of fractions of A. To this end, assume that A is a non-local principal ideal domain. We consider distinct prime ideals $\mathfrak{p}_1, \mathfrak{p}_2$ of A with respective generators ϖ_1, ϖ_2. In particular, ϖ_1 and ϖ_2 are irreducible (and prime). We can then define a valuation $\nu_1 \colon A \to \mathbb{Z}$ by setting $\nu(\varpi_1) = 1$, for some ϖ_1 which generates $\mathfrak{p}_1$. This indeed extends to a discrete valuation on K by unique factorization, as in the proof of the Lemma. However, A would not be the valuation ring of K with respect to ν_1. Indeed ϖ_1 does not appear in the prime factorization of ϖ_2, so $\nu_1(\varpi_2) = 0$. This means that $\nu(\varpi_2^{-1}) = 0$. Hence, ϖ_2^{-1} is in the valuation ring of K. By virtue of being prime, ϖ_2 is not invertible in A, so the valuation ring of K is strictly larger than A. Since $\mathfrak{p}_2 = \varpi_2 A$ was an arbitrary prime, we get that the valuation ring of K with respect to ν_1 contains the inverse of every prime element not contained in $\mathfrak{p}_1$, and therefore of every non-unit not contained in $\mathfrak{p}_1$. In other words, the valuation ring of the fraction field of A with respect to a discrete valuation induced by a uniformizer ϖ for $\mathfrak{p}$ is the localization $A_{\mathfrak{p}}$.

More generally, the example leads us to the following description of the valuation rings which occur when we embed a discrete valuation ring A in a finite separable extension of the fraction field of A.

Proposition 3.4.14. Let A be a discrete valuation ring with field of fractions K, and let B be the integral closure of A in a finite separable extension L of K. The assignment

$$\mathfrak{P} \mapsto B_{\mathfrak{P}}$$

induces a one-to-one correspondence between prime ideals of B dividing $\mathfrak{p}$ and discrete valuation rings of L with valuation extending the valuation of A.

Proof. Let $\mathfrak{P}$ be a prime of B dividing the maximal ideal $\mathfrak{p}$ of A. By Theorem 3.1.18, B is a Dedekind domain with field of fractions L, so that by Lemma 3.4.11, the localization $B_{\mathfrak{P}}$ is a discrete valuation ring. So, the codomain is correct. We define the inverse map. If $\mathcal{O}_E$ is a discrete valuation ring on L with valuation ν and maximal ideal $\mathfrak{q}$, we set $\varphi(\mathcal{O}_E) = \mathfrak{q} \cap B$. We immediately obtain that φ is a left inverse to that of the statement. Indeed, by Theorem 3.2.1, the intersection of the local ring $B_{\mathfrak{P}}$

with B must be the only maximal ideal of B which does not meet the multiplicative set $B\backslash\mathfrak{P}$, so that $\varphi(B_{\mathfrak{P}}) = \mathfrak{P}$. It remains to show that the codomain of φ is correct and that φ is a right inverse. Let ν be a discrete valuation on L extending that on A, and let $\mathfrak{q}$ be the maximal ideal of the valuation ring $\mathcal{O}_E$ of L. We must have that B is contained in $\mathcal{O}_E$. Since ν extends the valuation on A, then A is contained in $\mathcal{O}_E$, and as $\mathcal{O}_E$ is a discrete valuation ring, $\mathcal{O}_E$ is integrally closed. In particular, $\mathcal{O}_E$ contains all roots in L of monic polynomials with coefficients in A. Seeing as every element of B is integral over A, B must be contained in $\mathcal{O}_E$. We therefore get that $\varphi(\mathcal{O}_E) = \mathfrak{q} \cap B$ is a prime of B. Write $\mathfrak{P} = \varphi(\mathcal{O}_E)$. Since ν extends the valuation of A, we have that $\mathfrak{p} \subseteq \mathfrak{q}$ and therefore $\mathfrak{p} \subseteq \mathfrak{P}$. This means that $\mathfrak{P}$ divides $\mathfrak{p}$. To show that φ is a right inverse, we show $B_{\mathfrak{P}} = \mathcal{O}_E$. Seeing as the inclusion $B \to \mathcal{O}_E$ maps the multiplicative set $B\backslash\mathfrak{P}$ to a subset of $\mathcal{O}_E\backslash\mathfrak{q}$, the localization $B_{\mathfrak{P}}$ is a subring of $\mathcal{O}_E$ in L. By Corollary 3.4.12, $B_{\mathfrak{P}}$ is a discrete valuation ring. Applying Lemma 3.4.8, we must have that $B_{\mathfrak{P}}$ coincides with $\mathcal{O}_E$. ∎

To finish, we adapt the language of the previous section on ramification of primes in Dedekind domains to the case of extensions of local Dedekind domains. The proof of the following result will give us a glimpse into the theory of differentials of modules, which will serve to give us a criterion for separability. We give the proof from [Fo17, p.299], abstracting away the background for the relevant facts about differentials when necessary, for brevity. This is because the proof of the results about differentials requires the development of Hochschild cohomology, which is beyond the scope of this thesis. Hochschild cohomology is a tool used to classify different deformations, or infinitesimal transformations, of algebras by defining a series of groups. After presenting the proof of the following result, we will give a brief explanation of how Hochschild cohomology is used.

Definition 3.4.15. Let A be an R-algebra, and let M be an A-module. An R-derivation from A to M is an element $\delta \in \mathrm{Hom}_R(A, R)$ which satisfies the Leibniz rule

$$\delta(ab) = a\delta(b) + \delta(a)\, b$$

for all $a, b \in A$. The set of R-derivations from A to M is denoted $\mathrm{Der}_R(A, M)$.

Lemma 3.4.16. Let $A \to B$ be a finite homomorphism of local rings, and let $\mathfrak{p}$ be the maximal ideal of A. Then, B is a separable A-algebra if $B/\mathfrak{p}B$ is a separable $A/\mathfrak{p}$-algebra.

Proof. Write $J_{B/\mathfrak{p}}$ for the kernel of the exact sequence defining separability of the $A/\mathfrak{p}$-algebra $B/\mathfrak{p}B$, and define the $A/\mathfrak{p}$-derivation

$$\delta\colon B/\mathfrak{p}B \to J_{B/\mathfrak{p}}$$

$$a \mapsto a \otimes 1 - 1 \otimes a.$$

Write J_B for the separability kernel of the A-algebra B. Recall that the action of $B/\mathfrak{p}B$ on $J_{B/\mathfrak{p}}$ is given by $a \cdot z = (a \otimes 1 - 1 \otimes a)z$. By a series of Lemmas in [Fo17, §8.2], we have that the derivation δ is inner, so that there exists $z \in J_{B/\mathfrak{p}}$ such that

$$\delta(a) = az - za = (a \otimes 1 - 1 \otimes a)z = \delta(a)z.$$

Since we can write an element $x \in J_{B/\mathfrak{p}B}$ as $x = \sum_i x_i \otimes y_i$ with $\sum_i x_i y_i = 0$, we have $\sum_i x_i(1 \otimes y_i - y_i \otimes 1) = \sum_i x_i \otimes y_i = x$. So, the image of the differential δ generates $J_{B/\mathfrak{p}B}$ as a $B/\mathfrak{p}B$-module. By combining our expression for δ with this fact,

$$J_{B/\mathfrak{p}B} = B/\mathfrak{p}\delta(B/\mathfrak{p}) = B/\mathfrak{p}\delta(B/\mathfrak{p})z = J_{B/\mathfrak{p}B}z.$$

By some straightforward algebraic manipulations ([Fo17, p.270]), we have that

$$J_{B/\mathfrak{p}B} = J_B/(\mathfrak{p}J_B).$$

Since $B \otimes B$ is finitely generated as an A-module and the exact sequence of separability is split exact, J_B is finitely generated as an A-module. We apply Nakayama's Lemma 2.1.27 to lift the $A/\mathfrak{p}$-generator z of $J_{B/\mathfrak{p}B}$ to a generator Z of J_B. The B^e-module homomorphism $\phi\colon B^e \to J_B$ defined by

$$x \mapsto xZ$$

then splits the separability exact sequence for B on the left. Indeed, when restricted to the finitely generated A-module J_B, the relation $J_A = J_A Z$ means that ϕ is onto J_B and therefore bijective. ∎

We now briefly discuss Hochschild cohomology. Let A be an R-algebra, where A is not necessarily commutative, and let M be a two-sided A/R module, as in Definition 2.5.7. By elementary manipulations of the relevant formulas, it can be shown that the n^{th} Hochschild cohomology group with coefficients in M, denoted by $\mathrm{H}^n(A, M)$, $n \in \mathbb{Z}$, have the following interpretations in low dimension. We have that

$$\mathrm{H}^0(A, M) = \{m \in M \mid am = ma, \text{ for all } a \in A\}.$$

Then, $\mathrm{H}^1(A, M)$ is the quotient of the module of R-derivations from A to M by inner derivations, or derivations of the form $a \mapsto am - ma$ for some fixed $m \in M$. We then have that $\mathrm{H}^2(A, M)$ classifies twistings of the multiplication formula on A by the addition of an infinitesimal term. It can be shown that an R-algebra A is separable if and only if A has projective dimension 0, so that the cohomology groups $\mathrm{H}^n(A, M)$, computed by considering a minimal projective resolution, are 0 for all $n > 0$ and all two-sided A/R-modules M. In the proof of the Lemma, we used that $B/\mathfrak{p}B$ is separable over $A/\mathfrak{p}$ to get that $\mathrm{H}^1\left(B/\mathfrak{p}B, J_{B/\mathfrak{p}B}\right) = 0$, yielding that the derivation $a \mapsto a \otimes 1 - 1 \otimes a$ is inner, consistently with our interpretation of this cohomology group as the quotient of derivations by inner derivations. For details of the claims made above and further exposition, see [Wi19].

To finish the section, we state the following Corollary of Lem 3.4.16.

Corollary 3.4.17. Let $A \to B$ be a local homomorphism of discrete valuation rings. If the maximal ideal $\mathfrak{p}$ of A generates the maximal ideal $\mathfrak{P}$ of B and the residue field $B/\mathfrak{P}$ is separable over $A/\mathfrak{p}$, then B is a separable A-algebra.

Proof. It suffices to observe that $\mathfrak{p}B = \mathfrak{P}$ to apply the Lemma. ∎

The hypotheses of the Corollary might be familiar to specialists of non-Archimedean local fields. Indeed, these fields have a canonical discrete valuation. We say that a finite extension L/K of non-Archimedean local fields is unramified if the maximal ideal of the valuation ring of K does not ramify in the valuation ring of L, and if the corresponding extension of residue fields is separable. We meet these fields in the following section.

3.5 The p-adic numbers

We now introduce a well-known example of a discrete valuation ring — the ring of p-adic integers. We draw on material from the early chapters of [Cas86] and [Lor07] to give a quick exposition of the p-adic integers, with the goal of identifying their structure as discrete valuation rings and deducing some elementary properties with which we can work.

Definition 3.5.1. Let K be a field, and let $|\ |\colon K \to \mathbb{R}$ be a function. We say $|\ |$ is an absolute value or a norm on K if

1. $|xy| = |x||y|$ for all $x, y \in K$,
2. $|x| = 0$ if and only if $x = 0$,
3. and $|x + y| \leq |x| + |y|$ for all $x, y \in K$.

If further, $|x + y| \leq \max\{|x|, |y|\}$, we say the absolute value is non-Archimedean.

Definition 3.5.2. We define the p-adic valuation $\nu_p\colon \mathbb{Q}\backslash\{0\} \to \mathbb{R}$ by

$$\nu_p(x) = k,$$

where k is the exponent of p in the unique representation (up to ± 1) of $x = p^k \frac{a}{b}$, with

$$gcd(a, b) = 1$$

and p not dividing a or b. We then define the p-adic absolute value $|\ |_p\colon \mathbb{Q} \to \mathbb{R}$ by

$$|x|_p = \begin{cases} p^{-\nu_p(k)} & x \neq 0 \\ 0 & x = 0 \end{cases}.$$

Lemma 3.5.3. Let $p \in \mathbb{Z}$ be prime. The p-adic absolute value $|\ |_p$ satisfies

$$|xy|_p = |x|_p |y|_p.$$

Moreover, $|\ |_p$ satisfies the ultrametric inequality

$$|x + y|_p \leq \max\{|x|_p, |y|_p\},$$

for all $x, y \in \mathbb{Q}$.

Proof. Let $x = \frac{a}{b}$ and $y = \frac{c}{d}$ be reduced elements of $\mathbb{Q}$. If $|\frac{a}{b}|_p = p^k$ and $|\frac{c}{d}|_p = p^\ell$ for $k, \ell \in \mathbb{Z}$, then it follows from unique factorization of ac and bd in the integers that

$$|\frac{ac}{bd}|_p = p^{k+\ell} = |\frac{a}{b}|_p |\frac{c}{d}|_p.$$

We will show that $|\ |_p$ satisfies the ultrametric inequality. Let $a, b, c, d, k, \ell \in \mathbb{Z}$ be as above, with $p \nmid abcd$. Suppose without loss of generality that $k \leq \ell$, so that $|\frac{c}{d}|_p \leq |\frac{a}{b}|_p$. Then

$$x + y = p^k \frac{ad + p^{\ell-k}bc}{bd}.$$

It is immediate that p does not divide the denominator bd, so that $|bd|_p = 1$. Consider the p-adic valuation of the numerator $ad + p^{\ell-k}bc$. There exist $u, \lambda \in \mathbb{Z}$ with $p \nmid u$ and such that $up^\lambda = ad + p^{\ell-k}bc$. We have that λ is certainly non-negative, as $k \leq \ell$ guarantees that the expression $ad + p^{\ell-k}bc$ is an integer. This means that $|up^\lambda| \leq 1$. By the multiplicativity property that we have just proved,

$$|p^k \frac{ad + p^{\ell-k}bc}{bd}|_p = |p^k|_p \frac{|up^\lambda|_p}{|bd|_p} = \frac{1}{p^k}\frac{1}{p^\lambda} \leq \frac{1}{p^k} = \max\{|x|_p, |y|_p\},$$

as desired. ∎

Corollary 3.5.4. The p-adic absolute value $|\ |_p$ is a non-Archimedean absolute value on $\mathbb{Q}$.

Proof. It is immediate from the definition of $|\ |_p$ that $|x|_p = 0$ if and only $x = 0$. Since -1 is coprime to p and the p-adic absolute value is multiplicative, $\delta(x, y) = \delta(y, x)$. It is an instructive exercise from undergraduate real analysis to show that a function on $\mathbb{Q}$ satisfying the ultrametric inequality also satisfies the triangle inequality. We have therefore checked that $\mathbb{Q}$ endowed with the p-adic valuation is an absolute value satisfying the non-Archimedean inequality. ∎

Corollary 3.5.5. The p-adic valuation ν_p is a valuation in the sense of Definition 3.4.1. So, $\mathbb{Q}$ equipped with the p-adic valuation is a discretely valued field.

Proof. The operation $\nu \mapsto p^{-\nu}$ is an order-reversing bijection between the integers and the image of $|\ |_p$ on $\mathbb{Q}\backslash\{0\}$. Therefore, ν is additive and since $|x + y|_p \leq \max\{|x|_p, |y|_p\}$, we have that $\min\{\nu_p(x), \nu_p(y)\} \leq \nu_p(x + y)$ for all $x, y \in \mathbb{Q}$ such that $x + y \neq 0$. ∎

Remark 3.5.6. The ring $\mathbb{Z}_{(p)}$ is nothing but the localization of $\mathbb{Z}$ away from the prime ideal (p), explaining the choice of notation. Indeed, a reduced fraction $\frac{a}{b} \in \mathbb{Q}$ has non-negative p-adic valuation if and only if p does not divide b. This description is equivalent to that for the ring obtained from $\mathbb{Z}$ by inverting all elements of $\mathbb{Z}$ that

are not divisible by p, i.e., this ring is the localization of $\mathbb{Z}$ away from the prime ideal (p).

We perform the quintessential operation from real analysis and complete $\mathbb{Q}$ with respect to the absolute value $|\ |_p$.

Definition 3.5.7. We say a field K is complete with respect to an absolute value if every Cauchy sequence in K has a limit in K.

Remark 3.5.8. We can define the completion of a field equipped with an absolute value exactly as with the rationals and the usual absolute value, using equivalence classes of Cauchy sequences. The absolute value of the completion is defined as the limit taken in $\mathbb{R}$ of the absolute values of elements of the Cauchy sequence. For details, see [Cas86].

Definition 3.5.9. Let K be a field with absolute value $|\ |$ and let $\hat{K}$ be a field containing K, with absolute value $\|\ \|$. Suppose that $\|\ \|$ restricts to $|\ |$ on $K \subseteq \hat{K}$. We say that $\hat{K}$ is a completion of K if $\hat{K}$ is complete with respect to $\|\ \|$ and if K is dense in $\hat{K}$.

Definition 3.5.10. Let p be a prime number and endow $\mathbb{Q}$ with the p-adic absolute value $|\ |_p$. The completion of $\mathbb{Q}$ with respect to $|\ |_p$ is called the field of p-adic numbers and is denoted by $\mathbb{Q}_p$. We also denote the absolute value on $\mathbb{Q}_p$ by $|\ |_p$. The set of elements

$$\mathbb{Z}_p := \{x \in \mathbb{Q}_p \mid |x|_p \leq 1\}$$

is known as the ring of p-adic integers.

Remark 3.5.11. Since the absolute value of an element $\{x_n\}_{n\geq 0}$ of $\mathbb{Q}_p$ is defined as the (real-analytic) limit of the absolute values $|x_n|_p$, and all elements of $\mathbb{Q}_p$ are limits of elements of $\mathbb{Q}$, it can be worked out that the unit ball of $\mathbb{Q}_p$ is the completion of the unit ball of $\mathbb{Q}$. Therefore, the p-adic integers $\mathbb{Z}_p$ are the (analytic) completion of $\mathbb{Z}_{(p)} \subseteq \mathbb{Q}$. However, the operation of constructing $\mathbb{Z}_p$ and $\mathbb{Q}_p$ from $\mathbb{Z}$ and $\mathbb{Q}$ can also be performed algebraically. In the algebraic setting, $\mathbb{Z}_p$ is known as the formal completion of the ring $\mathbb{Z}_{(p)}$. Then, $\mathbb{Q}_p$ is defined as the field of fractions of $\mathbb{Z}_p$. This operation is explored in more detail in [Eis95, §8].

Let us return to the analytic theory. We state but do not prove the following technical result from [Cas86, p.16], and present its corollary: the p-adic integers are a discrete valuation ring.

Lemma 3.5.12. Let $|\ |$ be an absolute value on a field K. Then, $|\ |$ is non-Archimedean if and only if $|x| \leq 1$ for all x in the ring R generated by 1 in K.

Corollary 3.5.13. The extension of $\nu_p \colon \mathbb{Q} \to \mathbb{R}$ to $\mathbb{Q}_p$ endows $\mathbb{Q}_p$ with the structure of a discretely valued field. Therefore, the ring $\mathbb{Z}_p$ of p-adic integers is a discrete valuation ring.

Proof. As the completion of $\mathbb{Q}$ with respect to $|\ |_p$, any $\{x_n\}_{n\geq 0} \in \mathbb{Q}_p$ satisfies $|\lim_{n\to\infty} x_n|_p = \lim_{n\to\infty} |x_n|_p$. Any such convergent sequence of absolute values is necessarily eventually constant by the discreteness of $|\ |_p$. The corresponding sequence $\{\nu_p(x_n)\}_{n\geq 0}$ must then also be eventually constant. Therefore, we can extend $\nu_p \colon \mathbb{Q} \to \mathbb{R}$ to all of $\mathbb{Q}_p$ by setting $\nu_p\left(\lim_{n\to\infty}\right) = \lim_{n\to\infty}$. We have that $\nu_p \colon \mathbb{Q}_p \to \mathbb{R}$ is well defined since $\lim_{n\to\infty} \nu_p\left(x_n\right)$ is eventually constant whenever $\lim_{n\to\infty} |x_n|_p$ is eventually constant. The ring generated by 1 in $\mathbb{Q}_p$ is the image of $\mathbb{Z}$ in $\mathbb{Q}_p$. This means that the integers in $\mathbb{Q}_p$ have absolute value at most 1, since the absolute value of $\mathbb{Q}_p$ restricts to $|\ |_p$ on $\mathbb{Z}$, and $\nu_p(x) \geq 0$ for all $x \in \mathbb{Z}$ so $p^{-\nu_p(x)\leq 1}$ for all $x \in \mathbb{Z}$. We then apply Lemma 3.5.12 to obtain that $|\ |_p$ satisfies the non-Archimedean inequality. Since $\nu \mapsto p^{-\nu}$ is an order-reversing bijection between the integers and $\{p^{-k} \mid k \in \mathbb{Z}\}$, this shows by continuity that ν_p is a non-Archimedean valuation on $\mathbb{Q}_p$, in the sense of Definition 3.4.1. ∎

Discrete valuation rings are, in particular, local rings. If K is a discretely valued field, we define the residue field of K to be the residue field of its valuation ring. We will now describe the residue field of $\mathbb{Z}_{(p)}$, and then appeal to further theory to describe the residue field of $\mathbb{Q}_p$.

Lemma 3.5.14. Let p be a prime number and endow $\mathbb{Q}$ with the p-adic valuation $|\ |_p$. Then the residue field of $\mathbb{Q}$ is isomorphic to $\mathbb{Z}/p\mathbb{Z}$.

Proof. Using the fact that the integer b is always invertible mod p, it is straightforward to check that the ring homomorphism defined by the formula

$$\frac{a}{b} \mapsto [a][b]^{-1} + p\mathbb{Z}$$

is an isomorphism $\mathbb{Z}_{(p)} \to \mathbb{Z}/p\mathbb{Z}$. ∎

Proposition 3.5.15. Suppose that K is a field with non-Archimedean absolute value $|\ |$ and that $\hat{K}$ is the completion of K, endowed with the absolute value $\|\ \|$. The inclusion $K \to \hat{K}$ induces an inclusion of valuation rings $i\colon \mathcal{O} \to \hat{\mathcal{O}}$. Furthermore, the inclusion i induces an isomorphism of residue fields

$$\overline{K} \cong \overline{\hat{K}}.$$

Proof. See, for instance, [Cas86, p.41]. ∎

Corollary 3.5.16. The residue field of $\mathbb{Q}_p$ is the finite field $\mathbb{Z}/p\mathbb{Z}$.

Proof. The field $\mathbb{Q}_p$ is the completion of $\mathbb{Q}$ with respect to $|\ |_p$, which has residue field $\mathbb{Z}/p\mathbb{Z}$ by Lemma 3.5.14. ∎

So, we have just shown that $\mathbb{Q}_p$ is a discretely valued field with valuation ring $\mathbb{Z}_p$ and residue field isomorphic to $\mathbb{Z}/p\mathbb{Z}$. We now state some facts which will allow us to say that all finite extensions of $\mathbb{Q}_p$ behave in a similar manner. The following result, which we state without proof, will give us access to a broad class of examples of discrete valuation rings. The proof given in [Cas86, p.114-199] is done by explicitly constructing an absolute value on E using the algebraic field norm for the extension $E/\mathbb{Q}_p$, then showing that this absolute value is unique by using properties of normed vector spaces.

Lemma 3.5.17. Let E be a finite field extension of $\mathbb{Q}_p$. There exists a unique discrete valuation on E extending the p-adic valuation on $\mathbb{Q}_p$. The residue field of E is a finite extension of $\mathbb{Z}/p\mathbb{Z}$. In particular, E is a complete discretely valued field.

Proof. The proof, given over the course of several sections in [Cas86], is accessible to those with a background in undergraduate real analysis. However, it is too long to reproduce here. ∎

Definition 3.5.18. We say the field E is a p-adic field if it is a finite field extension of $\mathbb{Q}_p$, equipped with the unique extension of the p-adic valuation on $\mathbb{Q}_p$ to E. Let $K \subseteq E$ be a finite extension of p-adic fields, with respective valuation rings $\mathcal{O}_K$ and $\mathcal{O}_E$. We say the extension E/K is unramified if the maximal ideal of $\mathcal{O}_K$ is unramified in $\mathcal{O}_E$.

More generally, we call a field equipped with a discrete non-Archimedean absolute value, complete with respect to this absolute value, and having a finite residue field a *non-Archimedean local field*. Non-Archimedean local fields are precisely the p-adic fields and the rings of univariate formal Laurent series over a finite field. The proof of this fact is rather interesting, as it makes use of many different properties that are inherent to non-Archimedean local fields: Hensel's root lifting Lemma, the topology and normed structure of their field extensions, and Ostrovsky's theorem characterizing absolute values on $\mathbb{Q}$. For details, see [Neu99, p.135]. Many of the results in this section are versions of facts which hold more broadly for structures defined over non-Archimedean local fields. We note that since the valuation on a non-Archimedean local field is discrete, the valuation ring of such a field is a discrete valuation ring. Let ϖ be a uniformizer for the non-Archimedean local field K. A surprising fact about non-Archimedean local fields is that a power series in the field has Cauchy partial sums if and only if the terms of the power series tend to 0. So, any series of the form $\sum_{j\geq k} a_j \varpi^j$ for $a_j \in \mathcal{O}_K$ has Cauchy partial sums, and since we have imposed that our field is complete, the series is convergent in K. Conversely, we have the following description of the elements of K.

Lemma 3.5.19. Let K be a non-Archimedean local field, let ϖ be a uniformizer for the valuation ring $\mathcal{O}_K$ of K, and let $\{x_1, \ldots, x_q\}$ be a set of representatives in $\mathcal{O}_K$ for the equivalence classes of the residue field of $\mathcal{O}_K/\varpi\mathcal{O}_K$. Then, any non-zero element of K is of the form

$$\alpha = \sum_{j\geq k} a_j \varpi^j,$$

where $k = \nu(\alpha)$ and where each $a_j \in \{x_1, \ldots, x_q\}$ is uniquely determined.

Proof. By multiplying with an appropriate power of ϖ, we can assume $\nu(\alpha) = 0$. We give an expansion of α as a power series in ϖ with coefficients in the set $\mathcal{S} = \{x_1, \ldots, x_q\}$. First, put a_0 to be the representative in $\mathcal{S}$ for α mod $\varpi\mathcal{O}_K$. Then, $\alpha - a_0 = \varpi b_0$ for some $b_0 \in \mathcal{O}_K$. In the same way, can write $b_0 = a_1 + \varpi\mathcal{O}_K$ for some $a_1 \in \mathcal{S}$, so that $\alpha = a_0 + \varpi(a_1 + \varpi b_1)$, some $b_1 \in \mathcal{O}_K$. We can repeat this process ad infinitum for the successive b_i's, and see that the a_i thus obtained are uniquely determined by the process. As in the paragraph preceding the Lemma, the series $\sum_{j\geq 0} a_j \varpi^j$ converges in K, and by comparing the partial sums of this series with α, we see that the series converges to α. ∎

We now present Hensel's Lemma, a result which allows us to lift factorizations of polynomials and simple roots from the residue field to a non-Archimedean local field, given some mild hypotheses. The interested reader should see [Lor07, p.56] for details of the proof, which uses Newton's method (from calculus) in the p-adic setting.

Theorem 3.5.20 (Hensel's Lemma)**.** *Let K be a non-Archimedean local field with valuation ring $\mathcal{O}_K$ and residue field $\overline{K}$. Let $f(X) \in \mathcal{O}_K[X]$ have image $\overline{f}(X)$ in $\overline{K}[X]$. If there exist $\varphi, \psi \in \overline{K}[X]$ such that φ, ψ are relatively prime and*

$$\overline{f}(X) = \varphi\psi$$

over the residue field, then there exist $g(X), h(X) \in \mathcal{O}_K[X]$ such that

$$f(X) = g(X)h(X), \quad \overline{g} = \varphi, \quad \overline{h} = \psi,$$

$\deg(g) = \deg(\varphi)$, *and* $\deg(h) = \deg(\psi)$.

Corollary 3.5.21. Let $f(X) \in \mathcal{O}_K[X]$ be monic, with image $\overline{f}(X) \in \overline{K}[X]$. If $\overline{\alpha} \in \overline{K}$ is a simple root of $\overline{f}(X)$, then there exists $\alpha \in \mathcal{O}_K$ such that $f(\alpha) = 0$ and $\overline{\alpha}$ is the image of α under the residue map.

Hensel's Lemma allows us to give a description of all possible unramified extensions of a non-Archimedean local field, as in [Lor07, p.70].

Theorem 3.5.22. *Let K be a non-Archimedean local field, and let C be an algebraic closure of K. Then, $\overline{C}$ is an algebraic closure of the residue field $\overline{K}$ of K. The map*

$$E \mapsto \overline{E}$$

sending the non-Archimedean local field E to its residue field $\overline{E}$ is an inclusion preserving bijection between finite unramified extensions of K in $\overline{C}$ and finite separable extensions of $\overline{K}$ in $\overline{C}$. Every unramified extension of K in C is separable. The unramified extension E/K is Galois if and only if $\overline{E}/\overline{K}$ is Galois, and there is a natural isomorphism from the Galois group of E/K to the Galois group of $\overline{E}/\overline{K}$.

Proof. We give a sketch of the proof of [Lor07]. We define the inverse of the map of the theorem. Given an extension of non-Archimedean local fields, we know from Proposition 3.3.4 that the corresponding extension of residue fields is a finite extension

of finite fields, giving us the appropriate domain. Any such extension of finite fields $\overline{E}/\overline{K}$ is separable and is generated by a single $\overline{\alpha_E} \in \overline{C}$, with separable minimal polynomial $\overline{f} \in \overline{C}[X]$. By the relationship between the residue field and the valuation ring $\mathcal{O}_K$ of K, we can choose a monic $f(X) \in \mathcal{O}_K[X]$ which reduces to $\overline{f}$ over the residue field. By Hensel's Lemma 3.5.21, we can lift $\overline{\alpha_E}$ to a root α_E of $f(X)$ in C. This map

$$\overline{K}(\overline{\alpha_E}) \mapsto E := K(\alpha_E)$$

is the inverse we need. Since the valuation ring $\mathcal{O}_E$ of E is integrally closed, we have that α_E is contained in $\mathcal{O}_E$ and therefore that the residue field of E is contained in $\overline{K}(\overline{\alpha_E})$. By a clever inequality, we get all at once that f is irreducible, the residue field of E is equal to $\overline{K}(\overline{\alpha_E})$, and by further applying Proposition 3.3.4, we get that $K(\alpha_E)$ is unramified over K. Then, we get that $f(X)$ is a separable polynomial by applying Hensel's Lemma 3.5.21 to the roots of the polynomial $\overline{f}(X)$. So, E is separable over K. The statements about the Galois group are proved by manipulating roots in the residue field and applying Hensel's Lemma. See [Lor07] for details. ∎

Since the inverse map gives us the unique unramified extension of the base field by lifting a generator of the extension of (finite) residue fields, we can further describe unramified extensions of a non-Archimedean local field K by using our knowledge of finite degree extensions of finite fields. Indeed, it is a classic fact of field theory that finite extensions of finite fields are generated by the primitive roots of unity, and it will suffice to lift all these primitive roots of unity to our non-Archimedean algebraic closure to obtain all unramified extensions of K.

Corollary 3.5.23. Suppose the residue field of K has q elements. Under the correspondence of the Theorem, the unique unramified extension of K of degree n is obtained by adjoining to K a primitive $q^n - 1^{st}$ root of unity. Moreover, the Galois group of E/K is cyclic of order n, with canonical generator φ, where φ corresponds to the Frobenius automorphism of $\overline{E}/\overline{K}$ under the assignment $\tau \mapsto \overline{\tau}$.

To finish our discussion of p-adic fields, we give some concrete descriptions of the unramified extensions of $\mathbb{Q}_2$ in terms of familiar polynomials over the rationals and finite fields. We have seen in Corollary 3.5.23 that the unique unramified extension of degree n of $\mathbb{Q}_2$ is obtained by adjoining a primitive $2^n - 1^{st}$ root of unity to $\mathbb{Q}_2$. Set $k = 2^n - 1$. We denote by $\Phi_k(X)$ the minimal polynomial of a primitive k^{th} root of

unity over $\mathbb{Q}$. The basis of our approach is that $\mathbb{Q}$ is a subfield of $\mathbb{Q}_2$, so any primitive k^{th} root of unity ζ_k in an algebraic closure C of $\mathbb{Q}_2$ satisfies $\Phi_k(X) \in \mathbb{Q}[X]$. So, we can expect the minimal polynomial of ζ_k over $\mathbb{Q}_2[X]$ to be a divisor of $\Phi_k(X)$. We now describe the generator of the unique unramified quadratic extension of $\mathbb{Q}_2$, and give a basis for the corresponding extension of discrete valuation rings.

Example 3.5.24. Consider the case $n = 2$, so that $k = 3$. We claim that the third cyclotomic polynomial $\Phi_3(X) = X^2 + X + 1$ is irreducible over $\mathbb{Q}_2[X]$. We have that $\Phi_3(X)$ is irreducible in $\mathbb{Z}_2[X]$, as we can apply the mod $\mathfrak{p}$ irreducibility test to the reduction $\overline{X}^2 + \overline{X} + 1$ of $\Phi(X)$ in $\mathbb{F}_2[X]$, and since $\mathbb{Q}_2$ is the fraction field of the unique factorization domain $\mathbb{Z}_2$, we have that $\Phi(X)$ is irreducible as an element of $\mathbb{Q}_2[X]$ by the Corollary to Gauss' Lemma. Therefore, $\Phi_3(X)$ is irreducible over $\mathbb{Q}_2[X]$, so we have that the unique unramified extension of degree 2 of $\mathbb{Q}_2$ is given by $\mathbb{Q}_2(\zeta_3) \cong \mathbb{Q}_2[X]/\langle \Phi_3(X) \rangle$. Write $E = \mathbb{Q}_2(\zeta_3)$. If σ is the non-trivial Galois automorphism of $E/\mathbb{Q}_2$ we must have that $\sigma(\zeta_3) = \zeta_3^2$, the other distinct root of $\Phi_3(X)$. We can also say something about the valuation ring $\mathcal{O}_E$ of E. By Lemma 3.5.17, there is a unique extension of the 2-adic valuation on $\mathbb{Q}_2$ to E, and as a result of Proposition 3.4.14 and the ramification index formula of Theorem 3.3.4, the integral closure of $\mathbb{Z}_2$ in E must coincide with the valuation ring $\mathcal{O}_E$ of E. By looking at its minimal polynomial, we see that ζ_3 is integral over $\mathbb{Z}_2$. So, ζ_3 is an element of $\mathcal{O}_E$. Seeing as $E/\mathbb{Q}_2$ is unramified, Theorem 3.3.4 says that the extension of residue fields has degree 2, and since ζ_3 is a unit of $\mathcal{O}_E$ not contained in A, this means that the images of 1 and ζ_3 freely generate the residue field of $\mathcal{O}_E$ as an $\mathcal{O}_K/\mathfrak{p}\mathcal{O}_K$-vector space. By Nakayama's Lemma 2.1.27, $\{1, \zeta_3\}$ is an $\mathcal{O}_K$-basis of $\mathcal{O}_E$. Explicitly, the residue field of E is $\mathbb{F}_4$.

Unfortunately, $\Phi_k(X)$ is often a reducible polynomial over $\mathbb{Q}_2[X]$, as might be suggested by the fact that every unramified extension of degree n of $\mathbb{Q}_2$ is generated by a primitive $2^n - 1^{st}$ root of unity. For example, any root of the 7^{th} cyclotomic polynomial, which has degree 6 over $\mathbb{Q}$, will generate a degree 3 extension of $\mathbb{Q}_2$. This means that $\Phi_7(X)$ has proper irreducible factors over $\mathbb{Q}_2[X]$. It is therefore natural to obtain a description of the irreducible factors of $\Phi_k(X)$ in $\mathbb{Q}_p[X]$. To begin, we need the following theorem from field theory.

Theorem 3.5.25. *Let p be prime and set $q = p^n$, for some $n \in \mathbb{N}$. Then $\mathbb{F}_q$ is a splitting field for $X^q - X$ over $\mathbb{F}_p$. Moreover, $X^q - X$ is the product of the monic irreducibles of $\mathbb{F}_p[X]$ whose degree divides n.*

We will first describe the degree of every irreducible factor of the cyclotomic polynomial $\Phi_k(X)$ in $\mathbb{F}_p[X]$. We will be able to do so when the prime p does not divide k. Suppose k is a divisor of $p^n - 1$. Since finite extensions of finite fields are separable and $\Phi_k(X)$ is a divisor of $X^{q-1}-1$, Φ_k has coprime and separable irreducible factors in $\mathbb{F}_p[X]$, by the description of $X^q - X$ given in Theorem 3.5.25. Therefore, by Hensel's Lemma 3.5.20, giving a factorization of $\Phi_k(X)$ over $\mathbb{F}_p[X]$ into irreducibles of $\mathbb{F}_p[X]$ is equivalent to giving such a factorization over $\mathbb{Q}_p[X]$. Moreover, for the case that $p = 2$ and $k = 2^n - 1$, i.e., when we want to obtain the unique unramified extension of $\mathbb{Q}_2$ of degree n afforded by Corollary 3.5.23, $p = 2$ is always coprime to $k = 2^n - 1$, so our analysis will cover all cyclotomic polynomials which arise when we adjoin any $2^n - 1^{st}$ root of unity to $\mathbb{Q}_2$. Recall that the degree of $\Phi_k(X) \in \mathbb{Q}[X]$ is equal to the value of the Euler totient $\phi(k)$, the number of integers smaller than k which are also coprime to k.

Lemma 3.5.26. Let $k \in \mathbb{Z}$, and suppose that the prime p does not divide k. The cyclotomic polynomial $\Phi_k(X) \in \mathbb{F}_p[X]$ splits into $\frac{\phi(k)}{d}$ irreducible factors of degree d in $\mathbb{F}_p[X]$, where d is the multiplicative order of p in the ring $\mathbb{Z}/k\mathbb{Z}$.

Proof. Let $f(X)$ be an irreducible factor of $\Phi_k(X)$ and let α be a root of $f(X)$. Suppose that $f(X)$ has degree d and consider the extension $E := \mathbb{F}_p(\alpha)$ of $K := \mathbb{F}_p$. We know that the Galois group G of E/K is cyclic of order d with canonical generator

$$\sigma \colon E \to E, \quad x \mapsto x^p.$$

Then, the k-fold composition of σ is given by $\sigma^k(x) = x^{p^k}$. Since p does not divide k, and in particular, p is coprime to k, the class of p is a unit in $\mathbb{Z}/k\mathbb{Z}$. This allows us to define the map

$$\psi \colon G \to (\mathbb{Z}/k\mathbb{Z})^\times, \quad \psi(\sigma^k) = [p]^k.$$

We have that ψ is a group homomorphism, since

$$\begin{aligned} \psi(\sigma^m \circ \sigma^n) &= \psi(\sigma^{m+n}) \\ &= [p]^{m+n} \\ &= [p]^m [p]^n \\ &= \psi(\sigma^m)\psi(\sigma^n) \end{aligned}$$

and as G is cyclic with generator σ, ψ is surjective onto the cyclic group generated by $\psi(\sigma) = [p]$. Since E/K is Galois, and E/K contains a root α of $f(X)$, the d distinct

roots of $f(X)$ belong to E and are of the form

$$\alpha, \alpha^p, \ldots, \alpha^{p^{d-1}}.$$

Suppose that σ^m is in $\ker(\psi)$, for some $1 \leq m < d$. Then $\psi(\sigma^m) = [p]^m = 1$, so $p^m = 1 + rk$ for some $r \in \mathbb{Z}$. This means that $\sigma^m(\alpha) = \alpha^{p^m} = \alpha^{1+rk}$. Since α is a primitive k^{th} root of unity, we have $\alpha^{1+rk} = \alpha$. But G acts freely on the d distinct roots of $f(X)$, so σ^m must be the unique element of the Galois group fixing α, i.e., σ^m must be the identity of G. Hence, $\ker(\psi)$ is trivial. This shows that G is isomorphic to the cyclic group generated by p in $(\mathbb{Z}/k\mathbb{Z})^\times$, so that $|G| = d$ is equal to the multiplicative order of p in $\mathbb{Z}/k\mathbb{Z}$.

Now, this argument shows that each of the n irreducible factors of $f(X)$ in $\mathbb{F}_p[X]$ has the same degree d, so that $nd = \deg(\Phi_k(X)) = \phi(k)$, whence the formula for the number of factors. ∎

Corollary 3.5.27. The cyclotomic polynomial $\Phi_k(X)$ is irreducible in $\mathbb{F}_p[X]$ if and only if the group of units of $\mathbb{Z}/k\mathbb{Z}$ is cyclic with generator p.

Proof. In the notation of the Lemma, we have that $\Phi_k(X)$ is irreducible if and only if it has precisely one irreducible factor. This happens if and only if $d = \phi(n) = |(\mathbb{Z}/k\mathbb{Z})^\times|$, and if and only if the cyclic subgroup generated by p contains every element of $(\mathbb{Z}/k\mathbb{Z})^\times$. ∎

3.6 Functions on an algebraic curve

We introduce a second family of examples of discrete valuation rings, which comes from algebraic geometry. An excellent reference for the definitions and results presented in this section is [Fu89].

Definition 3.6.1. Let $V \subseteq \mathbb{A}^n$ be a set of points. We say V is an algebraic set if $V = V(S)$ for a set of polynomials $S \subseteq K[X_1, \ldots, X_n]$. We say V is a hypersurface if we can take $S = \{f(X_1, \ldots, X_n)\}$ for a polynomial $f(X_1, \ldots, X_n) \in K[X_1, \ldots, X_n]$.

We will only consider hypersurfaces in the case $n = 2$. These have a special name.

Definition 3.6.2. Let K be any field, and let $f(X,Y)$ be a polynomial in $K[X,Y]$. An affine algebraic curve C defined over K, is the set of all zeroes

$$C = C_f(K) = \{(a,b) \in \mathbb{A}^2 \mid f(a,b) = 0\}$$

of $f(X,Y)$ contained in K, so that $C_f(K) = V(f) \subseteq \mathbb{A}^2$. When the choice of base field and defining polynomial is clear, we simply write C.

An algebraic curve over K is the set of zeroes of a single polynomial in $K[X,Y]$. The notation $C_f(K)$ suggests that we reserve the right to consider the algebraic curve over field extensions of K. Indeed, we will be interested in looking at the behaviour of a ring associated to C_f as the base field varies.

Example 3.6.3. To come up with examples of plane curves, the sky is the limit! However, just like with univariate polynomials, one must be careful before asserting the existence of points on the curve when K is not algebraically closed. The polynomial $f(X,Y) = X^2 + Y^2 + 1 \in \mathbb{Q}[X,Y]$ defines an algebraic curve C_f. Then, $C_f(\mathbb{R})$ is empty. However, $C_f(\mathbb{C})$ has points, since $f((i,0)) = i^2 + 1 = 0$.

Definition 3.6.4. Suppose $V \subseteq \mathbb{A}^k$ is an algebraic set. The ideal

$$I(V) := \{f \in K[X_1, \ldots, X_k] \mid f(X_1, \ldots, X_k) \text{ vanishes on all } \vec{x} \in V\}$$

is the vanishing ideal of V.

Definition 3.6.5. Let I be an ideal in the commutative ring R. The radical of I is the ideal

$$\sqrt{I} = \{x \in A \mid x^n \in I \text{ for some } n \geq 0\}.$$

We say that I is a radical ideal if $I = \sqrt{I}$.

It can be shown by elementary manipulations that $V(I(V)) = V$. On the other hand, Example 3.6.3 highlights that in general, $I(V(I)) \neq I$, as when $V(I)$ is empty, $I(V(I))$ is trivially $K[X,Y]$. Univariate polynomials split into linear factors over an algebraic closure. In our example, passing to the algebraic closure also provides solutions to the equation $X^2 + Y^2 = -1$ in $\mathbb{A}^2$. This is a special case of a general fact, the content of which is outlined in Hilbert's Nullstellensatz: over an algebraically closed field K, $\sqrt{J} = I(V(J))$ for any ideal $J \in K[X_1, \ldots, X_n]$. Thus, the I operation is a left inverse to the V operation on radical ideals, setting up a bijection between radical ideals of $K[X_1, \ldots, X_n]$ and algebraic sets in $\mathbb{A}^n$.

Definition 3.6.6. Let $V \subseteq \mathbb{A}^n$ be an algebraic set. We say that V is an affine variety, or that V is irreducible, if V cannot be written as the union of two proper algebraic sets.

Let K be a field, and consider 2-dimensional affine space $\mathbb{A}^2$ over K. It is a general fact that an affine variety V contained in $\mathbb{A}^2$ which has a point over K uniquely defines a prime ideal $I(V) := \{f(X,Y) \in K[X,Y] \mid f(a,b) = 0 \text{ for all } (a,b) \in V\}$. By taking quotients, we obtain the integral domain $\Gamma(V) := K[X,Y]/I(V)$ of V, which is known as the *coordinate ring* of the variety V. When K is integrally closed, it is a consequence of the Nullstellensatz that integral domains that are quotients of $K[X,Y]$ are in one-to-one correspondence with affine varieties in $\mathbb{A}^2$. The coordinate ring $\Gamma(V)$ then has a natural interpretation as the ring of polynomial functions defined on points of the variety. Indeed, two polynomials define the same polynomial on the variety precisely when their difference is 0 on the variety.

Over any field, the vanishing set of a maximal ideal of the form

$$\langle X_1 - a_1, \ldots, X_k - a_k \rangle$$

is the singleton $\{(a_1, \ldots, a_k)\}$, and over an algebraically closed field, one can deduce from the Nullstellensatz that every maximal ideal is of this form. While there is no direct correspondence over a non-algebraically closed field, this motivates us to try to interpret prime ideals as points, lines, or higher-dimensional objects in $\mathbb{A}^n$. This can be done by introducing the Zariski topology, which we will have to disregard for now. It is clear that all solutions of a polynomial equation over the base field have a representative in the spectrum of the coordinate ring, but, as we will see, there are prime ideals of codimension 0, points in our analogy, which come from solutions which only exist over algebraic extensions. So, studying a geometric object through its coordinate ring unveils the hidden structure which would only otherwise appear over the algebraic closure when directly studying solutions in affine space. In this way, the construction of the coordinate ring gives us a richer view of the geometry of a curve. We mention that, formally, the *codimension* of the ideal $\mathfrak{p}$, an algebraic property, is the measure of the geometric dimension of the object corresponding to that obtained by considering all prime ideals which contain $\mathfrak{p}$. By the Correspondence Theorem, the codimension of an ideal is precisely the Krull dimension of the relevant coordinate ring. So, we should expect the coordinate ring of a curve to have Krull dimension 1.

Definition 3.6.7. Let $C_f(K)$ be an algebraic curve. We say the prime ideal $\mathfrak{p} \in \mathrm{Spec}\,(K[X,Y])$ is a geometric point of C if $\langle f(X,Y)\rangle$ is contained in $\mathfrak{p}$. We say this point is closed if $\mathfrak{p}$ is maximal. We say the geometric point $\mathfrak{p}$ of C is a rational point of C if there exists $(a,b) \in \mathbb{A}^2$ such that $\mathfrak{p} = \langle X-a, Y-b\rangle$. In this case, $f(a,b) = 0$.

In other words, the set of geometric points of an algebraic curve is the prime spectrum of its coordinate ring, and we consider these notions interchangeably.

Definition 3.6.8. Let V be a variety in $\mathbb{A}^2$, let $\Gamma(V) = K[X,Y]/I(V)$ be the coordinate ring of V, and let $\mathfrak{p} \in \mathrm{Spec}\,(K[X,Y])$ be such that $I(V) \subseteq \mathfrak{p}$. The localization

$$\mathcal{O}_{V,\mathfrak{p}} := \Gamma(V)_{\mathfrak{p}}$$

is known as the local ring of V at (the point) $\mathfrak{p}$.

We should view this local ring as that of all polynomial functions on V where fractions are allowed, so long as the denominator does not vanish at $\mathfrak{p}$. By Theorem 3.2.1, the ideals of $\mathcal{O}_{V,\mathfrak{p}}$ are the ideals of $\Gamma(V)$ which are contained in $\mathfrak{p}$. In this way, the local ring $\mathcal{O}_{V,\mathfrak{p}}$ gives us algebraic data on the geometric subobjects of V which contain our point. We have therefore touched on one of the main principles underpinning algebraic geometry: the study of geometric objects through the study of functions on their subobjects. We will not formally introduce the corresponding definitions, but one can then consider elements of $K[X,Y]$ (or of a coordinate algebra) as functions at every point $\mathfrak{p} \in \mathrm{Spec}\,(K[X,Y])$ at once, by imposing functoriality relations on the lattice of local rings. This yields the notion of an affine scheme.

We now investigate the algebraic properties of the coordinate ring $\Gamma(V)$ and its localizations in order to obtain new examples of discrete valuation rings. By the Hilbert Basis Theorem 2.2.7 and Lemma 2.2.4, we know that $\Gamma(V)$ is Noetherian, and by Lemma 3.2.3, we know that every localization of $\Gamma(V)$ is Noetherian. To find examples of discrete valuation rings among the localizations of $\Gamma(V)$, we will need to find a localization that is integrally closed and of dimension 1. We will see in the case that $V = C$ is an algebraic curve and $\mathfrak{p}$ is a non-singular point on the curve, the localization $\Gamma(C)_{\mathfrak{p}}$ is integrally closed. But first, we show that $\Gamma(C)$ has dimension 1, which will be an easy consequence of the following result, well-known to specialists.

Lemma 3.6.9. Let R be a principal ideal domain, and let $\mathfrak{p}$ be a non-zero prime ideal of $R[Y]$. Then $\mathfrak{p}$ is either a principal ideal generated by an irreducible element

$g \in R[Y]$ or $\mathfrak{p}$ is equal to $\langle p, f(Y)\rangle$, where p is prime in R and $f(Y)$ is irreducible in $R/\langle p\rangle[Y]$.

Proof. Put $A = R[Y]$, let $\mathfrak{p}$ be a non-zero prime ideal of A, and let pR be an arbitrary prime ideal of R. Suppose first that $\mathfrak{p} \subseteq pA$, so that $\mathfrak{p} \cap R \subseteq pR$. Seeing as $\mathfrak{p} \cap R$ is prime and R has dimension 1, then $\mathfrak{p} \cap R = pR$, so that $\mathfrak{p}$ is the principal ideal pA. Suppose then that $\mathfrak{p}$ is not a subset of pA. We consider the image of $\mathfrak{p}$ in the principal ideal domain $A/pA \cong R/\langle p\rangle[Y]$. Since $\mathfrak{p}$ is prime, $\mathfrak{p}$ is generated modulo pA by an irreducible $\tilde{f}(Y) \in R/\langle p\rangle[Y]$. Lift $\tilde{f}(Y)$ to an element $f(Y)$ of $\mathfrak{p}$. Then, if $g(Y) \in \mathfrak{p}$, we have that $g(Y) - f(Y)q(Y) \in pA$ for some $q(Y) \in A$. If $\mathfrak{p}$ intersects pA trivially, it follows that $\mathfrak{p} = f(Y)A$. Moreover, $f(y)$ satisfies $\deg(f(Y)) = \deg(\tilde{f}(Y))$, as the obvious lift $g(y)$ of $\tilde{f}(Y)$ belongs to $\mathfrak{p}$, and $f(y) - g(y) \in pA \cap \mathfrak{p} = 0$. Applying mod-$p$ irreducibility, we have that $f(Y)$ is irreducible. If instead $\mathfrak{p}$ meets pA non-trivially, $\mathfrak{p}$ is of the form $\langle f(Y), p\rangle$, where $f(Y)$ is irreducible in $R/\langle p\rangle[Y]$ and p is prime in R. ∎

Corollary 3.6.10. The ring $K[X, Y]$ has Krull dimension 2.

Proof. Let $A = K[X, Y] \cong (K[X])[Y]$. We first show the length of a chain of prime ideals of A is at most 2. Let $\cdots \subseteq \mathfrak{p}_k \subseteq \cdots \subseteq \mathfrak{p}_0$ be a chain of prime ideals contained in $\mathfrak{p}_0$. By the Lemma, $\mathfrak{p}_k$ is either generated by an irreducible $f(X, Y) \in K[X, Y]$ or $\mathfrak{p}_k$ is equal to $\langle p(X), g(X, Y)\rangle$ for $p(X)$ prime in $K[X, Y]$ and $g(X, Y)$ irreducible in $(K[X]/\langle p(X)\rangle)[Y]$. Suppose first that $\mathfrak{p}_k = \langle p(X), g(X, Y)\rangle$ is of the second type. We claim that any such ideal is maximal. Indeed, $p(X)$ is prime and therefore irreducible in $K[X]$, so that

$$K[X, Y]/\langle p(X)\rangle \cong (K[X]/\langle p(X)\rangle)[Y]$$

is a polynomial ring over a field. Put $K_p = K[X]/\langle p(X)\rangle$. Since the image of $g(X, Y)$ is prime and therefore irreducible in $K_p[Y]$, the quotient $K_p[Y]/\langle [g(X, Y)]\rangle$ is a field. So, $K[X, Y]/\mathfrak{p}_k$ is a field, which shows that $\mathfrak{p}_k$ is maximal. If $\mathfrak{p}_k = \langle f(X, Y)\rangle$ is of the first type, then $\mathfrak{p}_k$ cannot be properly contained in a principal ideal by the irreducibility of $f(X, Y)$. So, $\mathfrak{p}_1$ can only be properly contained in a prime ideal of the second type, which is maximal. This shows all chains of prime ideals have length at most 2. Since

$$\langle 0\rangle \subsetneq \langle X - a\rangle \subsetneq \langle X - a, Y - b\rangle \subsetneq A$$

is a chain of prime ideals of length 2, the Krull dimension of A is 2. ∎

Remark 3.6.11. Note that the proof of the Corollary sneakily makes use of the fact that $K[X,Y]$ is a univariate polynomial ring over the principal ideal domain $R = K[X]$. Since polynomial rings in more than one indeterminate have non-principal ideals, we cannot hope for an easy inductive result as in the proof of the Hilbert Basis Theorem. However, it is still true that $\dim(K[X_1,\ldots,X_r]) = r$ for any $r \geq 2$, but the proof requires different machinery. An argument using Noether's Normalization Lemma is given in [Eis95, §8].

As expected from our discussion of dimension and geometry, we get the following result.

Corollary 3.6.12. Let I be a non-maximal and non-zero prime ideal of $K[X,Y]$. Then $K[X,Y]/I$ has Krull dimension 1.

Proof. Since I is a non-maximal prime ideal in a 2-dimensional ring, the Lemma and the correspondence theorem for quotient rings tell us that chains of ideals containing I in $K[X,Y]/I$ have length at most 1. Since I is non-maximal in $K[X,Y]$, there is at least one chain of length 1 in the quotient. ∎

Corollary 3.6.13. Let $\mathcal{O}_{C,\mathfrak{p}}$ be the local ring of the closed point $\mathfrak{p}$ on an algebraic curve C. Then $\mathcal{O}_{C,\mathfrak{p}}$ has dimension 1.

Proof. Apply Theorem 3.2.1 to the coordinate ring $\Gamma(C)$ of C and the localization $\mathcal{O}_{C,\mathfrak{p}} = (\Gamma(C))_{\mathfrak{p}}$. ∎

Example 3.6.14. After discussing dimension in $K[X,Y]$, is worth returning to the definition of a rational point, which introduces a new subtlety. Since the coordinate ring of a curve has Krull dimension 1, the geometric points of C are precisely the maximal ideals of $K[X,Y]$ containing $\langle f(X,Y)\rangle$, and the so-called generic point $\langle f(X,Y)\rangle$, corresponding to the zero ideal in the coordinate ring. When K is algebraically closed, every non-zero geometric point of C is a rational point, as discussed when we introduced the Nullstellensatz. On the other hand, if K is not algebraically closed, there may be maximal ideals of $K[X,Y]$ which are not of the form $\langle X-a, Y-b\rangle$. In fact, it is easy to find an example of a curve with no rational points but with plenty of geometric points. For example, consider the ideal $\langle X, Y^2+1\rangle$ in $\mathbb{R}[X,Y]$. This ideal properly contains $\langle X^2+Y^2+1\rangle$, and we have seen that the polynomial X^2+Y^2+1 has

no solutions in $\mathbb{A}^2_{\mathbb{R}}$. On the other hand, the ideal $\langle X, Y^2+1\rangle$ is maximal in $\mathbb{R}[X,Y]$, as we can apply Lemma 3.6.9 with the ring $(\mathbb{R}[X]/\langle X\rangle)[Y]$. This shows that $\langle X, Y^2+1\rangle$ is a non-rational geometric point of the curve defined by X^2+Y^2+1.

We now introduce an abstract property of a geometric point on an algebraic curve, and show it is the correct one to impose so that the corresponding local ring is a discrete valuation ring.

Definition 3.6.15. Let C be an algebraic curve, and let $\mathfrak{p}$ be a geometric point on C. We say $\mathfrak{p}$ is a non-singular point if the local ring $\mathcal{O}_{C,\mathfrak{p}}$ satisfies $\dim(\mathcal{O}_{C,\mathfrak{p}}) = \dim_k(\mathfrak{p}/\mathfrak{p}^2)$. More generally, a Noetherian local ring A with maximal ideal $\mathfrak{m}$ and residue field K satisfying $\dim(A) = \dim_K(\mathfrak{m}/\mathfrak{m}^2)$ is called a regular local ring.

Lemma 3.6.16. The integral domain A is a regular local domain of dimension 1 if and only if A is a discrete valuation ring.

Proof. Since we have assumed that all other properties of Dedekind domains hold, it suffices to show A is integrally closed to conclude that A is a discrete valuation ring, by Lemma 3.4.11. To that end, we show that A is a unique factorization domain. By Nakayama's Lemma 2.1.27, the generator for $\mathfrak{m}/\mathfrak{m}^2$ lifts to a generator ϖ of the (finitely-generated) maximal ideal $\mathfrak{m}$ of A. Let $x \in A$ be non-zero. Then, $x = \varpi^k u$ for some $u \in A$ and $k \geq 0$. We can assume that this k is maximal. If, for every $k \geq 0$, we have that ϖ^k divides x, then $x \in \bigcap_{i=1}^{\infty} \mathfrak{m}^k$. By Krull's intersection theorem (see [Eis95, p.152]) we have that this intersection is 0. So any such x is 0. Hence, there is a maximal $k \in \mathbb{Z}$ such that $x = \varpi^k u$, where ϖ does not divide u, i.e., u is a unit. This establishes the existence of prime factorizations of non-units. The integral domain A has dimension 1 and has a unique maximal ideal. So, any prime element of A generates the maximal ideal and must be equal to ϖu, for $u \in A^\times$. Therefore, the only prime element of A up to units is ϖ. It follows that the prime factorization $x = \varpi^k u$ is unique. So, A is a unique factorization domain. Conversely, a discrete valuation ring is a local principal ideal domain, so it is a regular local domain of dimension 1. ∎

Corollary 3.6.17. Let $\mathfrak{p}$ be a non-singular point on the algebraic curve C. Then $\mathcal{O}_{C,\mathfrak{p}}$ is a discrete valuation ring.

Proof. By Corollary 3.6.12 and the discussion preceding Lemma 3.6.9, we know that $\mathcal{O}_{C,\mathfrak{p}} = (K[X,Y]/\langle f(X,Y)\rangle)_{\mathfrak{p}}$ is a Noetherian local domain of dimension 1. Further imposing that $\mathfrak{p}$ is non-singular (so that the local ring $\mathcal{O}_{C,\mathfrak{p}}$ is regular) allows us to use the Lemma. ∎

We give an easily applicable criterion for non-singularity when our geometric point is also rational.

Lemma 3.6.18. Suppose $\mathfrak{p} = \langle X-a, Y-b\rangle$ is a rational point on the irreducible curve $C = V(f)$. The polynomial $f(X,Y)$ satisfies the equation

$$f(X,Y) = f_X(a,b)(X-a) + f_Y(a,b)(Y-b) \pmod{\mathfrak{p}^2},$$

If the partial derivatives $f_X(a,b)$ and $f_Y(a,b)$ are not both zero, then $\mathfrak{p}$ is non-singular, and in this case, $\mathcal{O}_{C,\mathfrak{p}}$ is a discrete valuation ring.

Proof. Suppose without loss of generality that $f_Y(a,b)$ is not zero. Since $K[X,Y] \cong K[X-a, Y-b]$, we can write

$$f(X,Y) = f(a,b) + p(X-a) + q(Y-b) + r(X,Y),$$

where $r(X,Y) \in \mathfrak{p}^2$. Since (a,b) is a rational point, $f(a,b) = 0$, and since

$$r(X,Y) \in \mathfrak{p}^2 = \langle (X-a)^2, (Y-b)^2, (X-a)(Y-b)\rangle,$$

the partial derivatives $r_X(a,b)$ and $r_Y(a,b)$ are equal to 0. This, together with our expansion of $f(X,Y)$ in $X-a$ and $Y-b$ implies that $p = f_X(a,b)$ and $q = f_Y(a,b)$. So, we get that

$$f(X,Y) = f_X(a,b)(X-a) + f_Y(a,b)(Y-b) + \mathfrak{p}^2,$$

reduces to $0 = f_X(a,b)(X-a)+f_Y(a,b)(Y-b)+\mathfrak{p}^2$ in the coordinate ring $K[X,Y]/\langle f(X,Y)\rangle$ of C, and by the assumption that $f_Y(a,b)$ is non-zero, we have

$$Y - b = \frac{-f_X(a,b)}{f_Y(a,b)}(X-a).$$

This shows that $X-a$ is a generator of $\mathfrak{p}/\mathfrak{p}^2$. Upon localizing to $\mathcal{O}_{C,\mathfrak{p}}$, Nakayama's Lemma 2.1.27 says that $X-a$ generates the maximal ideal $\mathfrak{p}$ in $\mathcal{O}_{C,\mathfrak{p}}$. So, the dimension of $\mathfrak{p}/\mathfrak{p}^2$ is 1. On the other hand, $\mathcal{O}_{C,\mathfrak{p}}$ evidently has a non-zero prime ideal $\mathfrak{p}$, and therefore $\mathcal{O}_{C,\mathfrak{p}}$ is 1-dimensional by Corollaries 3.6.12 and 3.2.2. This means that $\dim(\mathcal{O}_{C,\mathfrak{p}}) = 1 = \dim(\mathfrak{p}/\mathfrak{p}^2)$ and $\mathfrak{p}$ is non-singular. ∎

Corollary 3.6.19. If $f_X(a,b) \neq 0$, then $Y - b$ generates $\mathfrak{p}$ in $\mathcal{O}_{C,\mathfrak{p}}$. If $f_Y(a,b) \neq 0$, then $X - a$ generates $\mathfrak{p}$ in $\mathcal{O}_{C,\mathfrak{p}}$.

Proof. Without loss of generality, assume that $f_Y(a,b) \neq 0$, as otherwise, we can switch the roles of $X-a$ and $Y-b$. As in the proof of the Lemma, we solve for $Y-b$ in terms of $X-a$, and use Nakayama's Lemma to obtain a generator for $\mathfrak{p}$. ∎

Example 3.6.20. Let K be a field of characteristic not 2, and consider the algebraic curve C defined over K by the polynomial $f(X,Y) = X^2+Y^2-1$. Then, $f_X(a,b) = 2a$ and $f_Y(a,b) = 2b$. Let us set $K = \mathbb{R}$. If $(a,b) \in \mathbb{R}^2$ satisfy $a^2 + b^2 = 1$, then a and b are not both 0. So, every rational point on C is non-singular by Lemma 3.6.18. This means that when $\mathfrak{p}$ is rational, the local ring $\mathcal{O}_{\mathfrak{p}}(C)$ is a discrete valuation ring, whose maximal ideal is generated by an element for which we can solve as in Corollary 3.6.19. We can describe the discrete valuation ring $\mathcal{O}_{\mathfrak{p}}(C)$ of the curve X^2+Y^2-1 at the north pole $\mathfrak{p} = \langle X, Y-1\rangle$. Formally, this local ring is equal to $\mathcal{O}_{C,\mathfrak{p}} = \left(\mathbb{R}[X,Y]/\langle X^2+Y^2-1\rangle\right)_{\mathfrak{p}}$. Since localization commutes with taking quotients, we can think of the invertible elements of $\mathcal{O}_{C,\mathfrak{p}}$ as the equivalence classes $[g(X,Y)]$ such that $g(0,1) \neq 0$. In general, the localization of a ring A away from a prime ideal $\mathfrak{p}$ has unique maximal ideal generated by the image of $\mathfrak{p}$. We cannot expect $\mathfrak{p}$ to be principal in the coordinate ring $\mathbb{R}[X,Y]/\langle X^2+Y^2-1\rangle$, but in the case of our local ring $\mathcal{O}_{C,\mathfrak{p}}$ at the point $(0,1)$, we know by Lemma 3.6.18 that the maximal ideal $\mathfrak{p}$ generated by the inclusion of $\mathfrak{p} = \langle X, Y-1\rangle$ is a principal ideal. In particular, $f_Y(0,1) = 2$, so Corollary 3.6.19 says that the maximal ideal $\mathfrak{p}$ of $\mathcal{O}_{C,\mathfrak{p}}$ is generated by the image of the indeterminate X in $\mathcal{O}_{C,\mathfrak{p}}$.

We now consider what happens to the local ring $\mathcal{O}_{C,\mathfrak{p}}$ of the Example when we change scalars to the algebraic closure of $\mathbb{R}$. In this way, we obtain an example of a separable quadratic extension of discrete valuation rings.

Example 3.6.21. Let $f(X,Y) = X^2+Y^2-1 \in \mathbb{C}[X,Y]$, and put $B = \mathbb{C}[X,Y]/\langle X^2+Y^2-1\rangle$. Write $A_{\mathfrak{p}}$ for the local ring of $C_f(\mathbb{R})$ at $(0,1)$. Then, as in Example 3.6.20, the local ring of the curve $C_f(\mathbb{C})$ at the non-singular point $\langle X, Y-1\rangle$ is the localization $B_{\mathfrak{P}}$ of B away from the prime ideal $\mathfrak{P} = \langle X, Y-1\rangle$. As before, $B_{\mathfrak{P}}$ has maximal ideal generated by the image of the indeterminate $X \in \mathbb{C}[X,Y]$. Recall from Example 2.5.6 that $\mathbb{C}$ is a separable $\mathbb{R}$-algebra. Since $B_{\mathfrak{P}} \cong \mathbb{C}\otimes_{\mathbb{R}}\left(\mathbb{R}[X,Y]/\langle X^2+Y^2+1\rangle\right)_{\mathfrak{p}}$, Lemma 2.5.18 says that $B_{\mathfrak{P}}$ is a separable $A_{\mathfrak{p}}$-algebra. It should not be surprising to

us that $\mathfrak{p}$ and $\mathfrak{P}$ are both generated by the same indeterminate X, as this is precisely the content of Corollary 2.5.16.

Example 3.6.22. Let $B_{\mathfrak{P}}$ be the separable $A_{\mathfrak{p}}$-algebra of Example 3.6.21. We saw in the previous example that these rings are discrete valuation rings and therefore local Dedekind domains. We would like to compute their residue fields. Abusing notation, let $\mathfrak{P}$ denote the maximal ideal of $B_{\mathfrak{P}}$ and let $\mathfrak{p}$ denote that of $A_{\mathfrak{p}}$. Since we are in characteristic 0 and the extension $\operatorname{Frac}(B_{\mathfrak{P}}) / \operatorname{Frac}(A_{\mathfrak{p}}) = \mathbb{C}(X,Y)/\mathbb{R}(X,Y)$ has degree 2, Lemma 3.3.4 says that $2 = ef$, where e is ramification index of $\mathfrak{p}$ in $\mathfrak{P}$ and f is the degree of the extension $[B_{\mathfrak{P}}/\mathfrak{P}B_{\mathfrak{P}} \colon A_{\mathfrak{p}}/\mathfrak{p}A_{\mathfrak{p}}]$. In the previous example, we saw that $e = 1$. Therefore, we should expect $B_{\mathfrak{P}}/\mathfrak{P}B_{\mathfrak{P}}$ to be a quadratic extension of $A_{\mathfrak{p}}/\mathfrak{p}A_{\mathfrak{p}}$. Now, the relations $X = 0$ and $Y = 1$ coming from $\mathfrak{p} = \langle X, Y-1\rangle = \langle X\rangle$ in $A_{\mathfrak{p}}/\mathfrak{p}A_{\mathfrak{p}}$ make it clear that the evaluation homomorphism

$$A_{\mathfrak{p}}/\mathfrak{p}A_{\mathfrak{p}} \to \mathbb{R}$$

$$X \mapsto 0, Y \mapsto 1$$

is an isomorphism. So, $A_{\mathfrak{p}}/\mathfrak{p}A_{\mathfrak{p}} = \mathbb{R}$, and since $f = 2$, $B_{\mathfrak{P}}/\mathfrak{P}B_{\mathfrak{P}} \cong \mathbb{C}$. Note that the same analysis applied to the relations in $B_{\mathfrak{P}}/\mathfrak{P}B_{\mathfrak{P}}$ shows that the map $X \mapsto 0, Y \mapsto 1$ is an isomorphism of the residue field with $\mathbb{C}$, as predicted by the ramification theory.

Chapter 4

Lattices and isometries

In Chapter 4, we present the main theorem of the thesis.

Theorem 4.0.1. *Let $\mathcal{O}_E/\mathcal{O}_K$ be a quadratic separable extension of discrete valuation rings with non-trivial standard involution σ. Let L be a Hermitian lattice over $\mathcal{O}_E$. The unitary group of L is generated by symmetries. Any minimal factorization of an isometry τ of L in terms of symmetries has length at most $2\,\mathrm{rk}\,(L_\tau) - \dim_a(\tau)$.*

Sections 4.1, 4.2, and 4.3 introduce rings with involution and Hermitian lattices over these rings. In Section 4.4, we define a quadratic extension of discrete valuation rings and delve deeply into the properties of the Hermitian lattices obtained by restricting ourselves to this setting. In Section 4.5, we verify that a generalized version of the Gram-Schmidt algorithm holds for lattices over quadratic extensions of discrete valuation rings, so that lattices are the orthogonal sum of rank 1 and rank 2 sublattices. In Section 4.6, we define an invariant of a Hermitian lattice, the Jordan decomposition. In Section 4.7, we explore properties of lattices over separable extensions of DVRs. In particular, we prove that such lattices have an orthogonal basis. In Section 4.8, we define the group of isometries of a Hermitian lattice, the unitary group, and in Section 4.9, we present our main result. As mentioned in Chapter 1, most proofs in this Chapter are verifications that results which hold over a non-Archimedean local field of characteristic zero also hold over a quadratic extension of discrete valuation rings. The relevant references are cited when extending a previously known result. However, our presentation of the anisotropic dimension of an

isometry and the part of the proof of Theorem 4.9.1 which describes the length of the factorization of a strictly diagonalizable isometry is entirely new.

4.1 Rings with involution

Following the presentation in [Knu91], we introduce the concept of a ring with involution, a broad generalization of the extension $\mathbb{C}/\mathbb{R}$ with its complex conjugation.

Definition 4.1.1. Let A be a ring. A function $\sigma\colon A \to A$ is an involution if for all $a, b \in A$, we have

1. $\sigma(a+b) = \sigma(a) + \sigma(b)$;
2. $\sigma(ab) = \sigma(b)\sigma(a)$;
3. $\sigma(1) = 1$;
4. $\sigma^2(a) = a$,

and additionally, that σ is a bijection. In other words, σ is an anti-automorphism of order at most 2.

Example 4.1.2. If A is a commutative ring, then any automorphism of order at most 2 is an involution. In particular, the identity is an involution.

Example 4.1.3. Let E/K be a Galois field extension of even degree. The automorphism group of E/K contains an automorphism σ of order 2. Then, σ is an involution.

Example 4.1.4. Let A be a ring. The matrix transpose $M \mapsto M^T$ on $M_n(A)$ is an involution.

Example 4.1.5. Let E/K be a quadratic Galois extension with non-trivial automorphism σ. We define a map

$$M \mapsto M^\sigma\colon M_n(E) \to M_n(E)$$

defined by

$$([m_{ij}])^\sigma = [\sigma(m_{ij})].$$

This map is an involution if and only if $n = 1$. Note that in the familiar setting $\mathbb{C}/\mathbb{R}$, the operation $M \mapsto M^\sigma$ is precisely complex conjugation of the entries of M.

On the other hand, we can generalize the conjugate transpose map of $M_n(\mathbb{C})$ to obtain an involution on the matrix algebra.

Example 4.1.6. Let A be a commutative ring with involution σ. The conjugate transpose map

$$M \mapsto M^* \colon M_n(A) \to M_n(A)$$

defined by

$$([m_{ij}])^* = [\sigma(m_{ji})]$$

is an involution on the matrix ring $M_n(A)$.

Recall that for any ring A we have the opposite ring A^{op}, where A^{op} has the same underlying additive group as A, and where A^{op} is equipped with the multiplication operation defined by

$$a^{op} \cdot b^{op} = (ba)^{op}.$$

Example 4.1.7. Let A be a ring. Consider the direct product $A \times A^{op}$ of A with its opposite ring A^{op}. Then $A \times A^{op}$ is equipped with the switch involution $(a, b^{op}) \mapsto (b, a^{op})$.

Definition 4.1.8. Let R be a commutative ring with involution σ, and let A be an R-algebra. We say that A is an R-algebra with involution if A is a ring with involution σ_A such that σ_A restricts to σ on R.

When we say that A is an R-algebra with involution (without mention of an involution on R), it will be implied that the involution of A restricts to an involution of R.

Definition 4.1.9. Let R be a ring with involution, and let A be an R-algebra with involution σ. We define the trace $\mathrm{Tr}_\sigma \colon A \to A$ by

$$\mathrm{Tr}_\sigma(a) = a + \sigma(a)$$

and the norm $\mathrm{N}_\sigma \colon A \to A$ by

$$\mathrm{N}_\sigma(a) = a\sigma(a).$$

Definition 4.1.10. Let A be an R-algebra with involution σ. We say that σ is a standard involution on A if $\sigma(r) = r$ for all $r \in R$ and if both Tr_σ and N_σ take values in R.

Example 4.1.11. If E/K is a quadratic Galois extension with non-trivial automorphism σ, we compare the formulas for N_σ and Tr_σ with those from Proposition 2.4.4 and observe that N_σ and Tr_σ are the usual norm and trace for a field extension. Considering E as a K-algebra, σ is a standard involution on E. This is because K is the fixed field of σ and the norm and trace have image in K.

Example 4.1.12. Consider $E = K \times K$ as an étale K-algebra. Since $K = K^{op}$, we can equip E with the involution $\sigma\colon E \to E$ given by $(a, b) \mapsto (b, a)$ as in Example 4.1.7. We have that K canonically embeds in E under the diagonal map. Since $(a, b)(b, a) = (ab, ab)$ and $(a, b) + (b, a) = (a + b, a + b)$, the norm and trace of E take values in E. Moreover, $\sigma((a, a)) = (a, a)$, so that σ fixes K. Therefore, σ is a standard involution on $K \times K$.

Indeed, non-trivial Galois automorphisms and the hyperbolic involution are the only two possible choices for standard involutions on a quadratic algebra over a field. However, it is not hard to show that uniqueness holds for a broader class of R-algebras. We present the following result and its proof, as in [Knu91, Prop. 1.3.4, p.3].

Proposition 4.1.13. Let A be an R-algebra which is finitely generated, projective, and faithful as an R-module. Then there exists at most one standard involution on A.

Proof. Since equality of module homomorphisms is a local property, it is sufficient to examine the case when R is local. Let σ_1, σ_2 be two standard involutions on A. Since A is free and finitely generated, we have by Lemma 2.1.23 that A has an R-basis $\{1, e_2, \ldots, e_n\}$. If σ is any standard involution, any $x \in A$ satisfies

$$x^2 + \mathrm{N}_\sigma(x) = x^2 + x\sigma(x) = x(x + \sigma(x)) = \mathrm{Tr}_\sigma(x)x.$$

From this, we must have

$$-\mathrm{N}_{\sigma_1}(e_i) \cdot 1 + \mathrm{Tr}_{\sigma_1}(e_i) \cdot e_i = e_i^2 = -\mathrm{N}_{\sigma_2}(e_i) \cdot 1 + \mathrm{Tr}_{\sigma_2}(e_i) \cdot e_i$$

for every $2 \leq i \leq n$. This means that

$$\mathrm{Tr}_{\sigma_1}(e_i) = \mathrm{Tr}_{\sigma_2}(e_i)$$

since 1 and the e_i are basis vectors for A over R, and the traces and norms of σ_i lie in R. As the involutions are standard, we must have that $\sigma_1(1) = 1 = \sigma_2(1)$. By linearity, $x + \sigma_1(x) = \mathrm{Tr}_{\sigma_1}(x) = \mathrm{Tr}_{\sigma_2(x)} = x + \sigma_2(x)$ for all $x \in A$. Cancelling, this means $\sigma_1(x) = \sigma_2(x)$ for all $x \in A$. ∎

Corollary 4.1.14. Let E be a quadratic étale K-algebra, where K is a field. If E is a field, then the non-trivial automorphism σ of E/K is the unique standard involution on E. If E is of the form $K \times K$, then the unique standard involution on E is given by the map $(a, b) \mapsto (b, a)$ of Example 4.1.7.

Proof. An étale K-algebra E is projective, finitely generated, and faithful as a K-module. We have seen that the above involutions are indeed standard involutions, so that by the Proposition, these are the only ones. ∎

Definition 4.1.15. We say that an R-algebra S is a quadratic R-algebra if S is a projective R-module of constant rank 2.

We give some properties of quadratic algebras over a local ring, as in [Knu91, $p.4$]. In particular, this gives an explicit formula for the standard involution on a quadratic algebra over a local ring.

Example 4.1.16. Let R be a local ring. Then, a quadratic R-algebra S is free of rank 2, as a finitely generated projective module over a local ring is free. By Lemma 2.1.23, R is a direct summand of S, so S has an R-basis $\{1, z\}$ for some $z \in S$. Write $z^2 = a + bz$ for some $a, b \in R$. Then, we endow S with an R-linear involution by extending

$$\sigma \colon S \to S$$

$$\sigma(z) = b - z$$

to all of S. Then, $\sigma(r + sz) = r + sb - sz$ for $x = r + sz$ so that

$$\mathrm{Tr}_\sigma(x) = 2r + sb \in R$$

and

$$\mathrm{N}_\sigma(x) = r^2 + rsb - s^2 a \in R$$

We note that $\mathrm{Tr}_\sigma(z) = b$ and $\mathrm{N}_\sigma(z) = -s^2 a$. Since the trace and norm take values in R, σ is the unique standard involution on S, as in Proposition 4.1.13. Note that if this

involution is trivial, then the characteristic of R is 2 and $b = 0$, so that $z^2 = a$. Let $r, s \in R$ be such that $r + sz$ is an arbitrary element of S fixed by σ. Then, $2sz = sb$. We have that $S = R \cdot 1 \oplus R \cdot z$ as an R-module, which means that $s = 0$. Hence, the set of fixed points S^σ of σ is R.

From our computations in the previous example, we can obtain some more information about the standard involution in the case that S is a domain and that R is a subring of S.

Corollary 4.1.17. Let R be a local subring of the integral domain S such that S is a quadratic R-algebra with standard involution σ. Then, σ is non-trivial if and only if the quadratic extension of fraction fields $\mathrm{Frac}\,(S) \,/\, \mathrm{Frac}\,(R)$ is separable.

Proof. Let E be the fraction field of S, and let K be the fraction field of R. Suppose that σ is trivial, and that $\{1, z\}$ is an R-basis of S. With the notation as in the computations of Example 4.1.16, this means that $b = 0$, so that z satisfies $z^2 = a \in R$. Therefore, $E = K(z)$, where z is inseparable over K. Conversely, if the quadratic extension E/K is inseparable, then the characteristic of R is 2. It then follows that every $z \in E$ satisfies $z^2 \in K$. As in the Example once more, this means that if $z^2 = a + bz$, for $a, b \in R$ and an R-basis vector $z \in S \backslash R$, then $b = 0$, so that $\sigma(z) = b - z = z$. So, the involution σ is trivial. ∎

Seeing as S/R is a free R-module of rank 2, we can describe the involution trace as in the case of fields, by appealing to a matrix embedding.

Example 4.1.18. Let S/R be a quadratic algebra over the local ring R. As we saw above, $S \cong R^2$ as R-modules and $\mathrm{Hom}_R(S, S) \cong M_2(R)$. Now, left multiplication by $\alpha \in S$ defines an R-module homomorphism $\ell_\alpha \colon S \to S$. The map

$$\rho_S \colon S \to M_n(R)$$

$$\alpha \mapsto \ell_\alpha$$

is called the left regular representation of S and gives us a matrix $M_\alpha \in M_2(R)$ representing ℓ_α, by the identification $\mathrm{Hom}_R(S, S) \cong M_2(R)$. We can therefore generalize the field trace to this setting by considering the trace and determinant of this matrix. We define

$$\mathrm{Tr}_\ell \colon S \to R$$

by

$$\alpha \mapsto Tr(M_\alpha).$$

Since $\mathrm{GL}(S)$ acts transitively on bases of S and R is commutative, this map is independent of the choice of R-basis for S. Then, with respect to our basis $\{1, z\}$ of Example 4.1.16, we have that $z^2 = a + bz$, so

$$M_z = \begin{bmatrix} 0 & a \\ 1 & b \end{bmatrix}$$

and therefore

$$\mathrm{Tr}_\ell(z) = b = z + b - z = z + \sigma(z) = \mathrm{Tr}_\sigma(z).$$

By R-linearity, $\mathrm{Tr}_\sigma = \mathrm{Tr}_\ell$.

Lemma 4.1.19. Let $R \subseteq S$ be local domains, and let S be a separable quadratic R-algebra with standard involution σ. The algebra trace map T_R^S is equal to the involution trace $\mathrm{Tr}_\sigma \colon S \to R$.

Proof. Let S/R be as above. If $E = \mathrm{Frac}(S)$ and $K = \mathrm{Frac}(R)$, we have that $E \cong K \otimes_R S$ is a separable field extension of K, by Corollary 2.5.18 and Theorem 2.5.15. By Corollary 4.1.17, σ induces the non-trivial Galois automorphism σ of E/K. Abusing notation, we denote by Tr_σ the two different traces induced by the involution σ on E and on S. By appealing to the existence of a normal K-basis $\{\alpha, \sigma(\alpha)\}$ for E/K and its trace-dual basis, and by applying Lemma 2.6.10, it can be shown that the algebra trace T_K^E coincides with the involution trace $\mathrm{Tr}_\sigma \colon E \to K$. Now, S is free as an R-module, and S has an R-basis $\{1, z\}$, as in Example 4.1.16. This basis is also a K-basis for E/K. If $\pi_1, \pi_z \colon E \to K$ are the K-linear projections on the coordinates 1 and z respectively, this defines a projective-dual basis for E/K. Lemma 2.6.10 then says that the algebra trace T_K^E is defined by $T_K^E(x) = \pi_1(x)) + \pi_z(xz)$ for all $x \in E$. Since $\{1, z\} \subseteq S$ and π_1, π_z can be interpreted as elements of $\mathrm{Hom}_R(S, R)$, the trace map T_K^E restricts to the algebra trace T_R^S. Since $\mathrm{Tr}_\sigma \colon E \to K$ restricts to $\mathrm{Tr}_\sigma \colon S \to R$ by construction, and $T_K^E = \mathrm{Tr}_\sigma$, we must have that $T_R^S = \mathrm{Tr}_\sigma$. ∎

This leads us to the proof of an important Lemma for our study of isometries. The following is a special case of a result stated in [Knu91, $p.130$].

Lemma 4.1.20. Let R be a local subring of the local ring S. If S is a separable quadratic R-algebra with involution σ, the sequence of R-modules

$$0 \longrightarrow \ker(\mathrm{Tr}_\sigma) \longrightarrow S \xrightarrow{\mathrm{Tr}_\sigma} R \longrightarrow 0$$

is split exact.

Proof. Since R is projective as an R-module, it suffices to show that Tr_σ is surjective. By the previous Example 4.1.18, we have that $\mathrm{Tr}_\ell = \mathrm{Tr}_\sigma$. But Tr_ℓ is also equal to trace map T_R^S of Definition 2.6.9. Since S is finitely generated, projective, and separable over R, Corollary 2.6.15 tells us that $T_R^S = \mathrm{Tr}_\sigma$ is surjective. ∎

To finish this section, we show that an involution on a discrete valuation ring preserves the valuation.

Lemma 4.1.21. Let S be a discrete valuation ring with involution σ and valuation ν. We have that

$$\nu(\sigma(x)) = \nu(x)$$

for all $x \in S$.

Proof. Let ϖ generate the unique maximal ideal $\mathfrak{P}$ of S, so that $\nu(\varpi) = 1$. Since σ is a ring automorphism, $\sigma(\varpi)$ is not a unit. Suppose that $\nu(\varpi) < \nu(\sigma(\varpi))$. Then, there exists $r \in S$ such that $r\varpi = \sigma(\varpi)$. Since σ is an involution, $\varpi = \sigma(r)\sigma(\varpi)$, so that $r\sigma(r)\sigma(\varpi) = \sigma(\varpi)$. Since S is a domain, $r\sigma(r) = 1$ and r is a unit. This shows $\varpi S = r\varpi S = \sigma(\varpi)S$ whence $\nu(\varpi) = \nu(\sigma(\varpi))$, a contradiction. Since every element of S is of the form $u\varpi^k$ for some $u \in S^\times$ and $k \in \mathbb{N}$, it follows that the ring automorphism σ is an isometry of ν. ∎

Corollary 4.1.22. Let K be the field of fractions of S equipped with its valuation ν extending that of S, together with the extension of the involution σ to K. Then

$$\nu(\sigma(x)) = \nu(x)$$

for all $x \in K$.

Proof. Let $x \in K \backslash S$. Since S is a valuation ring, $y = x^{-1} \in S$. Then,

$$\nu(x) = -\nu(y) = -\nu(\sigma(y)) = \nu(\sigma(y^{-1})) = \nu(\sigma(x)).$$

So, σ is an isometry of $\nu \colon K^\times \to \mathbb{Z}$. ∎

4.2 Hermitian forms, modules, and lattices

We continue our treatment of Hermitian forms over rings with involution, following [Knu91].

Definition 4.2.1. Let A be a commutative ring with involution σ, and let M be an A-module. A sesquilinear form on M is a map

$$\langle -, - \rangle \colon M \times M \to A$$

such that

1. $\langle x + y, z \rangle = \langle x, z \rangle + \langle y, z \rangle$ for all $x, y, z \in M$;
2. $\langle x, y + z \rangle = \langle x, y \rangle + \langle x, z \rangle$ for all $x, y, z \in M$;
3. $\langle ax, y \rangle = a\langle x, y \rangle = \langle x, \sigma(a)y \rangle$ for all $a \in A$ $x, y \in M$.

Note that $a\langle x, y \rangle = \langle x, \sigma(a)y \rangle$ is equivalent to $\sigma(a)\langle x, y \rangle = \langle x, ay \rangle$.

Definition 4.2.2. Let A be a commutative ring with involution σ, and let M be an A-module. A Hermitian form on M is a sesquilinear form

$$\langle -, - \rangle \colon M \times M \to A$$

satisfying the additional property that

$$\langle x, y \rangle = \sigma(\langle y, x \rangle)$$

for all $x, y \in M$. In other words, $\langle -, - \rangle$ is linear in the first argument and conjugate symmetric with respect to σ.

Definition 4.2.3. Let A be a commutative ring with involution. A Hermitian module M over A is a free and finitely generated A-module together with a Hermitian form $\langle -, - \rangle \colon M \times M \to A$. If A is a field, we say that M is a Hermitian space.

To reduce wordiness, we will often say that M is a Hermitian module without mention of the attached Hermitian form. Moreover, we will assume that the involution σ is a standard involution. Unless otherwise stated, involutions are non-trivial. In

particular, if E/K is a quadratic field extension, σ is the non-trivial automorphism of this extension.

Hermitian spaces generalize Euclidean vector spaces equipped with a positive definite inner product. Indeed, if we look at the standard inner product

$$\langle \vec{x}, \vec{y} \rangle = \sum_{i=1}^{n} x_i \overline{y_i}$$

on $\mathbb{C}^n$, we have that this product is linear in the first variable and conjugate symmetric with respect to complex conjugation, the non-trivial Galois automorphism of the quadratic field extension $\mathbb{C}/\mathbb{R}$.

If the characteristic of the field E is not equal to 2, Hermitian forms can be viewed as generalizations of quadratic forms as follows. Let us consider V, a finite-dimensional vector space over some field E where E is equipped with the trivial involution. Then, a Hermitian form $\langle ., . \rangle \colon V \times V \to E$ satisfies

$$\langle x, y \rangle = \langle y, x \rangle,$$

i.e. $\langle -, - \rangle$ is symmetric. From symmetry and linearity in the first argument of $\langle -, - \rangle$, it follows that $\langle -, - \rangle$ is linear in both arguments. In this case, we call V a symmetric bilinear form. Conversely, any symmetric bilinear form is a Hermitian form for the trivial involution. The map $\mathrm{Q} \colon V \to E$ defined by

$$\mathrm{Q}(x) = \frac{1}{2} \langle x, x \rangle$$

is a quadratic form over V. If $\mathrm{Q}(x)$ is any quadratic form over V, we have that the polarization identity

$$\langle x, y \rangle := \frac{1}{2} \left(\mathrm{Q}(x + y) - \mathrm{Q}(x) - \mathrm{Q}(y) \right)$$

defines a symmetric bilinear form over V. It can be checked that this pairing is two-sided inverse to the first, so that there is a set-theoretic bijection between quadratic forms and symmetric bilinear forms over V.

The theory of quadratic forms is given a thorough treatment in T.Y. Lam's classic text, Introduction to Quadratic Forms over Fields [Lam05]. Unfortunately, the correspondence between quadratic forms and symmetric bilinear forms breaks down

in characteristic 2. The astute reader will have noticed that the polarization identity makes use of division by 2, and indeed, even if we adopt a modified polarization identity to avoid the use of division by 2, there exist symmetric bilinear forms in characteristic 2 which do not arise as a polar of a quadratic form. For further details, see [Gro02, §12].

Next, we define Hermitian lattices.

Definition 4.2.4. Let M be a free and finitely generated module over the commutative ring E, with A a subring of E. A A-lattice L in M is a finitely generated A-submodule of M, with the A-module structure on L inherited from the action of E on M.

If it is clear which rings $A \subseteq E$ we are referring to, we call a finitely generated A-submodule of M a lattice in M.

Definition 4.2.5. Let $A \subseteq E$ be a commutative rings, and suppose L is an A-lattice in the E-module M. A sublattice of L is an A-submodule of L.

Definition 4.2.6. Let A be a subring of the ring E, and let L be an A-lattice in the E-module M. We say that L is a full lattice if extension of scalars to E yields an isomorphism of E-modules $E \otimes_A L \cong M$.

Recall the definitions of a discrete valuation ring and of a discretely valued field. We will be almost exclusively examining the case where M above is a finite-dimensional E-vector space V over the discretely valued field E and where L is a submodule of V over the valuation ring $\mathcal{O}_E$ of E. In particular, we will be interested in the case where E/K is a Galois extension of even degree, so that the Galois group of E/K has an order 2 automorphism which makes E into a ring with involution.

Definition 4.2.7. Let M be a Hermitian module over the commutative ring with involution E, and let A be a subring of E. A Hermitian lattice is an A-lattice L in M, together with the restriction to L of the Hermitian form $\langle ., . \rangle \colon M \times M \to E$. If the choice of subring A of E is unambiguous, we can simply call L a Hermitian lattice.

We note that the restriction of the Hermitian form on M to an A-submodule L will not necessarily take values in A. Therefore, L is not necessarily a Hermitian

A-module. However, the relation

$$a\langle \ell_1, \ell_2 \rangle = \langle a\ell_1, \ell_2 \rangle$$

for $a \in A$ and $\ell_1, \ell_2 \in L$ shows that the set of values taken by the Hermitian form on $L \times L$ is an A-submodule of the E-module E. In the context of chapter 1, i.e., if A is a Dedekind domain with field of fractions E, then $\langle L, L \rangle$ is a fractional ideal of A.

We now introduce some terminology common to both Hermitian modules and lattices.

Definition 4.2.8. Let A be a commutative ring with involution. Let M be a Hermitian module or, equally well, a Hermitian lattice. We say that $x \in M$ is isotropic if x is non-zero and $\langle x, x \rangle = 0$. Otherwise, we say that x is anisotropic. In particular, the zero vector is anisotropic. We say that M is isotropic if it contains an isotropic vector, and anisotropic otherwise.

Definition 4.2.9. Let A be a commutative ring with involution. Let M be a Hermitian module or, equally well, a Hermitian lattice. We say that x and y in M are orthogonal if $\langle x, y \rangle = 0$. The orthogonal complement of a subset S of M is the set

$$S^{\perp} := \{x \in M \mid \langle x, s \rangle = 0, s \in S\}.$$

We define the radical of M to be

$$\mathrm{rad}(M) := M^{\perp}.$$

We say that M is non-degenerate if $\mathrm{rad}(M) = M^{\perp} = 0$.

Remark 4.2.10. Denote by $M^{\vee}$ the dual module to M in the category of left A-modules. In the case that M is free and finitely generated over A, M is non-degenerate if and only if the map

$$h \colon M \to M^{\vee}$$

given by

$$m \mapsto (n \mapsto \langle m, n \rangle)$$

is injective. Indeed, if $m \in M$ is contained in $\mathrm{rad}(M)$, then $n \mapsto \langle m, n \rangle$ is identically 0, and conversely any $m \in \ker(h)$ is in $\mathrm{rad}(M)$.

Definition 4.2.11. Let M be a Hermitian module, and let N_1, N_2 be submodules of M. We say that M is the orthogonal sum of N_1 and N_2 if M is the direct sum of N_1 and N_2 and if $\langle N_1, N_2 \rangle = 0$. We then write $M = N_1 \oplus N_2$.

Definition 4.2.12. Let M be a Hermitian module, and let N be a submodule of M. We say that N is a component of M if there exists an orthogonal decomposition of M as

$$L = N \oplus N',$$

for some submodule N' of M.

The orthogonal sum of Hermitian modules $(M, \langle -, - \rangle_1)$ and $(N, \langle -, - \rangle_2)$ over the commutative ring A with involution σ always exists. We can take the direct sum of A-modules M and N and set $\langle ., . \rangle \colon M \oplus N \times M \oplus N \to A$ to be

$$\langle m + n, m' + n' \rangle = \langle m, m' \rangle_1 + \langle n, n' \rangle_2.$$

The map $\langle -, - \rangle$ is then a Hermitian form on the A-module $M \oplus N$, and $M \oplus N$ is an orthogonal decomposition for this module.

Definition 4.2.13. We say two Hermitian A-modules $(M_1, \langle -, - \rangle_1)$ and $(M_2, \langle -, - \rangle_2)$ are isometric if there exists an A-linear isomorphism $T \colon M_1 \to M_2$ such that

$$\langle T(x), T(y) \rangle_2 = \langle x, y \rangle_1$$

for all $x, y \in M_1$. In this case, we write $M_1 \cong M_2$.

We define the analogous notion for Hermitian lattices.

Definition 4.2.14. Let B be a subring of the ring with involution A. We say two Hermitian B-lattices $(L_1, \langle -, - \rangle_1)$ and $(L_2, \langle -, - \rangle_2)$ are isometric if there exists a B-linear isomorphism $T \colon L_1 \to L_2$ such that

$$\langle T(x), T(y) \rangle_2 = \langle x, y \rangle_1$$

for all $x, y \in L_1$. In this case, we write $L_1 \cong L_2$.

We have that isometry is an equivalence relation on Hermitian modules. Moreover, if L and M are Hermitian lattices, then $L \cong M$ implies that the Hermitian spaces $E \otimes_{\mathcal{O}_E} L$ and $E \otimes_{\mathcal{O}_E} M$ are isometric, but the converse is false.

4.3 The Gram matrix of a Hermitian lattice

We can represent Hermitian modules and lattices by matrices.

Definition 4.3.1. Let E be a commutative ring with involution, and let M be a Hermitian module over E. Suppose M has a basis $\{x_1, \ldots, x_n\}$. We define the Gram matrix of M with respect to this basis to be the matrix

$$H_M = [h_{ij}] := [\langle x_i, x_j \rangle].$$

In other words, the entries h_{ij} of the matrix H_M are computed by evaluating the Hermitian form $\langle -, - \rangle$ at each ordered pair of basis vectors x_i, x_j, for $1 \leq i, j \leq n$.

Definition 4.3.2. Suppose A is a subring of the ring with involution E, and that L is a Hermitian A-lattice in the Hermitian E-module M. If L is free as a A-module with basis $\{x_1, \ldots, x_n\}$, we define the Gram matrix of L to be

$$H_L = [h_{ij}] := [\langle x_i, x_j \rangle].$$

The entries of the Gram matrix H_L of a Hermitian A-lattice are elements of the ring E. These entries belong to the subring A if and only if L is a Hermitian A-module. We now discuss the properties of the matrix representation of a Hermitian form $\langle ., . \rangle$. Let E be a ring with involution, and let A be a subring of E. Since the Hermitian form $\langle -, - \rangle$ is conjugate symmetric with respect to the involution σ, we have the relation $h_{ij} = \sigma(h_{ji})$ for each matrix coefficient of H_M. This shows H_M is a σ-Hermitian matrix, i.e., we have

$$H_M = H_M^*$$

with $H \mapsto H^*$ the conjugate transpose involution defined in Example 4.1.6.

Definition 4.3.3. Suppose that A is a ring with involution σ and that $H \in M_n(A)$ is a matrix. We say that H is σ-Hermitian if $H = H^*$, for the conjugate transpose involution $H \mapsto H^*$ induced by σ on $M_n(A)$.

Any σ-Hermitian matrix H with entries in a commutative ring with involution A defines a Hermitian form

$$\langle -, - \rangle_H \colon A^n \times A^n \to A$$

given by

$$\langle x, y\rangle_H := x^t H y^\sigma,$$

when interpreting $x, y \in A^n$ as column vectors. This allows us to associate a Hermitian module M_H to every Hermitian matrix H in $M_n(A)$.

Definition 4.3.4. Suppose that A is a commutative ring with involution and that $H \in M_n(A)$ is a σ-Hermitian matrix. We define the Hermitian module M_H to be the free A-module A^n, together with the Hermitian form $\langle -, -\rangle_H$.

Similarly, given a subring B of A, every Hermitian matrix H defines a Hermitian B-lattice L_H.

Definition 4.3.5. Let E be a commutative ring with involution, and let A be a subring of E. If $H \in M_n(E)$ is a σ-Hermitian matrix, we can define the lattice L_H to be the free A-module A^n together with the restriction to A^n of the Hermitian form $\langle -, -\rangle_H$ defined on E^n.

Let us investigate isometries at the level of matrices. If there exists an isometry $T\colon M \to M'$ and $H_M, H_{M'}$ are Gram matrices for M and M' respectively, then, as with a Euclidean inner product on $\mathbb{C}^n$, we have the matrix equation

$$H_M = T^t H_{M'} T^\sigma$$

for the matrix T representing the isometry $T\colon M \to M'$ with respect to the chosen bases of M and M'.

Definition 4.3.6. Let E be a commutative ring with involution σ, with A a subring of E. Suppose that $H_1, H_2 \in M_n(E)$ are Hermitian matrices. We say that H_1 is congruent to H_2 over A if there exists an invertible matrix $T \in M_n(A)$ such that

$$H_1 = T^t H_2 T^\sigma.$$

Since A is commutative, we can check that congruence is an equivalence relation on Hermitian matrices. We say that Hermitian matrices H_1 and H_2 are in the same congruence class over B if H_1 and H_2 are congruent over B.

Proposition 4.3.7. Let A be a commutative ring with involution σ. There is a bijection ψ between isometry classes of Hermitian modules and congruency classes of Hermitian matrices over A. Explicitly, we have that the bijection ψ is defined on representatives by

$$M \mapsto H_M,$$

where H_M is the Gram matrix of the Hermitian module M.

Proof. We give a sketch. If $M_1 \cong M_2$, then H_{M_1} and H_{M_2} are congruent, so ψ is well defined. The map ψ is surjective, as the Gram matrix of $\langle -, - \rangle_H$ relative to the standard basis of A^n is H. If H_{M_1} and H_{M_2} are congruent, there is an A-isometry $T\colon M_2 \to M_1$ corresponding to the congruence, so ψ is injective. ∎

In a similar fashion, the isometry classes of free B-lattices can be identified with the B-congruency classes of Hermitian matrices in $M_n(A)$.

To finish the section, we introduce the determinant of a Hermitian module or lattice.

Definition 4.3.8. Let M be a Hermitian module with basis $\{x_1, \ldots, x_n\}$. We define the determinant $d(M)$ of M with respect to the basis $\{x_1, \ldots, x_n\}$ to be $\det([\langle x_i, x_j \rangle])$.

We see that isometric Hermitian modules and lattices have the same determinant up to a norm of the involution σ as

$$\det(H_{M'}) = \det(T^t)\det(T^\sigma)\det(H_M) = \mathrm{N}_\sigma(\det(T))\det(H_M).$$

Since T is invertible, $\det(T)$ is a unit of E in the case of Hermitian modules or a unit of A in that of Hermitian lattices. This means that the determinant of a Hermitian module or lattice is well defined up to the norm of a unit of E or A. We then write $d(M) \sim d(N)$ if $d(N) = d(M)\,\mathrm{N}_\sigma(u)$, for u a unit of A. To finish the section, we present the following from [Gro02, p.18, Prop 2.9], which tells us that non-degenerate subspaces of a non-degenerate Hermitian space over a field have orthogonal complements.

Lemma 4.3.9. Let V be a non-degenerate Hermitian space over the field E, and let W be a non-degenerate subspace of V, i.e., $W \cap W^\perp = \{0\}$. Then $V = W \oplus W^\perp$.

Proof. Let $w_1, \ldots, w_k$ be a basis for W and extend this basis with $w_{k+1}, \ldots, w_n$ to a basis $\mathcal{S}$ for V. Let H be the Gram matrix of the form on V with respect to the given ordered basis, and denote by H_W the $k \times n$ matrix obtained from H by truncating the last $n-k$ rows of H. Since H is invertible, H_W has k linearly independent rows. Then, note that $v \mapsto v^\sigma$ is a bijective semi-linear map of $V = E^n$. Hence, the map $v \mapsto v^\sigma$ restricts to a bijective semi-linear map on $W^\perp$. By the matrix calculus of Hermitian forms and the construction of H_W, elements of $W^\perp$ are mapped bijectively to $\ker(H_W)$ under this map. In particular, $\ker(H_W)$ is a subspace of dimension $n-k$ of V. Since $W^\perp$ is also a subspace of V, $W^\perp$ must have dimension at least $n-k$. Since $W \cap W^\perp = 0$, we must have that $V = W + W^\perp$, and moreover, that $V = W \oplus W^\perp$ by the definition of $W^\perp$. ∎

4.4 Hermitian lattices over discrete valuation rings

In this section, we extend the exposition given in [Jac62] for lattices over the valuation ring of a non-Archimedean local field to general discrete valuation rings. In particular, the definitions given in this section are extensions of those from [Jac62]. We indicate which results are extensions of those from the same paper. We now focus our attention on the class of rings of involution, which we will consider for the remainder of the thesis.

Definition 4.4.1. Let A and B be discrete valuation rings such that B is an A-algebra. We say that B/A is an extension of discrete valuation rings if the structure map $A \to B$ is an injective local homomorphism, or equivalently, if A is a subring of B. We say that B/A is a finite extension if B is finitely generated as an A-module. We say that B/A is a quadratic extension if B is a quadratic A-algebra.

Example 4.4.2. Consider a discrete valuation ring A with field of fractions K, equipped with the valuation induced by A. Suppose that E is a separable quadratic extension of K on which there is a unique extension of the valuation on K, and let B be the integral closure of A in E. By Proposition 3.4.14, there is a unique prime ideal $\mathfrak{P}$ of B lying over $\mathfrak{p}$, and the valuation ring of E is equal to the localization $B_{\mathfrak{P}}$. By Corollary 2.4.15, B is a free A-module of rank 2 containing A. Since B is integral over A, we apply Lemma 3.1.17 and get that every non-zero prime ideal of B lies over $\mathfrak{p}$. As $\mathfrak{P}$ is the only prime dividing $\mathfrak{p}$, we must have that B is a local

ring with maximal ideal $\mathfrak{P}$. In particular, $B = B_{\mathfrak{P}}$ and B is the valuation ring of E with respect to the extension of the valuation on K. This means that B/A satisfies our definition of a quadratic extension of discrete valuation rings. In particular, if E/K is a quadratic extension of p-adic fields, Lemma 3.5.17 says that the extension of the p-adic valuation to a discrete valuation of K and E is unique. So, the integral closure of the valuation ring $\mathcal{O}_K$ of K is the valuation ring $\mathcal{O}_E$ of E, and $\mathcal{O}_E/\mathcal{O}_K$ is a quadratic extension of discrete valuation rings.

We are now ready to fix the notation of the rest of the thesis. In what follows, we write $\mathcal{O}_E/\mathcal{O}_K$ for a quadratic extension of discrete valuation rings and E/K for the corresponding quadratic extension of fields of fractions. As in Example 4.1.16, $\mathcal{O}_E/\mathcal{O}_K$ is equipped with its unique standard involution σ. By Lemma 2.4.12, $\mathcal{O}_E$ contains a K-basis for E, so σ induces a K-automorphism of E/K, also denoted by σ. We begin with the definition of a Hermitian lattice over a quadratic extension of discrete valuation rings.

Definition 4.4.3. Let $\mathcal{O}_E/\mathcal{O}_K$ be a quadratic extension of discrete valuation rings equipped with its standard involution σ, let E/K be the corresponding extension of fraction fields equipped with the K-automorphism induced by σ, and let V be a finite-dimensional Hermitian space over E. A Hermitian lattice over $\mathcal{O}_E$ is a finitely generated $\mathcal{O}_E$-submodule L of V, together with the restriction of the Hermitian form $\langle -, - \rangle$ of V to L. We say L is a full lattice in V if $E \otimes L = V$.

Remark 4.4.4. As Lemma 3.4.8 says that $\mathcal{O}_E$ is a principal ideal domain, it is of course equivalent to take L to be a free $\mathcal{O}_E$-submodule of V of finite rank. As seen in an introductory abstract algebra course, a finitely generated module L over $\mathcal{O}_E$ is isomorphic to the direct sum

$$\mathcal{O}_E^k \oplus \mathcal{O}_E/I_{k+1} \oplus \cdots \oplus \mathcal{O}_E/I_n,$$

where $\mathcal{O}_E^k$ is free and the last $n-k-1$ summands are torsion. Since the $\mathcal{O}_E$-module structure on L is inherited from the E-vector space structure on V, the action of $\mathcal{O}_E$ on L is torsion-free. Therefore, every summand in the cyclic decomposition of L is isomorphic to $\mathcal{O}_E$, so there is an isomorphism of $\mathcal{O}_E$-modules from L to $\mathcal{O}_E^n$.

As mentioned earlier, it is worth noting that every Hermitian module is a Hermitian lattice, but not conversely, as the restriction of the Hermitian form on V to

L can take values in a fractional ideal of $\mathcal{O}_E$, which properly contains $\mathcal{O}_E$. When quantifying over lattices, we will usually mean full lattices. We quickly characterize full lattices as those of maximal rank.

Lemma 4.4.5. Suppose that A is an integral domain with field of fractions K, and that L is a free A-module with basis $x_1, \ldots, x_n$. Then $1 \otimes x_1, \ldots, 1 \otimes x_n$ is a basis for the K-vector space $K \otimes_A L$.

Proof. If $\sum_{i=1}^n \frac{a_i}{b_i} \otimes_A x_i = 0$ is an equation of linear dependence for the $1 \otimes x_i$, then $(\prod_{j=1}^n b_j) \sum_{i=1}^n a_i x_i = 0$ is an equation of linear dependence for the x_i over A. Hence, the $1 \otimes_A x_i$ are linearly independent. Since these elements generate $V = K \otimes_A L$, it results that V has a basis $1 \otimes_A x_1, \ldots 1 \otimes_A x_n$. ∎

Corollary 4.4.6. Let V be a Hermitian space over the discretely valued field E, with $\dim(V) = n$. A Hermitian lattice L in V is a full lattice if and only if $\mathrm{rk}(L) = n$.

One of our main goals in this chapter is to describe isometry classes of Hermitian lattices over the discrete valuation ring $\mathcal{O}_E$, and we will start by describing the orthogonal components of a Hermitian lattice. Let us then focus our attention on Hermitian lattices over the discrete valuation ring $\mathcal{O}_E$, with the goal of describing their orthogonal decompositions in general. The main obstruction to diagonalization is that $\mathcal{O}_E$ is not a division ring, so we may fail to find Fourier coefficients for the generalized Gram-Schmidt algorithm of Theorem 4.5.2. To begin describing vectors which do yield maximal values under application of the Hermitian form to L, we introduce the language of maximal vectors of L. We further characterize them as basis vectors for L. Recall the definition of the lattice $aL = \{ax \mid x \in L\}$ for any $a \in \mathcal{O}_E$.

Definition 4.4.7. Let ϖ be a uniformizer for $\mathcal{O}_E$. We say a vector $x \in L$ is maximal if it does not lie in ϖL.

Lemma 4.4.8. Let L and M be lattices, and let $\varphi\colon L \to M$ be an $\mathcal{O}_E$-module isomorphism. Then φ is a bijection between the maximal vectors of L and the maximal vectors of M.

Proof. We have that $\varphi(\varpi L) = \varpi\varphi(L) = \varpi M$, so by symmetry, $\varphi^{-1}(\varpi M) = \varpi L$. This means φ induces a bijection between the non-maximal vectors of ϖL and ϖM. But φ also induces a bijection between the set-theoretic complements of ϖL and ϖM,

so $x \in L$ is not in ϖL if and only if $\varphi(x)$ is not in ϖM, which coincides with the definition of maximal vectors in L and M respectively. ∎

We provide the details of the proof of the following result stated in [Jac62, Prop. 4.1].

Proposition 4.4.9. Let L be a Hermitian lattice.

1. Each vector in a basis for L is maximal.
2. A maximal vector in L can always be extended to a basis for L.
3. If x is any vector in L, there exists $a \in \mathcal{O}_E$ and a maximal vector $y \in L$ such that $x = ay$.

Proof. Since L is a free $\mathcal{O}_E$-module of rank n, there exists an $\mathcal{O}_E$-isomorphism $\varphi\colon L \to \mathcal{O}_E^n$. By Lemma 4.4.8, it is sufficient to show that any basis vector for $\mathcal{O}_E^n$ is maximal. Now, if $x_1 = (a_1, \ldots, a_n)$ is a basis vector for $\mathcal{O}_E^n$ and $\vec{x}$ is not maximal, then ϖ divides every a_i. But the matrix with columns corresponding to the chosen basis $\{x_1, \ldots, x_n\}$ is invertible in $M_n(\mathcal{O}_E)$, so we can write $1 = \sum_{j=i}^n a_j b_j$ for coefficients $b_j \in \mathcal{O}_E$. But $a_j b_j \in \varpi\mathcal{O}_E$ for every $1 \leq j \leq n$, since each a_j is in $\varpi\mathcal{O}_E$. We must then get $\nu(\sum_{j=i}^n b_j a_j) > 0$, which is a contradiction with $\sum_{j=i}^n b_j a_j = 1$. Hence, x_1 must be maximal. This proves the first claim. Because the isomorphism φ induces the bijection of Lemma 4.4.8 between maximal vectors, a maximal vector $x_1 \in L$ can be extended to a basis in L if and only if there exists a maximal vector $\varphi(x_1) \in \varphi(M)$ which can be extended to a basis of $\varphi(M)$. Therefore, we can assume $L = \mathcal{O}_E^n$. Consider a maximal vector $x_1 = (a_1, \ldots, a_n)$ of $\mathcal{O}_E^n$. Since x_1 is maximal, we have that ϖ does not divide every a_i, so we can assume without loss of generality that a_1 is a unit of $\mathcal{O}_E$. If e_i denotes the i^{th} standard basis vector of $\mathcal{O}_E^n$, then it is readily verified that $\{x_1, e_2, \ldots, e_n\}$ is a basis of $\mathcal{O}_E^n$. This proves the second claim. To prove the third property, we can assume once more that $L = \mathcal{O}_E^n$, since the isomorphism φ allows us to transfer an expression of the form $x = ay \in \mathcal{O}_E^n$ to L. If $(a_1, \ldots, a_n)$ is an arbitrary vector of $\mathcal{O}_E^n$, there exists a minimal j such that a_i is in $\varpi^j\mathcal{O}_E$ for all $1 \leq i \leq n$. In this case, $y := \varpi^{-j}(a_1, \ldots, a_n)$ is maximal and $\varpi^j y = (a_1, \ldots, a_n)$ with $\varpi^j \in \mathcal{O}_E$, which proves the third claim. ∎

We introduce some basic manipulations on the Hermitian form for the lattice L. These will be useful for making technical reductions later on.

Definition 4.4.10. Let $(L, \langle -, - \rangle)$ be a Hermitian lattice and $\alpha \in K^\times$. The scaled lattice $\alpha \cdot L$ is the $\mathcal{O}_E$-submodule L together with the Hermitian form $\alpha \cdot \langle -, - \rangle \colon L \times L \to E$ given by

$$\alpha \cdot \langle x, y \rangle = \alpha \langle x, y \rangle.$$

At the level of matrices, the Gram matrix of $\alpha \cdot L$ is simply αH, for H a Gram matrix for L.

We can also scale vectors by units of $\mathcal{O}_E$. Since $\mathcal{O}_E$ is a discrete valuation ring, every non-orthogonal pair of vectors $x, y \in L$ has the property that $\langle x, y \rangle = m\varpi^i$, for some $m \in \mathcal{O}_E^\times$ and $i \in \mathbb{Z}$. It is desirable to work directly with a uniformizer in arguments involving the Hermitian form, so we will often replace pairs of vectors $\{x, y\}$ by vectors $\{x_1, y_1\}$ with the property that $\langle x_1, y_1 \rangle = \varpi^j$ for some $j \in \mathbb{Z}$. We would like to impose a common language for this useful but simple operation.

Definition 4.4.11. Let $\{x, y\}$ be linearly independent vectors with $\nu(\langle x, y \rangle) = \nu(\varpi^i) = i$. Then, setting $x_1 := \frac{\varpi^i}{\langle x, y \rangle} x$ and $y_1 := y$, we get $\langle x_1, y_1 \rangle = \varpi^i$. Equivalently, we can set $x_1 := x$ and $y_1 := \frac{(\sigma(\varpi))^i}{\langle y, x \rangle} y$. We call the vectors $\{x_1, y_1\}$ a rescaling of $\{x, y\}$.

Since the coefficients of x and y are units of $\mathcal{O}_E$, we have $\operatorname{span}\{x_1, y_1\} = \operatorname{span}\{x, y\}$ and that $\{x_1, y_1\}$ is linearly independent over $\mathcal{O}_E$. This means that we can often perform rescaling as defined above without loss of generality. We would now like to take a measure of the image of our lattice under the Hermitian form $\langle -, - \rangle$.

Definition 4.4.12. The fractional ideal

$$\mathfrak{s}(L) := \{\langle x, y \rangle \mid x, y \in L\}$$

is called the scale of L and the fractional ideal $\mathfrak{n}(L)$ generated by $\{\langle x, x \rangle \mid x \in L\}$ is called the norm of L. We have $\mathfrak{n}(L) \subseteq \mathfrak{s}(L)$ and if equality holds, we say L is normal. If the inclusion is proper, we say L is subnormal.

We need the following lemma to describe some properties of the scale and norm ideals of a Hermitian lattice L.

Lemma 4.4.13. Suppose that $a, b \in E$ satisfy $\nu(a) = \nu(b)$. Then $aI = bI$ for every fractional ideal I of $\mathcal{O}_E$.

Proof. The relation $\nu(a) = \nu(b)$ means that we have that $a = ub$ for some $u \in \mathcal{O}_E^\times$. So, $\nu(ax) = \nu(ubx) = \nu(bx)$ for every $x \in I$. Since fractional ideals of $\mathcal{O}_E$ are uniquely determined by their valuation group, $aI = bI$. ∎

We list and prove some properties of the scale and norm of a lattice over a discrete valuation ring, as stated in [Jac62, pp.447-448, 4.1-4.3] for Hermitian lattices over the valuation ring of a local field.

Proposition 4.4.14. Let $\mathcal{O}_E/\mathcal{O}_K$ be a quadratic extension of discrete valuation rings with standard involution σ, let $\alpha \in K^\times$, and let $a \in E^\times$. If L is a Hermitian lattice over $\mathcal{O}_E$, the following statements hold.

1. $\mathfrak{s}(aL) = a^2\mathfrak{s}(L)$ and $\mathfrak{n}(aL) = a^2\mathfrak{n}(L)$ for $a \in E^\times$.
2. $\mathfrak{s}(\alpha \cdot L) = \alpha\mathfrak{s}(L)$ and $\mathfrak{n}(\alpha \cdot L) = \alpha\mathfrak{n}(L)$
3. $\mathfrak{s}(L \bigoplus M) = \mathfrak{s}(L) + \mathfrak{s}(M)$ and $\mathfrak{n}(L \bigoplus M) = \mathfrak{n}(L) + \mathfrak{n}(M)$ for lattices L and M.

Proof. Let $x, y \in L$ and let $a \in E^\times$. By Lemma 4.1.21, $\nu\left(a\sigma(a)\right) = 2\nu(a) = \nu(a^2)$, so we can apply Lemma 4.4.13 to get

$$\langle ax, ay\rangle = a\sigma(a)\langle x, y\rangle \in a^2\mathfrak{s}(L)$$

for all $x, y \in L$. Since this relation holds for a generator $\langle x, y\rangle$ of the principal ideal $\mathfrak{s}(L)$ and the ultrametric inequality holds in a discrete valuation ring, we have that $\mathfrak{s}(aL) = a^2\mathfrak{s}(L)$. Similarly, we have $a^2\mathfrak{n}(L) = a\sigma(a)\mathfrak{n}(L)$. The second property follows immediately from the definition of $\alpha \cdot L$.

If $x \in L$ and $z \in M$, then for $x + z \in L \oplus M$, we have, by orthogonality, that

$$\langle x + z, L \oplus M\rangle = \langle x, L\rangle + \langle z, M\rangle$$

and

$$\langle x + z, x + z\rangle = \langle x, x\rangle + \langle z, z\rangle.$$

These mean that $\mathfrak{s}(L \bigoplus M) = \mathfrak{s}(L) + \mathfrak{s}(M)$ and $\mathfrak{n}(L \bigoplus M) = \mathfrak{n}(L) + \mathfrak{n}(M)$. ∎

To finish this section, we introduce the concept of $\mathfrak{P}^i$-modularity for a Hermitian lattice L. Modularity will let us describe particularly nice orthogonal components of L in terms of what values are taken by the Hermitian form $\langle -, - \rangle$. In fact, our ultimate goal in the study of Hermitian lattices is to show that L admits an orthogonal decomposition in terms of modular components, and that this decomposition is unique up to rearrangement of the summands.

Definition 4.4.15. Let $i \in \mathbb{Z}$ and let $\mathfrak{P}$ be the unique maximal ideal of $\mathcal{O}_E$. We say that the lattice L is $\mathfrak{P}^i$-modular if we have $\langle x, L \rangle = \mathfrak{P}^i$ for every maximal vector $x \in L$. If L is $\mathfrak{P}^0$-modular, we say L is unimodular.

Remark 4.4.16. In light of Proposition 4.4.9, we have that a lattice L is $\mathfrak{P}^i$-modular if and only if $\langle x, L \rangle = \mathfrak{P}^i$ for every vector x in every basis of L. If L is modular, there always exists $i \in \mathbb{Z}$ such that the lattice $\varpi^i \cdot L$ is unimodular.

In essence, if a lattice is modular, then it cannot be broken up as an orthogonal sum of pieces of smaller scale. We give a few examples of small modular lattices, and introduce hyperbolic planes, which are an analogue of hyperbolic planes in a quadratic space (see [Lam05, §1]).

Example 4.4.17. The line L generated by $x \in L$ is ϖ^i-modular for an appropriate $i \in \mathbb{Z}$, since $\nu(\langle x, x \rangle) = \nu(\pi^i) = i$ and x is maximal in L.

Example 4.4.18. A plane with basis $\{x, y\}$ such that

$$\nu(\langle x, y \rangle) = \nu\left(\varpi^i\right), \quad \nu(\langle x, x \rangle) > \nu\left(\varpi^i\right), \quad \text{and } \nu(\langle y, y \rangle) > \nu\left(\varpi^i\right)$$

is ϖ^i-modular.

Definition 4.4.19. For each $i \in \mathbb{Z}$, we say the rank 2 lattice with Gram matrix

$$\begin{bmatrix} 0 & \varpi^i \\ \sigma(\varpi^i) & 0 \end{bmatrix}$$

is a hyperbolic plane of scale i.

The following example, which is worth keeping in mind until we prove the main result of the thesis, illustrates that not all lattices are normal. When we prove Theorem 4.9.1, we will heavily rely on the fact that our Hermitian lattices have orthogonal

bases, and since normality is a necessary condition to obtain an orthogonal basis for a lattice, subnormality is a major obstacle to extending the given proof of Theorem 4.9.1 to subnormal lattices. Prefacing Lemma 4.7.2, we mention that subnormal lattices only exist over inseparable extensions of discrete valuation rings, thus motivating the separability hypothesis of the main Theorem. Now, to obtain an inseparable extension of DVR's, we turn to the valuation ring of a ramified extension of $\mathbb{Q}_2$.

Example 4.4.20. Fix a uniformizer $\varpi \in \mathcal{O}_E$, and let L be a hyperbolic plane of rank i. If $\{x, y\}$ is a basis for L, as ordered for the matrix representation above, we have

$$\langle x, x\rangle = \langle y, y\rangle = 0, \text{ and } \langle x, y\rangle = \varpi^i = \sigma(\langle y, x\rangle).$$

In particular, L is a $\mathfrak{P}^i$-modular lattice. We have that

$$\langle u, u\rangle = \sigma(u_1)u_2\varpi^i + u_1\sigma(u_2)\sigma(\varpi^i) = \mathrm{Tr}_\sigma\left(\sigma(u_1)u_2\varpi^i\right)$$

for every $u = \begin{bmatrix} u_1 & u_2 \end{bmatrix}^t \in L$. However, not every hyperbolic plane is normal. From Example 4.4.2, recall that quadratic extensions of p-adic fields induce quadratic extensions of discrete valuation rings. If E is a (totally) ramified quadratic extension of $\mathbb{Q}_2$ and $\mathcal{O}_E/\mathcal{O}_K$ is the corresponding extension of valuation rings, then the residue field of $\mathcal{O}_E$ is $\mathbb{F}_2$ and every quotient space $\mathfrak{P}^i/\mathfrak{P}^{i+1}$ can be identified with $\mathbb{F}_2$. But $\nu(u_1\sigma(u_2)\varpi^i) = \nu(\sigma(u_1)u_2\sigma(\varpi^i))$. This means that writing $j = \nu(u_1\sigma(u_2)\varpi^i)$, then $u_1\sigma(u_2)\varpi^i$ and $\sigma(u_1)u_2\sigma(\varpi^i)$ are in $\mathfrak{P}^j\backslash\mathfrak{P}^{j+1}$. Since $\mathfrak{P}^j/\mathfrak{P}^{j+1} \cong \mathbb{F}_2$, we have

$$u_1\sigma(u_2)\varpi^i \equiv \sigma(u_1)u_2\sigma(\varpi^i) \pmod{\mathfrak{P}^{j+1}}$$

and $\mathrm{Tr}_\sigma(u_1\sigma(u_2)\varpi^i) \in \mathfrak{P}^{j+1} \subsetneq \varpi^i\mathcal{O}_E$ for any $u \in L$. So, L is not normal.

On the other hand, if $\mathcal{O}_E$ is a separable extension of $\mathcal{O}_K$, we have more luck. By Lemma 2.5.16, we can write $\varpi^i = c\pi^i$ for $c \in \mathcal{O}_E^\times$ and a generator $\pi \in \mathcal{O}_K$ of the maximal ideal $\mathfrak{p}$ of $\mathcal{O}_K$. Since the trace is $\mathcal{O}_K$-linear, we therefore have that $\mathrm{Tr}_\sigma(u_1\sigma(u_2)c\pi^i) = \mathrm{Tr}_\sigma(u_1\sigma(u_2)c)\pi^i$. But, by Lemma 4.1.20, the trace of $\mathcal{O}_E/\mathcal{O}_K$ is surjective. We therefore set $u_2 = \sigma(c^{-1})$ and choose an element $u_1 \in \mathcal{O}_E$ so that $\mathrm{Tr}_\sigma(u_1\sigma(u_2)c) = 1$, which means $\langle u, u\rangle = \pi^i$. By Lemma 4.1.21, L is normal.

The previous example, therefore, serves as proof of the following fact — our first direct application of the separability hypothesis.

Lemma 4.4.21. If $\mathcal{O}_E/\mathcal{O}_K$ is a separable quadratic extension of discrete valuation rings, then every hyperbolic plane over $\mathcal{O}_E$ is normal.

To finish this section, we extend some useful properties of modular Hermitian lattices over non-Archimedean local fields from [Jac62, p.448] to lattices over discrete valuation rings and provide the details of the proof.

Proposition 4.4.22. Let L be a $\mathfrak{P}^i$-modular lattice over the quadratic extension of discrete valuation rings $\mathcal{O}_E/\mathcal{O}_K$.

1. The scale $\mathfrak{s}(L)$ of L is the ideal $\mathfrak{P}^i$.
2. The lattice $\varpi^j L$ is $\mathfrak{P}^{i+2j}$ modular for any $j \in \mathbb{Z}$.
3. For $\alpha \in K^\times$ with $\nu(\alpha) = j$, the scaled lattice $\alpha \cdot L$ is $\mathfrak{P}^{i+j}$-modular.
4. The orthogonal sum $L_1 \oplus L_2$ is $\mathfrak{P}^i$-modular if and only if L_1 and L_2 are $\mathfrak{P}^i$-modular.

Proof. Let $\mathcal{U}$ be an $\mathcal{O}_E$-basis for L. By assumption, any $x \in \mathcal{U}$ satisfies

$$\langle x, L\rangle = \varpi^i \mathcal{O}_E = \mathfrak{P}^i$$

so it suffices to show $\langle y, L\rangle \subseteq \mathfrak{P}^i$ for every $y \in L$. If $\mathcal{U}$ is an $\mathcal{O}_E$-basis of L, then $y \in L$ can be written as $y = \sum_{k=1}^n a_k x_k$ for some $a_k \in \mathcal{O}_E$ and $x_l \in \mathcal{U}$. But by the ultrametric inequality and Proposition 4.4.14, we have for each k that

$$\nu(\langle y, L\rangle) \geq \max\{\nu(a_k), \nu(\langle x_k, L\rangle)\} \geq \nu\left(\varpi^i\right) = i,$$

so that $\langle y, L\rangle \subseteq \mathfrak{P}^i$. To see that the second property holds, notice that multiplication by ϖ^j is an $\mathcal{O}_E$-module isomorphism, so that by Lemma 4.4.8, the vector y of $\varpi^j L$ is maximal if and only if $y = \varpi^j x$ for a maximal $x \in L$. It follows that

$$\langle y, \varpi^j L\rangle = \varpi^{2j}\mathfrak{P}^i = \mathfrak{P}^{i+2j}$$

for every maximal $y \in \varpi^j L$. The third property holds since $\alpha \cdot L$ is forgetfully equal to L as an $\mathcal{O}_E$-module, so that the lattices L and $\alpha \cdot L$ have the same set of maximal vectors. Since

$$\nu(\alpha\langle x, L\rangle) = \nu\left(\omega^j\right) + \nu(\langle x, L\rangle) = j + \nu(\langle x, L\rangle),$$

we have $\alpha\langle x, L\rangle = \mathfrak{P}^{i+j}$ for every maximal $x \in L$. For the last point, set $L = L_1 \oplus L_2$. Suppose $x = x_1 + x_2$ is maximal in L for $x_1 \in L_1$ and $x_2 \in L_2$, and without loss of generality, suppose that x_1 is non-zero. If L_1 is $\mathfrak{P}^i$-modular, there exists $y_1 \in L_1$ such that $\nu(\langle x_1, y_1\rangle) = i$. Since $\langle x, L_1 \oplus L_2\rangle \subseteq \mathfrak{P}^i + \mathfrak{P}^i = \mathfrak{P}^i$, we have that $\langle x, L_1 \oplus L_2\rangle = \mathfrak{P}^i$ for an arbitrary maximal vector x of L. Conversely, suppose that L is $\mathfrak{P}^i$-modular, and let $x_1 \in L_1$ be maximal. Consider x_1 as an element of L. Since $\varpi L = \varpi L_1 \oplus \varpi L_2$, and we have that x_1 is not in ϖL_1, it follows that x_1 is maximal in L. In this case, there exists $y_1 + y_2 \in L_1 \oplus L_2$ with

$$\nu(\langle x_1, y_1 + y_2\rangle) = \nu\left(\varpi^i\right) = i.$$

But $\langle x_1, y_2\rangle = 0$, so $\nu(\langle x_1, y_1\rangle) = i$. This means

$$\langle x_1, L_1\rangle = \langle x_1, y_1\rangle \mathcal{O}_E = \mathfrak{P}^i,$$

so that L_1 is $\mathfrak{P}^i$-modular. In the same way, L_2 is $\mathfrak{P}^i$-modular. This shows L_1 and L_2 are $\mathfrak{P}^i$-modular. ∎

4.5 Orthogonal bases of lattices

Over a field, a Hermitian form is diagonalizable. As we saw for a subnormal hyperbolic plane over the valuation ring of a ramified quadratic extension of $\mathbb{Q}_2$, the same does not hold in general for a Hermitian lattice. The problem is that the Hermitian form $\langle -, -\rangle$ on the lattice L takes values in a fractional ideal of $\mathcal{O}_E$. Therefore, we cannot guarantee that L, merely an $\mathcal{O}_E$-module, will be closed under the multiplication by Fourier coefficients $\frac{\langle x_i, x_j\rangle}{\langle x_j, x_j\rangle}$ contained in some arbitrary fractional ideal $\mathfrak{P}^j$, $j \in \mathbb{Z}$. There is, however, a near-diagonalization algorithm, which allows us to express a Hermitian lattice in terms of an orthogonal sum of rank 1 and rank 2 sublattices. In Lemma 4.5.1 and Theorem 4.5.2, we provide a quick proof of the diagonalization result mentioned by Jacobowitz in [Jac62, Prop 4.2] for the setting of lattices over DVRs. We then verify in Lemma 4.5.3 and Corollaries 4.5.4 and 4.5.5 that the sufficient conditions of [Jac62, Prop. 4.4] for the existence of a decomposition of the lattice L as an orthogonal sum of rank 1 sublattices are also sufficient for lattices over a DVR, thereby recovering the Gram-Schmidt algorithm of an inner product space for certain cases.

Lemma 4.5.1. Let L be a Hermitian lattice with basis $\{x_1, \ldots, x_n\}$.

1. If the coefficients $\alpha_i := \frac{\langle x_i, x_1\rangle}{\langle x_1, x_1\rangle}$ are in $\mathcal{O}_E$ for $2 \leq i \leq n$, then L decomposes as an orthogonal sum
$$L = \mathrm{span}_{\mathcal{O}_E}\{x_1\} \oplus W$$
where $W = \mathrm{span}_{\mathcal{O}_E}\{x_i - \alpha_i x_1 \mid 2 \leq i \leq n\}$.

2. If the coefficients
$$\beta_i := \frac{\langle x_i, x_1\rangle\langle x_2, x_2\rangle - \langle x_i, x_2\rangle\langle x_2, x_1\rangle}{\langle x_2, x_1\rangle\langle x_1, x_2\rangle - \langle x_1, x_1\rangle\langle x_2, x_2\rangle}$$
and
$$\gamma_i := \frac{\langle x_i, x_2\rangle\langle x_1, x_1\rangle - \langle x_i, x_1\rangle\langle x_1, x_2\rangle}{\langle x_2, x_1\rangle\langle x_1, x_2\rangle - \langle x_1, x_1\rangle\langle x_2, x_2\rangle}$$
are in $\mathcal{O}_E$ for all $3 \leq i \leq n$, then L decomposes as an orthogonal sum
$$\mathrm{span}_{\mathcal{O}_E}\{x_1, x_2\} \oplus W$$
where $W = \mathrm{span}_{\mathcal{O}_E}\{x_i + \beta_i x_1 + \gamma_i x_2 \mid 3 \leq i \leq n\}$.

In particular, the submodule W of L is finitely generated.

Proof. We can check orthogonality by straightforward symbolic manipulations. The above expressions generate all of L, as, by adding $\alpha_i x_1$ to $x_i - \alpha_i x_1$ in the first case, or subtracting $\beta_i x_1 + \gamma_i x_2$ from $x_i + \beta_i x_1 + \gamma_i x_2$ in the second case, we recover the generators x_i of L. These operations are allowed since all the relevant elements belong to $\mathcal{O}_E L = L$ by assumption that α_i or β_i and γ_i are in $\mathcal{O}_E$. ∎

We now show that one of the cases 1 and 2 of the previous Lemma 4.5.1 always occurs. This yields an algorithm for decomposing a lattice into an orthogonal sum of rank 1 and rank 2 sublattices.

Theorem 4.5.2 (Generalized Gram-Schmidt Algorithm)**.** *If L is a Hermitian lattice, then L can be written as an orthogonal sum of rank* 1 *and rank* 2 *sublattices.*

Proof. We proceed by induction on $n = \mathrm{rk}(L)$. We see that the statement is trivially true for $n \leq 2$. We can suppose that $\mathrm{rk}(L) \geq 3$ and that the statement of the theorem is true for any lattice L' with rank strictly less than n. Choose a basis $\{x_1, \ldots, x_n\}$ for L and consider the set

$$S = \{\nu(\langle x_i, x_j \rangle) \mid 1 \leq i \leq j \leq n.\}.$$

Since S is finite, we have that S attains its minimum for some pair $\{x_i, x_j\}$ of basis vectors of L. If $i = j$, then

$$\nu(\langle x_k, x_i \rangle) \geq \nu(\langle x_i, x_i \rangle)$$

for all $1 \leq k \leq n$. In this case, the coefficient $\frac{\langle x_k, x_i \rangle}{\langle x_i, x_i \rangle}$ satisfies

$$\nu\left(\frac{\langle x_k, x_i \rangle}{\langle x_i, x_i \rangle}\right) = \nu(\langle x_k, x_i \rangle) - \nu(\langle x_i, x_i \rangle) \geq 0$$

and is therefore contained in $\mathcal{O}_E$. Applying Lemma 4.5.1, we can write

$$L = \mathrm{span}_{\mathcal{O}_E}\{x_i\} \oplus W$$

for a finitely-generated submodule W of L. Since $\mathcal{O}_E$ is a principal ideal domain, W is free. By the structure theorem for finitely generated modules over a principal ideal domain, the decomposition of L given above implies that W has rank $n - 1$, so that applying the induction hypothesis to W, L is an orthogonal sum of lines and planes.

Now, suppose that the case above does not apply, so there are distinct indices $i \neq j$ such that for all $1 \leq k \leq n$, we have $\nu(\langle x_k, x_k \rangle) > \nu(\langle x_i, x_j \rangle)$, and that $\nu(\langle x_i, x_j \rangle)$ is minimal in S. Without loss of generality, write $i = 1$, $j = 2$, and consider the coefficient β_k, for $3 \leq k \leq n$. First, note that the denominator of each β_k satisfies

$$\begin{aligned}\nu(\langle x_2, x_1 \rangle\langle x_1, x_2 \rangle - \langle x_1, x_1 \rangle\langle x_2, x_2 \rangle) &= \min\{\nu(\langle x_2, x_1 \rangle\langle x_1, x_2 \rangle), \nu(\langle x_1, x_1 \rangle\langle x_2, x_2 \rangle)\} \\ &= \nu(\langle x_2, x_1 \rangle\langle x_1, x_2 \rangle)\end{aligned}$$

because of the choice of x_1, x_2 and Lemma 3.4.6. Then, looking at the numerator of β_k, we get

$$\begin{aligned}\nu(\langle x_k, x_1 \rangle\langle x_2, x_2 \rangle - \langle x_k, x_2 \rangle\langle x_2, x_1 \rangle) &\geq \min\{\nu(\langle x_k, x_1 \rangle\langle x_2, x_2 \rangle), \nu(\langle x_k, x_2 \rangle\langle x_2, x_1 \rangle\} \\ &\geq \nu(\langle x_2, x_1 \rangle\langle x_1, x_2 \rangle)\end{aligned}$$

since $\nu(\langle x_1, x_2\rangle) = \nu(\langle x_2, x_1\rangle)$ is minimal in S. Combining the two inequalities, we must have that $\nu(\beta_k)$ is non-negative. Hence, β_k is contained in $\mathcal{O}_E$. Similarly, so is γ_k. Therefore, the hypotheses of Lemma 4.5.1 hold and we obtain an orthogonal decomposition of $L \cong (x_1, x_2) \oplus W$. As in the first case, we apply the induction hypothesis to W to finish the proof. ∎

We now verify that the proofs of the following result and its two corollaries go through as written in [Jac62, Prop 4.4] when extended to the setting of Hermitian lattices over discrete valuation rings.

Lemma 4.5.3. Let L be a modular Hermitian lattice. If L has odd rank, then L has an orthogonal basis.

Proof. If $\mathrm{rk}(L) = 1$, there is nothing to prove. Suppose $n = \mathrm{rk}(L) \geq 3$, with n odd, and let L be $\mathfrak{P}^i$-modular, for some $i \in \mathbb{Z}$.

By Theorem 4.5.2, L decomposes into an orthogonal sum of lines and planes, and since L has odd rank, an orthogonal decomposition for L must have a rank 1-summand N. If this decomposition also admits a rank 2 summand P, we have

$$L \cong P \oplus N \oplus M$$

for some submodule $M \subseteq L$ of rank $n - 3$ (possibly 0). If we can show that P is an orthogonal sum of two rank 1 submodules, we will be able to apply the induction hypothesis to $P^\perp = N \oplus M$ and get that L itself is an orthogonal sum of lines. Therefore, it is enough to show $P \oplus N$ admits a decomposition into rank 1 submodules.

Write $N = \mathrm{span}\{x_1\}$ and $P = \mathrm{span}\{x_2, x_3\}$. By Proposition 4.4.22, we have that N and P are $\mathfrak{P}^i$-modular. By the proof of Theorem 4.5.2, we can assume that $\nu(\langle x_2, x_2\rangle) > \nu(\langle x_2, x_3\rangle)$ and $\nu(\langle x_3, x_3\rangle) > \nu(\langle x_2, x_3)$. Since L is $\mathfrak{P}^i$-modular, we can further assume that $\nu(\langle x_1, x_1\rangle) = i$, as x_1 is a basis vector of the rank 1 $\mathfrak{P}^i$-modular lattice N, and that by rescaling, $\langle x_2, x_3\rangle = \varpi^i$ for a uniformizer $\varpi \in \mathfrak{P}$. Then, set

$$a = -\frac{\langle x_3, x_3\rangle}{\varpi^i}$$

and

$$b = \frac{1}{\langle x_1, x_1\rangle}\left(\frac{\langle x_2, x_2\rangle\langle x_3, x_3\rangle}{\varpi^i} - \sigma(\varpi)^i\right).$$

Looking at the valuation of the numerator and denominator of a, we see that a is contained in $\mathfrak{P}$. Since

$$\nu\left(\frac{\langle x_2, x_2\rangle\langle x_3, x_3\rangle}{\varpi^i}\right) > i$$

but

$$\nu\left(\sigma(\varpi)^i\right) = i,$$

we have by Lemma 3.4.3 that $\nu\left(\langle x_1, x_1\rangle b\right) = i$. Since $\nu\left(\langle x_1, x_1\rangle\right)$ is already equal to i, this means b is invertible. Now, we apply Corollary 3.4.6 and obtain that

$$\det\left(\begin{bmatrix} 1 & 1 & 0 \\ b^{-1}a & 1 & 0 \\ b^{-1} & 0 & 1 \end{bmatrix}\right) = 1 - b^{-1}a$$

is a unit of $\mathcal{O}_E$. This means that

$$y_1 = x_1 + b^{-1}\left(ax_2 + x_3\right), \quad y_2 = x_1 + x_2, \quad \text{, and} \quad y_3 = x_3$$

form a basis for $P \oplus N$, and seeing as $b^{-1}a$ is not equal to 1, y_1 is not an $\mathcal{O}_E$-linear combination of y_2 and y_3. So, if $P' = \text{span}_{\mathcal{O}_E}\{y_2, y_3\}$ and $N' = \text{span}_{\mathcal{O}_E}\{y_1\}$, we have that our lattice $P \oplus N$ is the direct sum $P' \oplus N'$. But $\nu\left(\langle y_2, y_2\rangle\right) = i$. So, we can apply Lemma 4.5.1 to extend y_2 to an orthogonal basis of P'. This shows $P \oplus N = P' \oplus N'$ has an orthogonal basis, as required. ∎

Corollary 4.5.4. If L is normal, then L has an orthogonal basis.

Proof. We can assume that the normal lattice L has even rank by the Lemma. By assumption, L contains a vector x with $\nu\left(\langle x, x\rangle\right) = i$. By the proof of Theorem 4.5.2, $N = \text{span}_{\mathcal{O}_E}\{x\}$ is a rank 1 orthogonal summand of L. We can apply the Lemma to the lattice of odd rank $N^\perp$ to obtain an orthogonal basis for $N^\perp$. Since $L \cong N \oplus N^\perp$, L has an orthogonal basis. ∎

Corollary 4.5.5. If L is unimodular and there exists an element $a \in \mathcal{O}_E$ such that $\nu(\text{Tr}_\sigma(a)) = 0$, then L has an orthogonal basis.

Proof. We show that any unimodular plane P with the given properties has an orthogonal basis, which will allow us to apply induction to a Gram-Schmidt decomposition of L. We can assume P is subnormal, so that for a basis $\{x_1, x_2\}$ of P, we have that $\nu\left(\langle x_i, x_i\rangle\right) > 0$ and that by rescaling, $\langle x_1, x_2\rangle = 1$.

Consider the element $ax_1 + x_2 \in P$. We have

$$\begin{aligned}\langle ax_1 + x_2, ax_1 + x_2\rangle &= a\sigma(a)\langle x_1, x_1\rangle + a\langle x_1, x_2\rangle + \sigma(a)\sigma(\langle x_1, x_2\rangle) + \langle x_2, x_2\rangle \\ &= a\sigma(a)\langle x_1, x_1\rangle + \mathrm{Tr}_\sigma(a) + \langle x_2, x_2\rangle.\end{aligned}$$

But we know that $\nu\left(a\sigma(a)\langle x_1, x_1\rangle\right) \geq \nu\left(\langle x_1, x_1\rangle\right) > 0$ and $\nu(\langle x_2, x_2\rangle) > 0$.

By the strengthening of the ultrametric inequality of Lemma 3.4.6,

$$\nu\left(\langle ax + y, ax + y\rangle\right) = \nu\left(\mathrm{Tr}_\sigma(a)\right) = 0,$$

which is a contradiction with P subnormal. Therefore, P is a normal lattice, to which we can apply Corollary 4.5.4. ∎

To finish this section, we provide information on the determinant of the Gram matrix of a modular Hermitian lattice.

Lemma 4.5.6. If L is a $\mathfrak{P}^j$-modular Hermitian lattice of rank n, then $\nu\left(\det(L)\right) = jn$.

Proof. We check that $\nu\left(\det(L)\right) = jn$ for a $\mathfrak{P}^j$-modular lattice L. Since ν satisfies $\nu(xy) = \nu(x) + \nu(y)$, L is an orthogonal sum of lines and planes by the Generalized Gram-Schmidt Algorithm 4.5.2, and a $\mathfrak{P}^j$-modular lattice with an orthogonal basis satisfies the desired expression, it suffices to check the case that L is a $\mathfrak{P}^j$-modular plane. Let $\{x, y\}$ be a basis for L, and let ϖ be a uniformizer for $\mathcal{O}_E$. By Lemma 4.5.4, we can assume that L is subnormal, so that $\langle x, x\rangle$ and $\langle y, y\rangle$ are contained in $\mathfrak{P}^{j+1}$. Since L is $\mathfrak{P}^j$-modular and x is a basis vector for L, we must have that there exists $a, b \in \mathcal{O}_E$ such that $\langle x, ax + by\rangle = \varpi^j$. Since $a\langle x, x\rangle$ is contained in $\mathfrak{P}^{j+1}$ and $\nu\left(\langle x, ax + by\rangle\right) = j$, we apply Corollary 3.4.6 to get $\nu\left(b\langle x, y\rangle\right) = j$. Since L is $\mathfrak{P}^j$-modular, this implies that b is a unit and $\langle x, y\rangle$ has valuation j. By Lemma 4.1.21, σ is an isometry of ν, and we have $\nu\left(\langle y, x\rangle\right) = \nu((\langle x, y\rangle) = j$. Consequently, if H is the Gram matrix of L, we have

$$\begin{aligned}\nu\left(\det(H)\right)) &= \nu\left(\langle x, x\rangle\langle y, y\rangle - \langle x, y\rangle\langle y, x\rangle\right) \\ &= \nu\left(\langle x, y\rangle\langle y, x\rangle\right) \\ &= 2\nu\left(\langle x, y\rangle\right) = 2j,\end{aligned}$$

by applying Corollary 3.4.6 where required. ∎

Corollary 4.5.7. Let L be a $\mathfrak{P}^j$-modular lattice. Then L is non-degenerate.

Proof. Let $m \in \mathrm{rad}(L)$, and let H be a Gram matrix for L. We know from the Lemma that the determinant of H is non-zero, so that considering the action of H on all of $K \otimes L$, we have $\ker(H) = \{0\}$. Because L contains a K-basis for $K \otimes L$, any K-linear map on $K \otimes L$ is determined by its action on L. This means that if the assignment $x \mapsto x^T H m^\sigma$ is 0 for all $x \in L$ then this map induces the zero element of the dual space of $K \otimes L$. So, this implies that Hm^σ is the zero matrix and m^σ is in the kernel of H. This kernel is $\{0\}$, so $m = m^\sigma = 0$ and $\mathrm{rad}(L) = \{0\}$. ∎

4.6 The Jordan splitting of a Hermitian lattice

We would now like to refine the orthogonal decomposition given by the Gram-Schmidt algorithm of Theorem 4.5.2 to one reflecting the structure of the modular components of L. In this section, we introduce and describe the main invariant reflecting the geometry of a Hermitian lattice L, the Jordan decomposition of L. The Jordan decomposition of L describes L as an orthogonal sum of distinct $\mathfrak{P}^j$-modular components, and is unique up to isometry. We extend with minor modifications the ideas and definitions of [Jac62] from the setting of valuation rings over non-Archimedean local fields to that of discrete valuation rings.

If L is a Hermitian lattice and j is any integer, we will first examine the subset defined in [Jac62, p.449] by

$$L^j := \{x \in L \mid \langle x, L\rangle \subseteq \mathfrak{P}^j\}.$$

We will prove the subset L^j forms a sublattice of L, and then demonstrate some of its properties. Then, for every relevant $j \in \mathbb{N}$, we will identify the $\mathfrak{P}^j$-modular components of L as distinguished sublattices contained in L^j, and show that any such distinguished sublattice of L^j corresponds to an orthogonal summand of a Jordan decomposition for L. Since the definition of the sublattice L^j does not depend on a choice of Jordan splitting for L, this approach will yield the desired uniqueness result.

Lemma 4.6.1. Let L be a Hermitian lattice, and let $j \in \mathbb{Z}$. Then L^j is a sublattice of L.

Proof. Since L is a free module of finite rank over the principal ideal domain $\mathcal{O}_E$, every submodule of L is also free of finite rank. It therefore suffices to show that L^j is an $\mathcal{O}_E$-submodule of L. Recall that the fractional ideal $\mathfrak{P}^j$ is by definition an $\mathcal{O}_E$-submodule of the field of fractions E of $\mathcal{O}_E$. If $x, y \in L^j$ and $a \in \mathcal{O}_E$, it follows from the previous sentence that

$$\langle x+y, L\rangle = \langle x, L\rangle + \langle y, L\rangle \subseteq \mathfrak{P}^j + \mathfrak{P}^j = \mathfrak{P}^j,$$

and

$$\langle ax, L\rangle = a\langle x, L\rangle \subseteq a\mathfrak{P}^j \subseteq \mathfrak{P}^j.$$

This shows $x+y$ and ay are elements of L^j, so that L^j is an $\mathcal{O}_E$-submodule of L. ∎

The following technical result was stated for Hermitian lattices over non-Archimedean local fields in [Jac62, p.449]. Our proof is valid in the more general setting of lattices over discrete valuation rings. For the proof of the third item, we follow that in [OMe57, p.160].

Lemma 4.6.2. Let L be a Hermitian lattice. Then:

1. $\mathfrak{s}(L^j) \subseteq \mathfrak{P}^j$.
2. $(L \oplus M)^j = L^j \oplus M^j$ for any Hermitian lattice M.
3. If M is a $\mathfrak{P}^j$-modular sublattice of L, then M is an orthogonal summand of L if and only if $M \subseteq L^j$.

Proof. For $x \in L^j$, we have $\langle x, L^j\rangle \subseteq \langle x, L\rangle \subseteq \mathfrak{P}^j$, which gives the first property. If $x \in L$ and $y \in M$ satisfy $\langle x+y, L\oplus M\rangle \subseteq \mathfrak{P}^j$, then we have $\langle x, L\rangle \subseteq \langle x, L\oplus M\rangle \subseteq \mathfrak{P}^j$, so that $x \in L^j$. By a similar argument, we have $y \in M^j$. This means that $x+y \in L^j \oplus M^j$. Conversely, for $x \in L^j$ and $y \in M^j$, we have by orthogonality that

$$\langle x+y, L\oplus M\rangle = \langle x, L\rangle + \langle y, M\rangle \subseteq \mathfrak{P}^j + \mathfrak{P}^j = \mathfrak{P}^j.$$

For the third property, write $L = M \oplus M^\perp$ for the $\mathfrak{P}^j$-modular sublattice M. Then, for $x \in M$, we have

$$\langle x, L\rangle = \langle x, M \oplus M^\perp\rangle \subseteq \mathfrak{P}^j + 0 = \mathfrak{P}^j$$

by Lemma 4.4.22. Conversely, let M be a $\mathfrak{P}^j$-modular submodule of L contained in L^j, let $\mathcal{O}_E$ be the field of fractions of $\mathcal{O}_E$, and let ϖ be a uniformizer for $\mathcal{O}_E$. Consider the Hermitian space $K \otimes M \subseteq K \otimes L$. Since M is modular, we have that M is non-degenerate by Corollary 4.5.7. Since M generates $K \otimes M$, $K \otimes M$ must also be non-degenerate. So, by Lemma 4.3.9, $K \otimes L = K \otimes M \oplus (K \otimes M)^{\perp}$. Set $N = (K \otimes M)^{\perp} \cap L$. We claim that $L = M \oplus N$. The $M \cap N = \{0\}$ and $M \perp N$, as these properties are preserved by restriction. Let $x \in L$, and write $x = rm + sn$ for some $m \in M, n \in N$, and $r, s \in K$. Since the maximal vectors of M generate $K \otimes M$, we can assume m is maximal. Since n is orthogonal to M, we then have that

$$\begin{aligned}\langle x, M\rangle &= \langle rm, M\rangle \\ &= r\langle m, M\rangle = r\mathfrak{P}^j.\end{aligned}$$

Since M is contained in L^j and $x \in L$, this implies that $r \in \mathcal{O}_E$, as otherwise, $\langle x, M\rangle$ would be a strictly larger fractional ideal. Therefore, $rm \in M$, and $x - rm \in L$. This shows $L = M + N$, which finishes the proof. ∎

To enumerate the summands in a particular orthogonal decomposition of L, we will write

$$L = \bigoplus_{i=1}^{n} L_i,$$

for submodules L_i, $1 \leq i \leq n$. This will prevent any interference with the notation for the integer-indexed modular components L^j of L.

We now introduce the Jordan splitting (or Jordan decomposition) of a Hermitian lattice L. In [Jac62], the Jordan splitting is the main invariant for Hermitian lattices over valuation rings of non-Archimedean local fields.

Definition 4.6.3. Consider a Hermitian lattice $L = \bigoplus_{i=1}^{r} L_i$, with each L_i a modular sublattice. If $\mathfrak{s}(L_r) \subsetneq \mathfrak{s}(L_{r-1}) \subsetneq \cdots \subsetneq \mathfrak{s}(L_1)$, then the decomposition

$$L = \bigoplus_{i=1}^{r} L_i$$

is called a Jordan splitting for L. We sometimes use the notation $[L]$ for a given Jordan decomposition of L.

Since each L_i is modular, we can define a function $s\colon \{1,\dots,r\} \to \mathbb{Z}$ by $s(i)$ by defining $s(i)$ such that $\mathfrak{s}(L_i) = \varpi^{s(i)}\mathcal{O}_E$, i.e., $s(i)$ identifies the fractional ideal $\mathfrak{P}^{s(i))}$ corresponding to $\mathfrak{s}(L_i)$. We see that $s(1) < s(2) < \cdots < s(r)$ for any Jordan splitting, so that s is injective.

Definition 4.6.4. We say that Jordan splittings

$$L = \bigoplus_{i=1}^{r} L_i \text{ and } K = \bigoplus_{j=1}^{s} K_j$$

for the Hermitian lattices L and K have the same type if $r = s$, and for each $1 \leq i \leq r$ we have

1. $\mathfrak{s}(L_i) = \mathfrak{s}(K_i)$,
2. $\mathrm{rk}(L_i) = \mathrm{rk}(K_i)$,
3. L_i and K_i are both normal or both subnormal.

Before proving the uniqueness theorem, we give three useful and deeply technical results on the sublattice L^j. We provide our own detailed proofs of Lemmas 4.6.5, 4.6.6, and 4.6.7, which are inspired by sketches of proofs given for quadratic forms in [OMe73, p.244].

Lemma 4.6.5. Suppose L is $\mathfrak{P}^i$-modular. The lattice L^j is $\mathfrak{P}^j$-modular if and only if $j = i$. If j is not equal to i, then L^j is $\mathfrak{P}^k$-modular for some $k \in \mathbb{Z}$, with $\mathfrak{P}^k \subsetneq \mathfrak{P}^j$.

Proof. First suppose that $j = i$. Then, Proposition 4.4.22 says $\mathfrak{s}(L) = \mathfrak{P}^i$ by the assumption that L is $\mathfrak{P}^i$-modular. This means $L^j = L^i = L$, so L^j is $\mathfrak{P}^j$-modular.

If $j < i$, we have $L^j \subseteq L$, so $\mathfrak{s}(L^j) \subseteq \mathfrak{s}(L) = \mathfrak{P}^i$ by Proposition 4.4.22, with $\mathfrak{P}^i \subsetneq \mathfrak{P}^j$. This means that L^j is not $\mathfrak{P}^j$-modular, as the scale of any maximal line in L^j fails to generate $\mathfrak{P}^j$. But since every line of L has scale contained in $\mathfrak{P}^j$, it follows from the definition of L^j that L^j coincides with all of L. Therefore, $L^j = L$ is $\mathfrak{P}^i$-modular, with $\mathfrak{P}^i \subsetneq \mathfrak{P}^j$.

Suppose now that $j > i$. We claim that $L^j = \mathfrak{P}^{j-i}L$. If $y \in L^j$, then there exists a maximal $x \in L$ such that $y = ax$, for some $a \in E^\times$. Since x is maximal and L is

$\mathfrak{P}^i$-modular we have

$$a\mathfrak{P}^i = a\langle x, L\rangle = \langle ax, L\rangle = \langle y, L\rangle = \mathfrak{P}^j$$

so that $a \in \mathfrak{P}^{j-i}$. This means $y \in \mathfrak{P}^{j-i}L$. Conversely, a vector $x = \varpi^{j-i}y$ in $\mathfrak{P}^{j-i}L$ satisfies $\langle x, L\rangle = \varpi^{j-i}\langle y, L\rangle \subseteq \mathfrak{P}^{j-i+i} = \mathfrak{P}^j$, so that $\mathfrak{P}^{j-i}L \subseteq L^j$. This proves the claim.

Now, let y be maximal in $L^j = \mathfrak{P}^{j-i}L$. By this equality, $y = ax$ for some $a \in \mathfrak{P}^{j-i}$ and $x \in L$. But since y is maximal in $\mathfrak{P}^{j-i}L$, y is not an element of $\mathfrak{P}^{j-i+1}L$, so that $x = a^{-1}y$ is not an element of $\mathfrak{P}^{i-j+j-i+1}L = \mathfrak{P}L$. Therefore, x is maximal in the $\mathfrak{P}^i$ modular lattice L, and in particular, a is an element of $\mathfrak{P}^{j-i}$ but not of $\mathfrak{P}^{j-i+1}$. Since $\mathfrak{P}^{j-i}$ is an integral ideal and therefore invariant under the involution σ, we have

$$\langle y, \mathfrak{P}^{j-i}L\rangle = a\mathfrak{P}^{j-i}\langle x, L\rangle = \mathfrak{P}^{2(j-i)+i} = \mathfrak{P}^{2j-i}.$$

Now y is an arbitrary maximal vector of L^j, and this shows L^j is $\mathfrak{P}^{2j-i}$ modular. ∎

Lemma 4.6.6. Let $i \in \mathbb{Z}$, and consider the lattice L with Jordan decomposition

$$L = \bigoplus_{j=1}^{r} L_j.$$

Then the scale $\mathfrak{s}(L^i)$ of L^i is equal to $\mathfrak{P}^i$ if and only if there is a $\mathfrak{P}^i$-modular component in the given Jordan decomposition of L. If there is no such component, then $\mathfrak{s}(L^i)$ is properly contained in $\mathfrak{P}^i$.

Proof. The Jordan decomposition of L induces a Jordan decomposition

$$L^i = \bigoplus_{j=1}^{r} L^i_j$$

by Lemma 4.6.2. In any case, the scale of L^i is contained in $\mathfrak{P}^i$ by definition.

Consider a component L_j in the Jordan decomposition of L.

If L_j is $\mathfrak{P}^i$-modular, then $L^i_j = L_j$, by definition of L^i_j. So, L^i_j is $\mathfrak{P}^i$-modular, which is a sufficient condition for the scale $\mathfrak{s}(L^i)$ to be equal to $\mathfrak{P}^i$.

If L_j is not $\mathfrak{P}^i$-modular, then L^i_j is not $\mathfrak{P}^i$-modular but $\mathfrak{P}^k$-modular for some $\mathfrak{P}^k \subsetneq \mathfrak{P}^i$ by Lemma 4.6.5. Hence, the scale of L^i_j is properly contained in $\mathfrak{P}^i$.

Therefore, if there is no $\mathfrak{P}^i$-modular component in the Jordan splitting of L, the scale of L^i_j is properly contained in $\mathfrak{P}^i$ for all $1 \leq j \leq r$. But, by Proposition 4.4.14, the scale of $L^i = \oplus_{j=1}^r L^i_j$ is equal to

$$\mathfrak{s}(\oplus_{j=1}^r L^i_j) = \sum_{i=1}^r \mathfrak{s}(L^i_j),$$

where $\mathfrak{s}(L^i_j)$ is properly contained in $\mathfrak{P}^i$ for all j. This means $\mathfrak{s}(L^i)$ is itself properly contained in $\mathfrak{P}^i$ whenever L has no $\mathfrak{P}^i$-modular components. ∎

Lemma 4.6.7. If L_k is a $\mathfrak{P}^i$-modular component in some Jordan decomposition of L, then L_k is the first component in some Jordan decomposition of L^i.

Proof. Consider the orthogonal sum $L^i = \bigoplus_{j=1}^r L^i_j$ induced by taking the power L^i of the Jordan decomposition $L = \bigoplus_{j=1}^r L_j$, and suppose that L_k is the $\mathfrak{P}^i$-modular component of the Jordan decomposition of L. By Lemma 4.6.5, we have that L^i_k is $\mathfrak{P}^i$-modular because $L_k = L^i_k$ is $\mathfrak{P}^i$-modular, and if $j \neq k$, then L^i_k is $\mathfrak{P}^\ell$-modular with $\mathfrak{P}^\ell \subsetneq \mathfrak{P}^j$. Therefore, up to reordering of the summands, $\bigoplus_{j=1}^r L^i_j$ is a Jordan splitting of L^i with first component L_k. ∎

We therefore have our main theorem for the section, generalized from that stated and proved in [OMe73, p.244] for quadratic forms. We provide the minor modifications necessary for the setting of Hermitian lattices over discrete valuation rings and verify that the proof given above goes through without difficulty.

Theorem 4.6.8. *Let L be a Hermitian lattice over the discrete valuation ring $\mathcal{O}_E$. Then L has a Jordan decomposition, and the Jordan type of the decomposition is uniquely determined by the isometry class of L.*

Proof. Existence is given by the generalized Gram-Schmidt algorithm of Theorem 4.5.2, as an orthogonal sum of modular lines and planes is a Jordan splitting for L (of course, when the summands are ordered to satisfy the inclusion condition).

To prove uniqueness, we modify the proof given for quadratic forms in [OMe73]. Let

$$L = \bigoplus_{i=1}^r L_i \text{ and } L = \bigoplus_{j=1}^s M_j$$

be two Jordan splittings.

Suppose that there is a $\mathfrak{P}^i$-modular component in the first Jordan decomposition. Then, by Lemma 4.6.6, the scale $\mathfrak{s}(L^i)$ of L^i is equal to $\mathfrak{P}^i$, so that applying the Lemma to the second Jordan decomposition, one of the M_j must be $\mathfrak{P}^i$-modular. By symmetry, $\mathfrak{P}^i$-modular components of the first and second decomposition are in one-to-one correspondence, so we must have $r = s$ and $\mathfrak{s}(L_j) = \mathfrak{s}(M_j)$ for all $1 \leq j \leq r$.

It remains to show that $\mathrm{rk}(L_j) = \mathrm{rk}(M_j)$ for all $1 \leq j \leq r$, and that L_j is normal if and only if K_j is normal. Let us use a clever trick. By Lemma 4.6.7, the $\mathfrak{P}^j$-modular components L_k and M_k of L each appear as the first component in some Jordan decomposition of L^j. But L_k and M_k have the same rank as $\mathcal{O}_E$-submodules of L and as $\mathcal{O}_E$-submodules of L^j. It is therefore sufficient to show that the rank of the first component of any given Jordan decomposition is an invariant of the lattice. We will then apply this fact to L^j, to show that $\mathrm{rk}(L_k) = \mathrm{rk}(M_k)$. Therefore, consider the first components L_1 and M_1, of two Jordan decompositions of L. We can further assume that L_1 and M_1 are both unimodular, since multiplication by a scalar preserves the rank of submodules.

Denote by V the E-vector space

$$E \otimes_{\mathcal{O}_E} \bigoplus_{i=1}^{r} L_i = E \otimes_{\mathcal{O}_E} L = E \otimes_{\mathcal{O}_E} \bigoplus_{j=1}^{s} M_j.$$

Since the tensor product distributes over the direct sum of submodules and the extension of scalars preserves orthogonality relations, we can consider the subspaces $V_L := E \otimes_{\mathcal{O}_E} L_1$ and $V_M := E \otimes_{\mathcal{O}_E} M_1$ of V, and the canonical projection $V \to V_M$. The restriction $\varphi\colon V_L \to V_M$ of the canonical projection is a linear map. As E is the field of fractions of $\mathcal{O}_E$, we have $\dim(V_L) = \mathrm{rk}(L_1)$ and $\dim(V_M) = \mathrm{rk}(M_1)$ by Lemma 4.4.5. Therefore, if we can show that φ is an injection, we will have $\mathrm{rk}(L_1) = \dim(V_L) \leq \dim(V_M) = \mathrm{rk}(M_1)$. By a symmetrical argument, we will then have $\mathrm{rk}(M_1) \leq \mathrm{rk}(L_1)$, forcing equality. Suppose towards contradiction that there is a vector $x \in V_L$ such that $\varphi(x) = 0$. After multiplying by an appropriate scalar in $E^\times$, we can take x to be a maximal vector in the lattice L_1. By the Jordan decomposition $L = \oplus_{j=1}^{r} M_j$, we can write $x = y + z$ for a unique $y \in M_1$ and $z \in M_1^\perp = \oplus_{j=2}^{r} M_j$. In particular, $\varphi(x) = y = x - z$. Say $x' \in L_1$ is also written in a unique way as $x' = y' + z'$

for $y' \in M_1$ and $z' \in M_1^\perp$. We then obtain the identity

$$\langle \varphi(x), \varphi(x') \rangle = \langle x - z, x' - z' \rangle \equiv \langle x, x' \rangle \qquad (\mathrm{mod}\ \mathfrak{P}),$$

as $\langle M_1^\perp, L \rangle$ is contained in $\mathfrak{P}$ by definition of the Jordan decomposition. Since L_1 is unimodular and x is maximal, there is $w \in L_1$ with $\langle x, w \rangle = 1$. But as $\varphi(x) = 0$, we must have

$$1 = \langle x, w \rangle \equiv \langle \varphi(x), \varphi(w) \rangle \equiv 0 \pmod{\mathfrak{P}},$$

a contradiction. This means φ is injective, as desired.

Finally, suppose L_k is normal, and consider the map $\varphi\colon L_k \to M_k$ given above. Once again, we can assume that L_k and M_k are unimodular. But as above, we have

$$\langle x, x \rangle \equiv \langle \varphi(x), \varphi(x) \rangle \qquad (\mathrm{mod}\ \mathfrak{P}),$$

so that taking x to have $\langle x, x \rangle \in \mathcal{O}_E^\times$, we must have that $\langle \varphi(x), \varphi(x) \rangle$ is a generator of $\mathcal{O}_E$. Hence, $\mathfrak{n}(M_k) = \mathcal{O}_E$ and M_k is a normal sublattice of L. ∎

4.7 Isometry classes of lattices over separable extensions

We will now prove that every Hermitian lattice over a quadratic separable extension $\mathcal{O}_E/\mathcal{O}_K$ of discrete valuation rings admits an orthogonal basis. In the preceding section, we saw that the Jordan splitting of a lattice over a discrete valuation ring is uniquely determined by its isometry class. For a Hermitian lattice over a separable extension of discrete valuation rings, we will obtain a closely related sufficient condition for isometry: the Jordan type of L uniquely determines the isometry class of the lattice up to some norm group factors. We show that every modular Hermitian lattice over a separable quadratic extension of discrete valuation rings admits an orthogonal basis. Recall from Corollary 2.5.16 that the maximal ideal $\mathfrak{p}$ of $\mathcal{O}_K$ generates $\mathfrak{P}$ in $\mathcal{O}_E$, and in particular, that $\pi\mathcal{O}_E = \mathfrak{P}$ for a generator π of $\mathfrak{p}$. We extend the following result of [Jac62, Thm. 7.1] to our more general setting, taking care to account for the possibility that the norm map $\mathrm{N}_\sigma\colon \mathcal{O}_E^\times \to \mathcal{O}_K^\times$ fails to surject onto $\mathcal{O}_K^\times$ for an arbitrary separable extension of discrete valuation rings.

Proposition 4.7.1. Suppose $\mathcal{O}_E/\mathcal{O}_K$ is a quadratic separable extension of discrete valuation rings with standard involution σ, and let $\nu\colon \mathcal{O}_E \to \mathbb{Z}$ be the discrete valuation on $\mathcal{O}_E$. Let π be a uniformizer for both $\mathcal{O}_K$ and $\mathcal{O}_E$. The π^i-modular lattice L has an orthogonal basis. We have that L is represented by the Gram matrix

$$a_1\pi^i \oplus \cdots \oplus a_j\pi^i \oplus \cdots \oplus a_n\pi^i,$$

for some $a_j \in \mathcal{O}_K^\times$, $1 \leq j \leq n$, where the a_j can be chosen up to a norm of $\mathcal{O}_E^\times$. In other words, L has an orthogonal basis $\{x_1, \ldots, x_n\}$ such that

$$\langle x_j, x_j\rangle = a_j\pi^i,$$

for all $1 \leq j \leq n$.

Proof. Since $L = \pi^i \cdot \pi^{-i} \cdot L$, we can find an orthogonal basis for the unimodular lattice $\pi^{-i} \cdot L$ and rescale it by a factor of π^i. The rescaled basis will be an orthogonal basis for L. So, without loss of generality, let us assume L is unimodular (of rank n).

By Lemma 4.1.20, there exists an element $a \in \mathcal{O}_E$ that is a unit of $\mathcal{O}_E$ and has unit trace, so that $\nu(a) = \nu(\mathrm{Tr}_\sigma(a)) = 0$. We can apply Corollary 4.5.5 to obtain an orthogonal basis $x_1, \ldots, x_n$ for L. This shows the first part.

Let us now work in all generality, and suppose that L is π^i-modular. Consider the orthogonal basis $S = \{x_1, \ldots, x_n\}$ of L. Since L is π^i-modular, Proposition 4.4.22 says that we have that $\langle x_j, \mathcal{O}_E x_j\rangle = \pi^i\mathcal{O}_E$ for any $j \in \{1, \ldots, n\}$. But then

$$\nu(\langle x_j, rx_j\rangle) = \nu(\sigma(r)\langle x_j, x_j\rangle) \geq \nu(\langle x_j, x_j\rangle) \geq \nu\left(\pi^i\right) = i$$

for any $r \in \mathcal{O}_E$. By the assumption that L is π^i-modular, we must have $\nu(\langle x_j, x_j\rangle) = i$. We have that $\langle x_j, x_j\rangle = \sigma(\langle x_j, x_j\rangle)$, so the scalar $\langle x_j, x_j\rangle$ is fixed by σ. Since σ is the standard involution on the quadratic $\mathcal{O}_K$-algebra $\mathcal{O}_E$, we have by Example 4.1.16 that $\mathcal{O}_E^\sigma = \mathcal{O}_K$, and therefore that $\langle x_j, x_j\rangle \in \mathcal{O}_K$. Write $\langle x_j, x_j\rangle = a_j\pi^i$ for some $a_j \in \mathcal{O}_K^\times$. Let $b_j = a_j\,\mathrm{N}_\sigma(\alpha_j)$ for some $\alpha_j \in \mathcal{O}_E^\times$. We have that $\langle \alpha x_j, \alpha x_j\rangle = a_j\,\mathrm{N}_\sigma(\alpha_j) = b_j$. Since multiplication by a unit is an $\mathcal{O}_E$-module isomorphism, we can replace the basis vector x_j with $\alpha_j x_j$, $\alpha_j \in \mathcal{O}_E^\times$, so that $\langle \alpha_j x_j, \alpha_j x_j\rangle = b_j$. This change does not affect the orthogonality of the basis. Therefore,

$$L \cong \bigoplus_{j=1}^{n} b_j\pi^i,$$

where $b_j \in a_j\,\mathrm{N}_\sigma(\mathcal{O}_E^\times)$ can be chosen arbitrarily in $a_j\,\mathrm{N}_\sigma(\mathcal{O}_E^\times)$. ∎

Corollary 4.7.2. If $\mathcal{O}_E$ is a separable $\mathcal{O}_K$-algebra, every modular Hermitian lattice is normal.

Proof. By the Lemma, L is isometric to the normal lattice $a_1\pi^i \oplus \cdots \oplus a_n\pi^i$. ∎

We mention but do not prove the following result, which is well known in number theory. This result will allow us to show that over an unramified extension of local fields, every lattice is isometric to a direct sum of powers of a fixed uniformizer.

Lemma 4.7.3. Let E/K be an unramified extension of local fields. Then the restriction

$$\mathrm{N_{E/K}}\colon \mathcal{O}_E^\times \to \mathcal{O}_K^\times$$

of the norm map to the group of units of the valuation ring $\mathcal{O}_E$ is a surjection onto $\mathcal{O}_K^\times$.

Proof. See, for instance, [Se79, p.82]. ∎

Corollary 4.7.4. Let $i \in \mathbb{Z}$. Every $\mathfrak{P}^i$-modular lattice over an unramified extension of non-Archimedean local fields is isometric to

$$\pi^i \oplus \cdots \oplus \pi^i.$$

Proof. By Lemma 4.7.3, the norm map $\mathcal{O}_E \to \mathcal{O}_K$ is surjective, so a_j is a norm of $\mathcal{O}_E$. Therefore, we can replace each basis vector x_j for L with $\alpha_j x_j$ for some $\alpha_j \in \mathcal{O}_E^\times$ such that $\mathrm{N_{E/K}}(\alpha_j) = a_j^{-1}$. ∎

To finish the section, we diagonalize the Gram matrix of a hyperbolic plane over a quadratic separable extension of 2-adic valuation rings.

Example 4.7.5. Consider the unique unramified extension E of $\mathbb{Q}_2$ of degree 2, with valuation ring $\mathcal{O}_E$. As in Example 3.5.24, E is equal to $\mathbb{Q}_2(\zeta_3)$, where ζ_3 is a primitive 3^{rd} root of unity and the Galois involution on E is given by linearly extending

$$\zeta_3 \mapsto \zeta_3^2$$

over $K = \mathbb{Q}_2$. Let us choose 2 to be a uniformizer for K, and let ϖ be an arbitrary uniformizer for E, i.e., we have $|\varpi| = |2| = \frac{1}{2}$ but that ϖ is not necessarily in K.

We will compute an orthogonal basis for the hyperbolic plane

$$L = \begin{bmatrix} 0 & \varpi^i \\ \sigma(\varpi)^i & 0 \end{bmatrix}$$

over $\mathcal{O}_E^2$, for any given $i \in \mathbb{Z}$. First, we can write $\varpi^i = 2^i u$, for u a unit of $\mathcal{O}_E^\times$ not contained in $\mathcal{O}_K^\times$, so that the lattice

$$2^{-i}L = \begin{bmatrix} 0 & u \\ \sigma(u) & 0 \end{bmatrix}$$

is unimodular by Proposition 4.4.22. Denote by H the given Gram matrix for the unimodular $2^{-i}L$ and by H_L the $\mathfrak{P}^i$-modular lattice L. If e_1, e_2 denotes the standard basis of $\mathcal{O}_E^2$, we have that $\langle e_1, e_2\rangle = u$, so that we can set $x_1 := \frac{1}{u}e_1$ and $x_2 := e_2$ to get $\langle x_1, x_2\rangle = 1$. In particular, x_1, x_2 is a basis of $2^{-i}L$, but it is not orthogonal. In particular, both x_1 and x_2 are isotropic. Note that $\mathrm{Tr_{E/K}}(\zeta_3) = \zeta_3 + \sigma(\zeta_3) = -1$, by the identity $1 + \zeta_3 + \zeta_3^2 = 0$. We can therefore consider the vector $y_1 := \zeta_3 x_1 + x_2$ as in the proof of Corollary 4.5.5. Now, we can extend y_1 to a basis of $2^{-i}L$ with the vector $y_2 := x_2$. We can simply check directly that y_1 and y_2 are linearly independent over $\mathcal{O}_E$ and generate $2^{-i}L$. We will then be able to apply the Gram-Schmidt process to the basis $\{y_1, y_2\}$. The vector y_1 satisfies

$$\langle y_1, y_1\rangle = \begin{bmatrix} \frac{\zeta_3}{u} & 1 \end{bmatrix} \begin{bmatrix} 0 & u \\ \sigma(u) & 0 \end{bmatrix} \begin{bmatrix} \frac{\sigma(\zeta_3)}{\sigma(u)} \\ 1 \end{bmatrix} = \mathrm{Tr_{E/K}}(\zeta_3) = -1,$$

so we know the Fourier coefficient

$$\frac{\langle y_2, y_1\rangle}{\langle y_1, y_1\rangle} = -\langle y_2, y_1\rangle = -\zeta_3^2 = 1 + \zeta_3$$

belongs to $\mathcal{O}_E$. As in Theorem 4.5.2, the vector

$$y_2 - (1 + \zeta_3)y_1 = \frac{1}{u}e_1 - \zeta_3 e_2$$

extends y_1 to an orthogonal basis of $2^{-i}L$. In particular, we compute

$$\langle y_1, y_1\rangle = -1 \quad \text{and} \quad \langle y_2, y_2\rangle = 1,$$

by using the matrix representation H.

Now, we can diagonalize the matrix H with respect to the orthogonal basis $\{y_1, y_2\}$. By this, we mean providing an invertible matrix $T \in M_n(E)$ such that

$D = T^t H T^\sigma$ is a diagonal matrix with $D_{ii} = \langle y_i, y_i \rangle$. Note that since σ commutes with the action of matrices on vectors, a good guess for T is the (invertible) change of basis matrix with first and second columns y_1 and y_2 respectively. Explicitly,

$$T = \begin{bmatrix} \frac{\zeta_3}{u} & \frac{1}{u} \\ 1 & -\zeta_3 \end{bmatrix}.$$

We then compute

$$T^t H T^\sigma = \begin{bmatrix} \frac{\zeta_3}{u} & 1 \\ \frac{1}{u} & -\zeta_3 \end{bmatrix} \begin{bmatrix} 0 & u \\ \sigma(u) & 0 \end{bmatrix} \begin{bmatrix} \frac{\sigma(\zeta_3)}{\sigma(u)} & \frac{1}{\sigma(u)} \\ 1 & -\sigma(\zeta_3) \end{bmatrix} = \begin{bmatrix} -1 & 0 \\ 0 & 1 \end{bmatrix},$$

which is diagonal. Finally, we need to obtain a diagonalization for the original lattice L. But, all we did to obtain H from H_L was to multiply H_L by 2^{-i}. Since $2^i I$ commutes with every matrix in $M_n(E)$, we have

$$T^t H_L T^\sigma = \begin{bmatrix} -2^i & 0 \\ 0 & 2^i, \end{bmatrix}$$

so that the isometry T also diagonalizes H_L, and that an orthogonal basis for L is given by the vectors Te_1, Te_2.

4.8 The unitary group

We now examine the unitary group of a Hermitian lattice with the goal of describing its generators. In this section, we extend the terminology and some basic results of [BHM22], which are used to describe Hermitian lattices over non-Archimedean local fields, to Hermitian lattices over discrete valuation rings with minimal change. We also extend some terminology of Wall used to describe Hermitian spaces over fields to that of Hermitian lattices over discrete valuation rings. We set $\mathcal{O}_E/\mathcal{O}_K$ to be a quadratic extension of discrete valuation rings equipped with the standard involution σ, and we set e/k to be the corresponding extension of residue fields. Let $\mathcal{O}_E$ have field of fractions E and $\mathcal{O}_K$ have field of fractions K. We further assume that σ is non-trivial, so that σ extends to a K-automorphism of E.

Definition 4.8.1. Let V be a Hermitian space over E with $\dim(V) = n$. The unitary group $\mathrm{U}(V)$ of V is the group of isometries preserving the Hermitian form $\langle -, - \rangle$. In

other words,

$$\mathrm{U}(V) = \{T \in \mathrm{GL}(V) \mid \langle Tx, Ty\rangle = \langle x, y\rangle \text{ for all } x, y \in V\}.$$

If H is the Gram matrix of the Hermitian form $\langle -, -\rangle$ on V, we define

$$\mathrm{U}(n) = \{T \in \mathrm{GL}(V) \mid H = T^t H T^\sigma\}.$$

When $T \in \mathrm{U}(V)$ or $T \in \mathrm{U}(n)$, we say that T is an isometry of V.

If $T \in U(n)$, the unitary group of the Hermitian space with Gram matrix H, then

$$H = T^t H T^\sigma,$$

and taking determinants, we have

$$\mathrm{N}_\sigma(\det(T)) = 1.$$

Definition 4.8.2. Let V be a Hermitian space. Consider $t \in V \setminus \mathrm{rad}\,(V)$ and $\tau \in E^\times$ with $\langle t, t\rangle = \mathrm{Tr}_{\mathrm{E/K}}(\tau)$. The linear map

$$S_{t,\tau} \colon V \to V, \quad x \mapsto x - \langle x, t\rangle \tau^{-1} t$$

is called a symmetry of V with respect to t.

From the formula for a symmetry and the requirement that the defining vector is not in the radical of V, we see that we can think of symmetries as rank 1 isometries, in the sense that the fixed subspace of a symmetry has codimension 1. We now assume the Hermitian space V is non-degenerate. We quickly state and prove a Lemma for V, which will allow us to compute the determinant of a symmetry of V defined with respect to an isotropic vector.

Lemma 4.8.3. Let V be a non-degenerate Hermitian space over the field E, and let $t \in V$ be isotropic. Then, there exists $y \in V$ such that $\langle t, y\rangle = 1$. In this case, $W = \mathrm{span}\{t, y\}$ is non-degenerate.

Proof. If every $y \in V$ satisfies $\langle t, y\rangle = 0$, then $t \in \mathrm{rad}(V) = 0$. So, rescaling y if necessary, there exists $y \in V$ such that $\langle t, y\rangle = 1$. From the relation

$$\langle at + bx, ct + dx\rangle = \sigma(b)c + a\sigma(d) + b\sigma(d)\langle x, x\rangle$$

it follows that any $at + bx \in W \cap W^\perp$ is zero. ∎

We now give the details of the proof of the following result mentioned in [BHM22] on the properties of symmetries of a Hermitian space. In particular, we show that the symmetries of V can be characterized as the one-dimensional isometries of V.

Proposition 4.8.4. Let t, τ, and $S_{t,\tau}$ be as in Definition 4.8.2. Then

1. $S_{t,\tau}$ is an isometry of V.
2. If t is isotropic, then $\det(S_{t,\tau}) = 1$. Otherwise, $\det(S_{t,\tau}) = \frac{-\sigma(\tau)}{\tau}$.
3. The inverse of $S_{t,\tau}$ is $S_{t,\sigma(\tau)}$.
4. The symmetries of V are the non-trivial elements of $\mathrm{U}(V)$ which restrict to the identity on hyperplanes.

Proof. We can check by straightforward symbolic manipulations that $S_{t,\tau}$ is an isometry of V and that the formula for the inverse holds. Suppose now that t is anisotropic and non-zero. Then, the subspace $W := \mathrm{span}\{t\}$ of V is non-degenerate, so we apply Proposition 4.3.9 and obtain that $V \cong W \oplus W^\perp$. Looking at the formula for $S_{t,\tau}$, all vectors orthogonal to t (so $W^\perp$) are 1-eigenvectors for $S_{t,\tau}$. Then, $S_{t,\tau}(t) = \left(1 - \frac{\langle t,t\rangle}{\tau}\right) t = \frac{-\sigma(\tau)}{\tau}$. Therefore, $\det(S_{t,\tau}) = \frac{-\sigma(\tau)}{\tau}$. Suppose then that t is isotropic. By Lemma 4.8.3, we can find $x \in V$ such that $\langle t, x\rangle = 1$ and by Proposition 4.3.9, we have that $V = W \oplus W^\perp$ for $W = \mathrm{span}\{t, x\}$. Then, $W^\perp \subseteq \mathrm{span}\{t\}^\perp$, so that $W^\perp$ is contained the 1-eigenspace for $S_{t,\tau}$. We then see from the formula for $S_{t,\tau}$ that a matrix for $S_{t,\tau}$ restricted to W is unitriangular, as x maps to $x - t$. Since there is a direct sum of vector spaces $V = W \oplus W^\perp$, we get that $\det(S_{t,\tau}) = 1$.

To finish, let U be a non-trivial isometry which restricts to the identity on an $n-1$-dimensional subspace W of V. Then, W is the kernel of a non-trivial E-linear map $\varphi\colon V \to E$. As the form $\langle -, -\rangle$ is non-degenerate on V, φ is of the form $x \mapsto \langle x, v\rangle$ for some non-zero $v \in V$. If v is anisotropic, then, by Proposition 4.3.9, V is the orthogonal direct sum $\mathrm{span}\{v\} \oplus W$. Since U is an isometry and W is the 1-eigenspace of U, U maps each orthogonal summand to itself. Therefore, we must have that U maps v to a multiple bv of v, for $b \in E^\times$, so that

$$Uv = bv = (1 - a\langle v, v\rangle)v = v - a\langle v, v\rangle v$$

for some $a \in E^\times$. From the condition that $\langle Uv, Uv\rangle = \langle v, v\rangle$ and the formula for Uv, we get that $\mathrm{Tr}_\sigma(a^{-1}) = \langle t, t\rangle$. Since, additionally, $\langle x, v\rangle = 0$ for all $x \in W$ and U is the identity on W, we obtain that the formula

$$Ux = x - a\langle x, v\rangle$$

holds for all $x \in V$. Therefore, for $t = v$ and $\tau = a^{-1}$, we have $U = S_{t,\tau}$.

Now, let v above be isotropic. For an isometry U of V, it can be shown by symbolic manipulations that the 1-eigenspace of U is orthogonal to the subspace $\mathrm{im}(I - U)$ of V. So, there exists $x \in V$ such that the non-zero vector $Ux - x$ is orthogonal to every vector of $W = v^\perp$. Since the map assigning v to the linear map $x \mapsto \langle x, v\rangle$ is injective but $Ux - x$ is non-zero, this implies $Ux - x$ is a non-zero multiple of v. By a similar argument as that of the first case, we obtain that

$$Ux = x - a\langle x, v\rangle v$$

for some $a \in E^\times$, and further, that U is a symmetry of V. ∎

We are now ready to define the unitary group of a Hermitian lattice L in V.

Definition 4.8.5. Let L be an $\mathcal{O}_E$-lattice in the Hermitian space V. The unitary group of L is

$$\mathrm{U}(L) := \mathrm{U}(V) \cap \mathrm{GL}(L).$$

Lemma 4.8.6. Let L be a Hermitian lattice in V, and let $S_{t,\tau}$ be a symmetry of V, with respect to $t \in L$. If $\langle L, t\rangle$ is contained in the fractional ideal $\tau\mathcal{O}_E$, then $S_{t,\tau}$ is an element of $\mathrm{U}(L)$.

Proof. We know that $S_{t,\tau}$ preserves the Hermitian form of L, and we know that $S_{t,\tau}$ is injective, so it suffices to prove that $S_{t,\tau}$ is surjective. Let $\ell \in L$. By looking at the formula for $S_{t,\tau}^{-1}$ and seeing as L is an $\mathcal{O}_E$-module, we have that $S_{t,\tau}^{-1}(\ell)$ is contained in L if $\frac{\langle \ell, t\rangle}{\sigma(\tau)}$ is in $\mathcal{O}_E$. But $\nu(\sigma(\tau)) = \nu(\tau)$ by Lemma 4.1.21, and therefore $\sigma(\tau)\mathcal{O}_E = \tau\mathcal{O}_E$, so that $\langle \ell, t\rangle \subseteq \tau\mathcal{O}_E$. This shows $S_{t,\tau}$ is surjective. ∎

Lemma 4.8.7. Let L be a lattice in the Hermitian space V, with $x, x' \in L$ such that $\langle x, x\rangle = \langle x', x'\rangle$ and $\langle x, x - x'\rangle \neq 0$. Setting $t = x - x'$ and $\tau = \langle x, x - x'\rangle$, we have that $S_{t,\tau}(x) = x'$. If in addition $\frac{\langle L, t\rangle}{\langle x, t\rangle} t \subseteq L$, then $S_{t,\tau}$ is an isometry of L.

Proof. We can verify by direct computation that $\mathrm{Tr}_{\mathrm{E/K}}(\tau) = \langle t, t\rangle$ and that $S_{t,\tau}(x) = x'$. The last statement is an application of Lemma 4.8.6. ∎

We now introduce a second type of isometry.

Definition 4.8.8. Let L be a Hermitian lattice and $\{u, v\}$ a hyperbolic pair splitting L with $\mathrm{span}\{x, y\} = P$. Let $y \in P^\perp$ with $\langle L, y\rangle \subseteq \langle u, v\rangle\mathcal{O}_E$, and let $\mu \in \mathcal{O}_E$ with

$$\mathrm{Tr}_{\mathrm{E/K}}(\mu\langle u, v\rangle) = -\langle y, y\rangle.$$

Then the rescaled Eichler isometry $\mathrm{E}_{\mathrm{y}}^{\mu}$ is defined by

$$\mathrm{E}_{\mathrm{y}}^{\mu}\colon L \to L, \quad x \mapsto x + \frac{\langle x, u\rangle}{\langle v, u\rangle}y + \left(\frac{\mu\langle x, u\rangle}{\langle v, u\rangle} - \frac{\langle x, y\rangle}{\langle u, v\rangle}\right)u.$$

We define a well-behaved class of isometries of L.

Definition 4.8.9. Let L be a Hermitian lattice, and let $\tau \in \mathrm{U}(L)$. We say τ is Jordan-semisimple if there exists a Jordan decomposition $[L] = \bigoplus_{i=1}^{r} L_i$ for L such that $\tau L_i \subseteq L_i$ for each $1 \leq i \leq r$. We call such a decomposition of L a Jordan-invariant decomposition.

Example 4.8.10. Any modular lattice L is Jordan-semisimple, as Theorem 4.6.8 implies that L is the only modular component appearing in a Jordan decomposition for L.

Definition 4.8.11. Let L be a Hermitian lattice, and let $\tau \in \mathrm{U}(L)$. We say τ is strictly diagonalizable if there exists an orthogonal $\mathcal{O}_E$-basis $\{x_1, \ldots, x_n\}$ for L consisting of eigenvectors of τ.

Example 4.8.12. From the definition, we obtain that a strictly diagonalizable isometry τ of L is Jordan-semisimple. Indeed, for any $i \in \mathbb{Z}$, the subset of the eigenbasis

$$\{x_1, \ldots, x_n\}$$

defined by $\{x_k \mid \nu(\langle x_k, x_k\rangle) = i\}$ generates a $\mathfrak{P}^i$-modular sublattice L_i. Let $\mathcal{I} \subseteq \mathbb{Z}$ index the modular components of L by scale. Since the eigenbasis vectors are pairwise orthogonal, we have that $\bigoplus_{i\in I} L_i$ is a Jordan decomposition for L constructed in such a way that L_i is τ-invariant for each i.

We now extend the definitions pertaining to Hermitian and sesquilinear spaces over fields found in [Wa59] and further exposed in [Gro02, §11] to our analysis of the unitary group of a Hermitian lattice over an extension of DVRs. We then make use of Wall's work to introduce an invariant for certain elements of the unitary group of a Hermitian lattice.

Definition 4.8.13. Let L be a Hermitian lattice over a quadratic extension of discrete valuation rings, and let $\tau \in \mathrm{U}(L)$. Put $\hat{\tau} = \mathrm{id} - \tau$. We define L_τ to be the image of the $\mathcal{O}_E$-linear map $\hat{\tau} \in \mathrm{Hom}_{\mathcal{O}_E}(L, L)$.

Definition 4.8.14. Suppose that L is a Hermitian lattice over a quadratic extension of discrete valuation rings, and let $\tau \in \mathrm{U}(L)$. For L_τ and $\hat{\tau}$ as above, we define the σ-sesquilinear form

$$\langle -, - \rangle_\tau \colon L_\tau \times L_\tau \to E$$

by

$$\langle u, v \rangle_\tau := \langle x, \hat{\tau}(y) \rangle$$

for any choice of $x, y \in L$ such that $u = \hat{\tau}(x)$ and $v = \hat{\tau}(y)$.

It is readily checked that $\langle u, v \rangle_\sigma$ is well-defined, independently of the choice of preimage for $u \in L_\tau$ under $\hat{\tau}$. We say a sesquilinear form on L_τ is non-degenerate if the left radical of L_τ is zero, i.e., if there is no $u \in L_\tau$ such that $\langle v, u \rangle_\tau = 0$ for all $v \in L_\tau$. The module L_τ and its sesquilinear form enjoy the following properties.

Lemma 4.8.15. Let L_τ and $\hat{\tau}$ be as above. We have

$$\langle u, v \rangle = \langle u, v \rangle_\tau + \sigma\left(\langle v, u \rangle_\tau\right)$$

for all $u, v \in L_\tau$. If L is a non-degenerate Hermitian lattice, the sesquilinear form $\langle -, - \rangle_\tau$ on L_τ is non-degenerate.

Proof. See [Gro02, §11] for details of the computation of the identity, which holds for any Hermitian lattice over an integral domain. Then, if $v \in \mathrm{rad}_L(L_\tau)$, so that $\langle u, v \rangle_\tau = 0$ for all $u \in L_\tau$, we have for any choice of $x \in L$ that $0 = \langle x - \tau x, v \rangle_\tau = \langle x, v \rangle$, so that $v \in \mathrm{rad}(L)$. Since L is non-degenerate, $v = 0$. ∎

Since $\mathcal{O}_E$ has residue field $e = \mathcal{O}_E / \mathfrak{P}\mathcal{O}_E$, we can consider L_τ as an e-vector space in the natural way. If we further assume that $\langle -, - \rangle$ is $\mathfrak{P}^i$-modular, we can then read

the sesquilinear form $\langle -, - \rangle_\tau$ as taking values in the field $\mathfrak{P}^i\mathcal{O}_E/\mathfrak{P}^{i+1}\mathcal{O}_E \cong e$ by rescaling, as Lemma 4.1.21 tells us that the involution σ preserves each fractional ideal $\mathfrak{P}^i$ of $\mathcal{O}_E$. We introduce the following definition, which will allow us to describe the length of a factorization of certain isometries of L in terms of symmetries.

Definition 4.8.16. Let L be a Hermitian lattice, and let $\tau \in \mathrm{U}(L)$ be Jordan-semisimple. Let $[L] = \bigoplus_{i \in \mathcal{I}} L_i$ be a Jordan decomposition for L such that $\tau L_i \subseteq L_i$. Further assume that each L_i is $\mathfrak{P}^i$-modular. Identifying the quotient $\mathfrak{P}^i\mathcal{O}_E/\mathfrak{P}^{i+1}\mathcal{O}_E$ with the residue field e, we define the anisotropic dimension of τ with respect to L_i, denoted by

$$\dim_a(\tau, L_i),$$

to be the rank of any representing matrix for the e-sesquilinear form $\langle -, - \rangle_\tau$ on the e-vector space

$$L_{\tau|_{L_i}} = \mathrm{im}\left(I - \tau|_{L_i}\right).$$

If $[L] = \bigoplus_{i=1}^k L_i$ is a Jordan-invariant decomposition for L, we define the anisotropic dimension of τ with respect to $[L]$ to be the sum

$$\dim_a(\tau, [L]) = \sum_{i=1}^{k} \dim_a(\tau, L_i).$$

In other words, $\dim_a(\tau, [L])$ is the sum of the anisotropic dimensions of τ with respect to each modular component of the Jordan-invariant decomposition $[L]$.

Remark 4.8.17. The sublattice L_τ consists of all elements of L of the form $x - \tau x$, $x \in L$. The $\mathcal{O}_E$-sesquilinear form on $L_\tau \times L_\tau$ is then given by the formula

$$\langle x - \tau x, y - \tau y \rangle_\tau = \langle x, y - \tau y \rangle.$$

If G is a Gram matrix for $\langle -, - \rangle_\tau$ and H is a Gram matrix for $\langle -, - \rangle$, this shows that the column space of G is the same as the column space of $H - HT$, where T is a matrix representing τ. Therefore, if L is modular, we can compute the anisotropic dimension of τ by computing the rank of $H - HT$ over the residue field.

Definition 4.8.18. Let L be a Hermitian lattice and let $\tau \in \mathrm{U}(L)$ be Jordan-semisimple. We define the anisotropic dimension of τ to be

$$\dim_a(\tau) = \max \dim_a(\tau, [L]),$$

where the maximum is taken over all Jordan-invariant decompositions $[L]$ for L.

4.9 The Cartan-Dieudonné Theorem

We now extend the techniques of [BHM22] used to prove the Cartan-Dieudonné Theorem for Hermitian lattices over the valuation ring of an unramified extension of non-Archimedean local fields to all Hermitian lattices over separable quadratic extensions of DVRs. We obtain a new upper bound on the minimum number of symmetries required to factorize a strictly diagonalizable isometry. We then present a conjecture on the minimal number of symmetries required to factorize a Jordan-semisimple isometry of a Hermitian lattice. To this end, we make extensive use of the notation and definitions of [BHM22], and of Wall's terminology just previously introduced.

We set $\mathcal{O}_E/\mathcal{O}_K$ to be a separable quadratic extension of discrete valuation rings with non-trivial standard involution σ, together with the corresponding quadratic extension of residue fields e/k, equipped with the standard involution $\overline{\sigma}$. We set E to be the field of fractions of $\mathcal{O}_E$, and K that of $\mathcal{O}_K$.

Theorem 4.9.1. *Let $\mathcal{O}_E/\mathcal{O}_K$ be a quadratic separable extension of discrete valuation rings with non-trivial standard involution σ. Let L be a Hermitian lattice over $\mathcal{O}_E$ of rank n. The unitary group of L is generated by symmetries. Any minimal factorization of a strictly diagonalizable isometry U of L has length at most $2\,\mathrm{rk}(L_U) - \dim_a(U, [L])$, where $[L]$ is the Jordan decomposition afforded by the strict diagonalizability of U.*

We make a quick remark on the hypotheses of the Theorem. It is not sufficient to assume that the field extension E/K is separable to get that the corresponding extension of valuation rings is separable. Indeed, if E/K is any totally ramified quadratic extension of a 2-adic field, the restriction of the trace map to the valuation ring is multiplication by 2 when considered modulo the maximal ideal of $\mathcal{O}_E$ and is therefore not surjective. By Lemma 2.1.23, $\mathcal{O}_E/\mathcal{O}_K$ cannot be separable. Since the technique for obtaining a factorization of an isometry makes use of the surjectivity of the trace in an essential way, we must exclude such cases. However, using techniques adapted to the case where the maximal ideal of $\mathcal{O}_K$ ramifies, it is shown in [BHM22] that a version of the theorem holds for almost every Hermitian lattice over inseparable quadratic extensions of DVR's which arise from extensions of 2-adic local fields when the residue field is not $\mathbb{F}_2$. This counterexample had already been computed in [Hay68], using the fact that a Hermitian lattice over such a ramified quadratic extension collapses under the residue map to a *quadratic* space over $\mathbb{F}_2$ because of the

trivialization of the involution. In sum, our work builds on [BHM22] by extending their algorithm to all separable quadratic extensions of discrete valuation rings, and by describing the length of the resulting factorizations using an entirely new invariant based on the definitions of [Wa59].

To prove the Theorem, we need a few lemmas. We set $\mathcal{O}_E/\mathcal{O}_K$ to be a separable quadratic extension of discrete valuation rings with standard involution σ, together with the corresponding quadratic extension of fraction fields E/K. Since $\mathcal{O}_E$ is separable over $\mathcal{O}_K$, the standard involution on $\mathcal{O}_E$ extends to a non-trivial K-automorphism of E, so E/K is separable. Write e/k for the quadratic extension of residue fields, equipped with the standard involution $\overline{\sigma}$. We choose $\pi \in K$ to be a uniformizer for both $\mathcal{O}_K$ and $\mathcal{O}_E$. This is allowed by Lemma 2.5.16. Since $\mathcal{O}_E/\mathcal{O}_K$ is separable, we have from Lemma 4.1.20 that the trace map is surjective. Therefore, there exists an element $\rho \in \mathcal{O}_E$ with $\mathrm{Tr}_\sigma(\rho) = 1$. Since the valuation is non-Archimedean, ρ is a unit of $\mathcal{O}_E$.

The following result is an extension of Lemma 4.1 from [BHM22].

Lemma 4.9.2. Let $u, u' \in L$ be such that

$$\langle u, L\rangle = \langle u', L\rangle = \mathfrak{s}(L).$$

Then there exists $x \in L$ such that

$$\langle x, u\rangle\mathcal{O}_E = \langle x, u'\rangle\mathcal{O}_E = \mathfrak{s}(L).$$

If u and u' are isotropic and $\langle u, u'\rangle\mathcal{O}_E \neq \mathfrak{s}(L)$, we can choose x above to be isotropic.

Proof. By assumption we have that there exist elements $y, y' \in \mathcal{O}_E$ such that

$$\langle u, y\rangle\mathcal{O}_E = \langle u', y'\rangle\mathcal{O}_E = \mathfrak{s}(L).$$

Suppose that no $z \in \{y, y'\}$ satisfies $\langle z, u\rangle\mathcal{O}_E = \langle z, u'\rangle\mathcal{O}_E = \mathfrak{s}(L)$. We then have that

$$\langle y, u'\rangle\mathcal{O}_E \subsetneq \langle y, u\rangle\mathcal{O}_E = \mathfrak{s}(L)$$

and

$$\langle y', u\rangle\mathcal{O}_E \subsetneq \langle y', u'\rangle\mathcal{O}_E = \mathfrak{s}(L).$$

Consider the element $x := y + y'$.

By the strict inclusions above, we must have that both $\nu(\langle y, u' \rangle)$ and $\nu(\langle y', u \rangle)$ are strictly greater than $\nu(\langle y, u \rangle) = \nu(\langle y', u' \rangle)$, so that by the strengthening of the ultrametric inequality of Corollary 3.4.6,

$$\nu(\langle x, u \rangle) = \min\{\nu(\langle y, u \rangle), \nu(\langle y', u \rangle)\} = \nu(\langle y, u \rangle).$$

Similarly, we have $\nu(\langle x, u' \rangle) = \nu(\langle y', u' \rangle)$. This implies that

$$\langle x, u \rangle \mathcal{O}_E = \langle x, u' \rangle \mathcal{O}_E = \mathfrak{s}(L).$$

Let us show that we can choose the above $x \in L$ to be isotropic if u and u' are isotropic. For any $\rho \in \mathcal{O}_E$ with $\mathrm{Tr}_\sigma(\rho) = 1$, we consider the element $x' = x - \rho \frac{\langle x,x \rangle}{\langle u,x \rangle} u$. For some scalar $\lambda \in \mathcal{O}_E$, we have

$$\langle x', x' \rangle = \langle x, x \rangle (1 + \mathrm{Tr}_\sigma(-\rho) + \lambda \langle u, u \rangle) = 0.$$

Since u is isotropic, we have

$$\langle x', u \rangle = \langle x, u \rangle - \rho \frac{\langle x, x \rangle}{\langle u, x \rangle} \langle u, u \rangle = \langle x, u \rangle,$$

so that $\langle x', u \rangle \mathcal{O}_E = \langle x, u \rangle \mathcal{O}_E = \mathfrak{s}(L)$.

We then have $\langle x', u' \rangle \mathcal{O}_E = \langle x, u' \rangle \mathcal{O}_E = \mathfrak{s}(L)$, by using Corollary 3.4.6 with

$$\langle x', u' \rangle = \langle x, u' \rangle - \rho \frac{\langle x, x \rangle}{\langle u, x \rangle} \langle u, u' \rangle$$

and the assumption that $\nu(\langle u', u \rangle)$ is strictly greater than $\min\{\nu(\mathfrak{s}(L))\}$. ∎

We modify [BHM22, Lem 4.2] for our use.

Lemma 4.9.3. Suppose that L is unimodular, and let $x, x' \in L$ with $\langle x, x \rangle = \langle x', x' \rangle$. If $\langle x, x \rangle$ is a unit of $\mathcal{O}_E$, there exists a product of symmetries of L mapping x' to x.

Proof. If $\langle x, x - x' \rangle$ is a unit of $\mathcal{O}_E$, then so is $\langle x', x' - x \rangle = \sigma(\langle x, x - x' \rangle)$. We can therefore apply Lemma 4.8.7 with $t = x' - x$ and $\tau = \langle x', x' - x \rangle$ to get $S_{t,\tau}(x') = x$.

Suppose now that $\langle x, x - x'\rangle$ is not a unit of $\mathcal{O}_E$, so that $\nu(\langle x, x - x'\rangle) > 0$. Taking an element $\rho \in \mathcal{O}_E$ with $\mathrm{Tr}_\sigma(\rho) = 1$, necessarily a unit, the symmetry $S_{t,\tau}$ with $t = x'$ and $\tau = \rho\langle x', x'\rangle \in \mathcal{O}_E^\times$ is in $\mathrm{U}(L)$ by Lemma 4.8.6. Now, we have

$$\langle x, x - S_{t,\tau}(x')\rangle = \langle x, x - x'\rangle + \frac{\langle x', x'\rangle}{\rho\langle x, x'\rangle}\langle x, x'\rangle = \langle x, x - x'\rangle + \langle x, x'\rangle\sigma(\rho)^{-1}. \quad (4.9.1)$$

By assumption, $\langle x, x\rangle$ is a unit of $\mathcal{O}_E$. In this case, we must have that $\langle x, x'\rangle$ is a unit of $\mathcal{O}_E$, as

$$\nu(\langle x, x\rangle - \langle x, x'\rangle) = \nu(\langle x, x - x'\rangle) > 0$$

which can only happen if $\nu(\langle x, x'\rangle) = 0$, by Corollary 3.4.6. Therefore, $\langle x, x'\rangle\sigma(\rho)^{-1}$ must be a unit of $\mathcal{O}_E$, so, we apply the Lemma to Equation 4.9.1 to obtain that

$$\omega := \langle x, x - S_{t,\tau}(x')\rangle$$

is a unit of $\mathcal{O}_E$. We set $w = x - S_{t,\tau}(x')$ and, by a quick computation, we have $S_{w,\omega}(x) = S_{t,\tau}(x')$. By Lemma 4.8.6, we must have $S_{w,\omega} \in \mathrm{U}(L)$. This shows

$$S_{w,\omega}^{-1} \circ S_{t,\tau}(x') = x,$$

where both symmetries are symmetries of L. ∎

Proof of theorem.

Let us assume that L satisfies $\langle L, L\rangle = \mathcal{O}_E$, since replacing the form $\langle -, -\rangle$ by $\pi^i\langle -, -\rangle$ for an appropriate $i \in \mathbb{Z}$ does not change the unitary group. We proceed by induction on $n = \mathrm{rk}(L)$. Let $U \in \mathrm{U}(L)$ and write L in its Jordan decomposition $[L] = \bigoplus_{i=1}^k L_i$ with L_1 unimodular. Since $\mathcal{O}_E/\mathcal{O}_K$ is separable, we have that L_1 decomposes as an orthogonal sum $a_1x_1 \oplus \cdots \oplus a_mx_m$ for some $a_j \in \mathcal{O}_K^\times$, $j \in \{1, \ldots, m\}$, by Proposition 4.7.1. Up to an appropriate reordering of the orthogonal basis, we can assume that $Ux_1 \neq x_1$.

Let us consider $y_1 := Ux_1$. By assumption, we have that $\langle x_1, x_1\rangle = a_1$ is a unit of $\mathcal{O}_E$. We apply Lemma 4.9.3 to x_1 and y_1 to obtain a product S of at most two symmetries of L mapping y_1 to x_1.

Set $N = \mathrm{span}\{x_1\}^\perp$ and let $w \in N$. We have

$$0 = \langle x_1, w\rangle = \langle SUx_1, SUw\rangle = \langle x_1, SUw\rangle,$$

so that SUw is in N. Hence, the isometry SU of L restricts to an isometry of N, a Hermitian lattice of rank strictly less than n. In particular, $N_{SU} \subseteq N$. Applying induction, SU is a product of symmetries of N. Since we have the orthogonal direct sum $L = \text{span}\{x_1\} \oplus N$, the factorization of SU into symmetries of N extends trivially to all of L. This establishes the existence of a factorization of U as a product of symmetries.

It remains to describe the length of the factorization of a strictly diagonalizable isometry U of L. Let $\{x_1, \ldots, x_n\}$ be an ordered orthogonal basis of L consisting of eigenvectors of U. By rescaling the Hermitian form, we can assume that the scale of L is $\mathcal{O}_E$, and up to reordering, we can assume that both $\langle x_1, x_1 \rangle$ is a unit and that no x_k is a 1-eigenvector. This is because the orthogonality of the eigenspaces allows us to describe U as an isometry of the lattice generated by its non-trivial eigenvectors. As above, we write $N = \text{span}\{x_2, \ldots, x_n\}$.

We show that U is an isometry of N and that $U|_N = SU|_N$. Since U is an isometry of L and x_1 is an eigenvector, we have that $Ux_1 = ax_1$ for some $a \in E^\times$ of norm 1. Since x_1 is maximal and $ax_1 \in L$, we must have that $a \in \mathcal{O}_E$. In particular, a is a unit of $\mathcal{O}_E$. As above,

$$\begin{aligned} \langle x_1, Uy \rangle &= \langle \frac{1}{a} Ux_1, Uy \rangle \\ &= \frac{1}{a} \langle x_1, y \rangle \\ &= 0 \end{aligned}$$

for any $y \in N$. So, UN is contained in the orthogonal complement N of x_1. Therefore, $U = a \oplus U|_N$. Since a is a unit of $\mathcal{O}_E$ and $\det(U) = a \det(U|_N)$, we must have that $\det(U|_N)$ is also a unit of $\mathcal{O}_E$. So, $U|_N$ is an isometry of N. In the case that $\langle x_1, x_1 - Ux_1 \rangle$ is a unit of $\mathcal{O}_E$, S is defined with respect to the vector $t = (a-1)x_1$. By construction, S is trivial on $t^\perp = x_1^\perp = N$. As $UN = N$, $U|_N = SU|_N$. If $\langle x_1, x_1 - Ux_1 \rangle$ is not a unit of $\mathcal{O}_E$, then $S = S_{w,\omega}^{-1} \circ S_{t,\tau}$, where $t = Ux_1 = ax_1$, and $w = x_1 - S_{t,\tau} Ux_1 = bx_1$, some $b \in \mathcal{O}_E$. So, S is a product of symmetries defined with respect to multiples of x_1, and, as a result, S is trivial on $x_1^\perp = N$. As in the first case, we have that $U|_N = SU|_N$.

We claim that in both cases,

$$\text{rk}(L_U) = \text{rk}(N_{SU}) + 1.$$

By the decomposition of L as $x_1 \oplus N$ and the equality $U|_N = SU|_N$, it suffices to show that $x_1 - Ux_1$ is not contained in N_{SU}. In this case, $x_1 - Ux_1$ extends any basis of N_{SU} to a basis $\mathcal{S}$ of L_U. Suppose, towards contradiction, that $\langle x_1, x_1 - Ux_1\rangle = 0$. Let $1 \leq i, j \leq n$ be such that $i \neq j$. We have that

$$\begin{aligned}\langle x_i - Ux_i, x_j - Ux_j\rangle_U &= \langle x_i, x_j - Ux_j\rangle \\ &= \langle x_i, x_j\rangle - \langle x_i, a_j x_j\rangle \\ &= 0.\end{aligned}$$

Then, $\langle x_1 - Ux_1, x_1 - Ux_1\rangle_U = 0$, and seeing as

$$\langle N_{SU}, x_1 - Ux_1\rangle_{SU} = \langle N_U, x_1 - Ux_1\rangle_U = 0,$$

we have that $x_1 - Ux_1$ is in the left radical of L_U. We have that L_U is non-degenerate by Lemma 4.8.15. So, $x_1 - Ux_1 = 0$. This means that x_1 is a 1-eigenvector for U, a contradiction with our reduction. Hence, $\langle x_1, x_1 - Ux_1\rangle \neq 0$, and we have that $x_1 - Ux_1$ is not contained in $N = x_1^{\perp}$. Seeing as $N_{SU} = N_U \subseteq N$, $x_1 - Ux_1$ is not contained in N_{SU}. Let H be the Gram matrix of the e-sesquilinear form $\langle -, -\rangle_U$ on the quotient $L_U/\mathfrak{P}L_U$, with respect to the image of the basis $\mathcal{S}$ for L_U just previously found. Note that, by the equality $N_U = N_{SU}$, the matrix G obtained by deleting the first row and column of H is a Gram matrix for N_{SU} under the projection to the residue field. To finish, we will compare the ranks of H and G and count the number of symmetries obtained in a Cartan-Dieudonné decomposition of U with respect to the given orthogonal basis. Inductively, $SU|_N = U|_N$ is a strictly diagonalizable isometry of N, and is therefore a product of $2\operatorname{rk}(N_{SU}) - \dim_a(SU, [N])$ symmetries of N, where the Jordan decomposition of N is that induced by the orthogonal basis.

Suppose that S was obtained as in the first case of Lemma 4.9.3. Then, $\langle x_1, x_1 - Ux_1\rangle$ is a unit of $\mathcal{O}_E$ and the first entry of H is a unit of e. By the orthogonality relation proved above, every other entry of the first row is 0. This shows that the e-rank of H is one more than the e-rank of G. So, $\dim_a(U, [L]) = \dim_a(SU|_N, [L]) + 1$, and thus $\dim_a(U, L_1) = \dim_a(SU, L_1) + 1 = \dim_a(SU, N_1) + 1$. We have therefore shown

$$\dim_a(U, [L]) = \dim_a(SU, [N]) + 1.$$

Since S is a symmetry and SU is a product of

$$2\operatorname{rk}(N_{SU}) - \dim_a(SU, [N]) = 2(\operatorname{rk}(L_U) - 1) - (\dim_a(U, [L]) - 1)$$

$$= 2\left(\mathrm{rk}\left(L_U\right)\right) - \left(\dim_a\left(U, [L]\right) - 1\right.$$

symmetries, U is a product of

$$2\left(\mathrm{rk}\left(L_U\right)\right) - \dim_a\left(U, [L]\right)$$

symmetries.

If S was obtained as in the second case of Lemma 4.9.3, then S is a product of two symmetries and $\langle x_1, x_1 - Ux_1\rangle$ is 0 over the residue field. Applying the first and second facts proved above, we have that the first row and column of H are identically zero. Comparing the ranks of G and H, this shows that $\dim_a\left(U, [L]\right) = \dim_a\left(SU, [N]\right)$. So, SU is a product of

$$2\left(\mathrm{rk}\left(L_U\right) - 1\right) - \dim_a\left(U, [L]\right)$$

symmetries. As S is a product of two symmetries, U must be a product of

$$2\,\mathrm{rk}\left(L_U\right) - \dim_a\left(U, [L]\right)$$

symmetries of L. ∎

We will now turn to computing some examples.

Example 4.9.4. We can extract an algorithm for computing a decomposition of an isometry $U \in \mathrm{U}(L)$ into a product of symmetries from the proof of Theorem 4.9.1. In the proof of the Theorem, we proceeded by induction, where each step consisted of splitting off a suitable basis vector x_1 of L, then giving a product of symmetries S which mapped Ux_1 back to x_1. To apply induction, we invoked the orthogonal decomposition of L into $\mathrm{span}_{\mathcal{O}_E}\{x_1\}$ and its orthogonal complement, which we denoted by N. We then observed that SU was an isometry of N, then claimed that we could repeat the process for SU and the lattice N. As mentioned in the proof of the Theorem, multiplying the Hermitian form by a power of a uniformizer $\pi \in \mathcal{O}_K$ for $\mathcal{O}_K$ does not change the embedding of $\mathrm{U}(L)$ in $M_n(\mathcal{O}_E)$. So, we rescale the lattice if necessary and assume that the scale of L is $\mathcal{O}_E$.

We first want to find a vector x_1 in L such that $\langle x_1, x_1\rangle$ is a unit, and such that x_1 can be extended to an orthogonal basis of L. Since the scale of L is $\mathcal{O}_E$ and $\mathcal{O}_E$ is separable over $\mathcal{O}_K$, finding such an orthogonal basis is possible because of Proposition 4.7.1 and Theorem 4.6.8. To find this effective basis of induction, we will

use the main tool we have introduced in this thesis for computing orthogonal bases: the Gram-Schmidt package of Lemma 4.5.1 and Theorem 4.5.2. As L is a finitely generated free module over the principal ideal domain $\mathcal{O}_E$, we can assume that we have access to some basis for L, for instance, by computing the Smith normal form of L. By the proof of Theorem 4.5.2, repeatedly computing the Fourier coefficients of Lemma 4.5.1 for this basis effectively decomposes L into an orthogonal sum of lines and planes. Since $\mathcal{O}_E$ is separable, we can then repeatedly apply the proof of Corollary 4.5.5 to any remaining rank 2 summand of L, thus diagonalizing any planes encountered. In other words, we apply the proof of Proposition 4.7.1 to the given basis of L, and obtain an orthogonal basis $\{x_1, \ldots, x_n\}$ for L. In particular, each line spanned by x_i is modular, and grouping these lines by scale gives a Jordan decomposition for L. Since $\{x_1, \ldots, x_n\}$ is an orthogonal basis for L and the scale of L is $\mathcal{O}_E$, we can apply the uniqueness portion of Theorem 4.6.8 to say that one of the x_i must satisfy $\langle x_i, x_i\rangle = \mathcal{O}_E$, as otherwise, the image of $\langle -, -\rangle$ would be properly contained in $\mathcal{O}_E$. Without loss of generality, this vector is x_1, and we have our desired basis of induction. To define the symmetries of the Theorem, we need an element $\rho \in \mathcal{O}_E$ of trace 1. Since $\mathcal{O}_E$ is a free $\mathcal{O}_K$-module of rank 2 and the involution trace is an $\mathcal{O}_K$-linear map, we can effectively find such $\rho \in \mathcal{O}_E$ with the Smith Normal Form.

We can now examine the action of U on x_1, and there are three cases. Set $N_0 = \{x_1, \ldots, x_n\}$ to be the orthogonal basis of L we computed, and more generally, set $N_i := \{x_j \mid j > i\} \subseteq N_0$.

1. If U fixes x_1, then U is entirely determined by its action on $N_1 := \{x_2, \ldots, x_n\}$. In this case, put $S_1 = \mathrm{id}$.

2. If $x_1 \neq Ux_1$, but $\langle x_1, x_1 - Ux_1\rangle$ is a unit, we can set $t = Ux_1 - x_1$ and $\tau = \langle Ux_1, Ux_1 - x_1\rangle$ to get $S_{t,\tau}Ux_1 = x_1$. Put $S_1 = S_{t,\tau}$.

3. If $x_1 \neq Ux_1$ and $\langle x_1, x_1 - Ux_1\rangle$ is not a unit, we set $t := Ux_1$, $\tau = \rho\langle x_1, x_1\rangle$, $w = x_1 - S_{t,\tau}Ux_1$, and $\omega = \langle x_1, x_1 - S_{t,\tau}Ux_1\rangle$. Then,

$$S_{w,\omega}^{-1}S_{t,\tau}(Ux_1) = x_1.$$

   Put $S_1 = S_{w,\omega}^{-1} \circ S_{t,\tau}$.

So, we obtain a (possibly empty) product of symmetries S_1 and an isometry S_1U

of $\mathrm{span}_{\mathcal{O}_E}\{x_2, \ldots, x_n\} = \mathrm{span}_{\mathcal{O}_E} N_1$. Multiplying the Hermitian form on N_1 with an appropriate power of π, the isometry S_1U satisfies one of the three cases above.

So, we loop over these three cases. At step i, we input the isometry $S_{i-1} \ldots S_1U$ obtained at the last iteration, which fixes $\mathrm{span}\{N_0 \backslash N_{i-1}\}$. By selecting the appropriate case (with respect to the vector x_i), we obtain a product S_i of symmetries of L such that $S_iS_{i-1} \ldots S_1U$ fixes $N_0 \backslash N_i$. Therefore, $S_iS_{i-1} \ldots S_1U$ is an isometry of both L and $\mathrm{span}_{\mathcal{O}_E} N_i$. By repeating this process, we eventually obtain an equality $S_n \ldots S_1U = \mathrm{id}$. Seeing as Proposition 4.8.4 says that the inverse of an arbitrary symmetry $S_{t,\tau}$ is the symmetry $S_{t,\sigma(\tau)}$, we know how to write $U = S_1^{-1} \ldots S_n^{-1}$ as a product of symmetries.

We now give two examples of a Cartan-Dieudonné decomposition.

Example 4.9.5. Consider the unique unramified quadratic extension E of $\mathbb{Q}_2$, given by adjoining a primitive third root of unity ζ_3 to $\mathbb{Q}_2$, i.e., $E = \mathbb{Q}_2(\zeta_3)$. Let 2 be our chosen uniformizer for E/K and define the lattice L to be the direct sum of hyperbolic planes

$$L = H(0) \oplus H(1)$$

with Gram matrix

$$H_L = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 2 \\ 0 & 0 & 2 & 0 \end{bmatrix}.$$

Identifying an appropriate ordered basis $\{u, v, z, w\}$ of L with the standard basis of $\mathcal{O}_E^4$, we have that $u := e_1$ and $v := e_2$ are isotropic vectors, with $\langle u, v \rangle = 1$. We see that

$$L \cong \mathrm{span}\{u, v\} \oplus \mathrm{span}\{u, v\}^{\perp}$$

so that (u, v) is a hyperbolic pair splitting L. Set $z := e_3$, $w := e_4$ and $y = z + w$, so that $\langle y, y \rangle = 4$. Then, we have that $\mathrm{Tr_{E/K}}(-2) = -4 = -\langle y, y \rangle$. Therefore, the Eichler isometry E_y^{-2} with respect to the hyperbolic pair (u, v) exists. We then have

$$\mathrm{E_y^{-2}}(u) = u, \quad \mathrm{E_y^{-2}}(v) = -2u + v + z + w, \quad \mathrm{E_y^{-2}}(z) = -2u + z, \quad \mathrm{E_y^{-2}}(w) = -2u + w,$$

so that the corresponding matrix for E_y^{-2} is given by

$$E := \begin{bmatrix} 1 & -2 & -2 & -2 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{bmatrix}.$$

We will compute the decomposition of E_y^{-2} into a product of symmetries as in the proof of Theorem 4.9.1. Indeed, the proof gives a method to compute this decomposition, and at each step, there are two key ingredients. First, we need to find an appropriate vector $x \in L$ and an appropriate product S of symmetries that maps $\mathrm{E}_y^{-2}(x)$ to x. The second ingredient is an orthogonal basis for L extending $\{x\}$. In particular, this shows that we must take x to be anisotropic, as an isotropic vector cannot be extended to an orthogonal basis of the entire (non-degenerate) lattice. Let $x \in L$ have $\langle x, x\rangle \in \mathcal{O}_E^\times$. We can always find a product of symmetries mapping $\mathrm{E}_y^{-2}(x)$ to x, by Lemma 4.9.3. We would like to see if it is possible to find a minimal product of symmetries mapping $\mathrm{E}_y^{-2}(x)$ to x, i.e., a single symmetry taking $\mathrm{E}_y^{-2}(x)$ directly to x. We are therefore in the case of Lemma 4.9.3. We would then like to take x such that $\langle x, x\rangle$ and $\langle \mathrm{E}_y^{-2}(x), \mathrm{E}_y^{-2}(x) - x\rangle$ are both invertible, since Lemma 4.9.3 will give us a single symmetry that maps $\mathrm{E}_y^{-2}(x)$ to x. Let us see if this is possible.

By expanding the Hermitian form, the condition

$$\nu(\langle \mathrm{E}_y^{-2}(x), \mathrm{E}_y^{-2}(x) - x\rangle) = 0$$

is equivalent to

$$\nu(\langle x, x\rangle - \langle x, \mathrm{E}_y^{-2}(x)\rangle) = 0. \tag{4.9.2}$$

Since $\langle x, x\rangle$ is assumed to be invertible, Equation 4.9.2 fails to hold precisely when

$$\langle x, x\rangle \equiv \langle x, \mathrm{E}_y^{-2}(x)\rangle \pmod{\mathfrak{P}}.$$

Since reduction mod $\mathfrak{P}$ is a homomorphism of commutative rings, multiplication of matrices commutes with reduction mod $\mathfrak{P}$, and for the same reason, so does the Galois involution σ. We can therefore reduce H_L and E_y^{-2} modulo $\mathfrak{P}$ and investigate the solutions to the congruence 4.9.2. We have that the form

$$y \mapsto \langle y, y\rangle \pmod{\mathfrak{P}}$$

is represented by the matrix $\overline{H_L}$ and

$$y \mapsto \langle y, \mathrm{E}_{\mathrm{y}}^{\text{-}2}(y)\rangle \pmod{\mathfrak{P}}$$

can be represented by the matrix $\overline{H_L E}$. We see that

$$\overline{H_L} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} = \overline{H_L E},$$

so that every $y \in L$ in the unimodular component of the lattice is a solution to the congruence 4.9.2. This means that we will have to settle for a product of two symmetries taking $\mathrm{E}_{\mathrm{y}}^{\text{-}2}(x)$ to x for the first step of the induction. In particular, there is no optimal choice of vector x for the first step.

The second main ingredient required in the proof of the theorem is a basis of induction: an orthogonal decomposition of the lattice L. For example, an orthogonal basis of L will work. As we calculated in Example 4.7.5, an orthogonal basis for L is given by

$$y_1 := \zeta_3 u + v, \quad y_2 := u - \zeta_3 v, \quad y_3 := \zeta_3 z + w, \quad y_4 := z - \zeta_3 w.$$

As the first step, we compute a product of symmetries $S_1 \in \mathrm{U}(L)$ taking $\mathrm{E}_{\mathrm{y}}^{\text{-}2}(y_1)$ to y_1. To do this, we apply Lemma 4.9.3 to y_1, and seeing that

$$\langle \mathrm{E}_{\mathrm{y}}^{\text{-}2}(y_1), \mathrm{E}_{\mathrm{y}}^{\text{-}2}(y_1) - y_1\rangle = 2$$

is not a unit of $\mathcal{O}_E$ (as predicted above), there are symmetries $S_{w,\omega}$ and $S_{t,\tau}$ such that

$$y_1 = S_{w,\sigma(\omega)} \circ S_{t,\tau} \circ \mathrm{E}_{\mathrm{y}}^{\text{-}2}(y_1). \tag{4.9.3}$$

If we set $\rho := -\zeta_3$ so that $\mathrm{Tr}_{\mathrm{E/K}}(\rho) = -\langle \mathrm{E}_{\mathrm{y}}^{\text{-}2}(y_1), \mathrm{E}_{\mathrm{y}}^{\text{-}2}(y_1)\rangle$, the symmetries satisfying 4.9.3 are determined by the values

$$t = \mathrm{E}_{\mathrm{y}}^{\text{-}2}(y_1), \quad \tau = \rho\langle t, t\rangle$$

and

$$w = y_1 - S_{t,\tau} \circ \mathrm{E}_{\mathrm{y}}^{\text{-}2}(y_1), \quad \omega = \langle y_1, y_1 - w\rangle.$$

Then, for $S_1 := S_{w,\sigma(\omega)} \circ S_{t,\tau}$ we have the equality

$$y_1 = S_1 \circ \mathrm{E}_{\mathrm{y}}^{\text{-}2}(y_1).$$

We have reached the inductive step. If S_2 is any product of symmetries with respect to lines orthogonal to y_1, we will also have that $S_2 \circ S_1 \circ \mathrm{E}_{\mathrm{y}}^{\text{-}2}(y_1) = y_1$, as symmetries with respect to a line fix the hyperplane orthogonal to that line. Since y_1 is fixed by $S_1 \circ \mathrm{E}_{\mathrm{y}}^{\text{-}2}$, any symmetry obtained as in Lemma 4.9.3 will be with respect to a vector orthogonal to y_1. We apply the Lemma once more and get a product of symmetries S_2 with respect to lines orthogonal to y_1 fixing both $S_1 \circ \mathrm{E}_{\mathrm{y}}^{\text{-}2}(y_1)$ and $S_1 \circ \mathrm{E}_{\mathrm{y}}^{\text{-}2}(y_2)$. In fact, the product $S_2 \circ S_1 \circ \mathrm{E}_{\mathrm{y}}^{\text{-}2}$ is equal to the identity of L. So,

$$\mathrm{E}_{\mathrm{y}}^{\text{-}2} = S_1^{-1} \circ S_2^{-1},$$

as desired.

Example 4.9.6. Let $\mathcal{O}_E = (\mathbb{C}[X,Y]/\langle X^2+Y^2-1\rangle)_{\mathfrak{P}}$ be the separable quadratic $\mathcal{O}_K = (\mathbb{R}[X,Y]/\langle X^2+Y^2-1\rangle)_{\mathfrak{p}}$-algebra of Example 3.6.21, and let $\mathfrak{P}$ and $\mathfrak{p}$ be the respective maximal ideals. Define the $\mathcal{O}_E$-lattice $L = \mathcal{O}_E^3$ by the Gram matrix

$$H = \begin{bmatrix} 1 & 0 & 0 \\ 0 & Y & 0 \\ 0 & 0 & X+Y \end{bmatrix}$$

with standard basis $\{e_1, e_2, e_3\}$. We observe that the lattice is unimodular, as its Gram matrix is diagonal with unit determinant. We define elements $v_1, v_2, v_3 \in L$ by

$$v_1 = e_1 + e_2, \quad v_2 = e_1 + e_3, \quad v_3 = e_2 + e_3.$$

Then,

$$\langle v_1, v_1\rangle = 1+Y, \quad \langle v_2, v_2\rangle = 1+X+Y, \quad \langle v_3, v_3\rangle = X+2Y,$$

so choosing units

$$\alpha_1 = \frac{1}{2}(1+Y), \quad \alpha_2 = \frac{1}{2}(1+X+Y), \quad \alpha_3 = \frac{1}{2}(X+2Y)$$

of $\mathcal{O}_E$ allows us to define symmetries $\tau_i = S_{v_i,\alpha_i}$ by the formula of Definition 4.8.2. By Lemma 4.8.6, the τ_i are symmetries of L.

Consider the isometry

$$\tau = \tau_1 \circ \tau_2 \circ \tau_3.$$

A matrix for τ with respect to the standard basis is given by

$$T = \begin{bmatrix} \frac{XY+Y^2-X-2Y+1}{XY+Y^2+X+2Y+1} & \frac{-2X^2Y+2XY^2+4Y^3-6XY-4Y^2}{X^2Y+3XY^2+2Y^3+X^2+4XY+4Y^2+X+2Y} & \frac{6X^2Y+10XY^2+4Y^3-2X^2+2XY+4Y^2}{X^2Y+3XY^2+2Y^3+X^2+4XY+4Y^2+X+2Y} \\ \frac{-2X-2Y+2}{XY+Y^2+X+2Y+1} & \frac{-X^2Y-XY^2+X^2-8XY-8Y^2+X}{X^2Y+3XY^2+2Y^3+X^2+4XY+4Y^2+X+2Y} & \frac{2X^2Y+4XY^2+2Y^3-6X^2-6XY-2X-2Y}{X^2Y+3XY^2+2Y^3+X^2+4XY+4Y^2+X+2Y} \\ \frac{-2}{X+Y+1} & \frac{2XY+2Y^2-2Y}{X^2+3XY+2Y^2+X+2Y} & \frac{X^2+XY-X}{X^2+3XY+2Y^2+X+2Y}. \end{bmatrix}.$$

This is quite unwieldy, but we can decide which symmetries should appear in our factorization by working over the residue field $\mathbb{C}$. Applying the evaluation map $B_{\mathfrak{P}} \to \mathbb{C}$ of Example 3.6.22, we get that T has matrix

$$\overline{T} = \begin{bmatrix} 0 & 0 & 1 \\ 0 & -1 & 0 \\ -1 & 0 & 0 \end{bmatrix}$$

over the residue field. Therefore, the $\mathbb{C}$-sesquilinear space $\overline{L}_\tau$ is the column space of

$$I - \overline{T} = \begin{bmatrix} 1 & 0 & -1 \\ 0 & 2 & 0 \\ 1 & 0 & 1 \end{bmatrix},$$

i.e., $\overline{L}_\tau = \mathbb{C}^3$. Since the evaluation homomorphism maps the Gram matrix H for L to the identity matrix, a Gram matrix $\overline{H}$ for the $\mathbb{C}$-sesquilinear form $\langle -, - \rangle_\tau$ on $\overline{L}_\tau$ has the same column space as

$$\overline{H} - \overline{H}\overline{T} = I - \overline{T} = \begin{bmatrix} 1 & 0 & -1 \\ 0 & 2 & 0 \\ 1 & 0 & 1 \end{bmatrix}.$$

So, as in Remark 4.8.17, the Gram matrix for $\overline{L}_\tau$ has full rank. Since L is unimodular, τ has anisotropic dimension $\dim_a(\tau) = 3$. We should then expect the algorithm of the Cartan-Dieudonné Theorem to give us a factorization of τ of length $2\,\mathrm{rk}(L) - \dim_a(\tau) = 3$. Our computation of $I - \overline{T}$ gives us that $\langle e_i, e_i - Te_i \rangle$ is a unit for every $1 \leq i \leq 3$. Since the trace $\mathcal{O}_E \to \mathcal{O}_K$ is surjective, we have all the ingredients necessary to form symmetries of L with respect to the vectors $z_1 = e_1 - \tau e_1$, $z_2 = e_2 - \tau e_2$, and $z_3 = e_3 - \tau e_3$. Put

$$\beta_1 := \frac{\langle z_1, z_1 \rangle}{2} = \frac{2X + 4Y}{XY + Y^2 + X + 2Y + 1},$$

$$\beta_2 = \frac{\langle z_2, z_2 \rangle}{2} = \frac{2X^2Y^2 + 4XY^3 + 2Y^4 + 12XY^2 + 2Y^2}{X^2Y + 3XY^2 + 2Y^3 + X^2 + 4XY + 4Y^2 + X + 2Y},$$

and

$$\beta_3 = \frac{\langle z_3, z_3\rangle}{2} = \frac{2X^2Y + 4XY^2 + 2Y^3 + 2X^2 + 4XY + 2Y^2}{X^2 + 3XY + 2Y^2 + X + 2Y}.$$

We can therefore define symmetries of L

$$S_1 = S_{z_1,\beta_1}, \quad S_2 = S_{z_2,\beta_2}, \text{ and } S_3 = S_{z_3,\beta_3}$$

with respect to each pair z_i, β_i.

As in the proof of the Theorem, $S_1 \circ \tau(e_1) = e_1$. Writing $N = \{e_1\}^\perp \subseteq L$, we therefore obtain

$$S_1 \circ \tau = I_{\langle e_1\rangle} \oplus \tau|_N.$$

We can do the same for $\tau|_N$: we get that $S_2 \circ \tau|_N(e_2) = S_2 \circ \tau(e_2) = e_2$, and since S_2 fixes e_1, we have that

$$S_2 \circ S_1 \circ \tau = I_{\langle e_1,e_2\rangle} \oplus \tau|_{\langle e_3\rangle}.$$

Lastly, $S_3 \circ \tau(e_3) = e_3$, and since $\text{span}\{e_1, e_2\} = \{e_3\}^\perp$, we get that

$$S_3 \circ S_2 \circ S_1 \circ \tau = I.$$

The inverse of a symmetry is a symmetry, so τ is a product of 3 symmetries.

4.10 Discussion and further directions

Let L be a Hermitian lattice over the quadratic separable extension of discrete valuation rings $\mathcal{O}_E/\mathcal{O}_K$. Let ν be the discrete valuation on $\mathcal{O}_E\backslash\{0\}$. In Theorem 4.9.1, we were able to show that an arbitrary isometry U of $\mathrm{U}(L)$ can be written as a product of symmetries, and, provided that if U is strictly diagonalizable, that there exists a factorization of length $2\,\mathrm{rk}\,(L_U) - \dim_a(U, [L])$, where $[L]$ is the Jordan decomposition induced by the orthogonal basis $\{x_1, \ldots, x_n\}$ of eigenvectors of U for L. We have the following Corollary to Theorem 4.9.1.

Corollary 4.10.1. Suppose that $U \in \mathrm{U}(L)$ is strictly diagonalizable with eigenbasis $\{x_1, \ldots, x_n\}$, and further, that $\nu(\langle x_j, x_j - Ux_j\rangle) = \nu(\langle x_j, x_j\rangle)$ for any $1 \leq j \leq n$ such that $Ux_j \neq x_j$. Any minimal factorization of U in terms of symmetries has length $2\,\mathrm{rk}\,(L_U) - \dim_a(U) = \mathrm{rk}\,(L_U)$.

Proof. Let $\{x_1, \ldots, x_n\}$ be the orthogonal basis of eigenvectors of U for L. The hypothesis that U is strictly diagonalizable means that $\{x_i - Ux_i \mid x_i - Ux_i \neq 0\}$ is a basis for L_U. Since the orthogonal basis of eigenvectors induces a Jordan-invariant decomposition for L, we have that the rank of L_U is the sum

$$\mathrm{rk}\,(L_U) = \sum_{i \in \mathbb{Z}} \mathrm{rk}\,((\mathrm{id} - U)\, L_i)\,.$$

On the other hand, the condition that

$$\nu\,(\langle x_j, x_j - Ux_j\rangle) = \nu\,(\langle x_j, x_j\rangle)$$

for all $1 \leq j \leq n$ such that $x_j \neq Ux_j$, together with the fact that $(\mathrm{id} - U)\, L_i$ has basis

$$\{x_j - Ux_j \mid x_j \neq Ux_j, \nu\,(\langle x_j, x_j\rangle) = i, 1 \leq j \leq n\}$$

shows that the resulting Gram matrix for $\langle -, -\rangle_U$ restricted to $(\mathrm{id} - U)\, L_i$ is diagonal and invertible over $\mathfrak{P}^i/\mathfrak{P}^{i+1} \cong e$. Therefore, $\dim_a (U, L_i) = \mathrm{rk}\,((\mathrm{id} - U)\, L_i)$. Combining the above facts, we have that $\dim_a (U, [L]) = \mathrm{rk}(L_U)$. But $\dim_a (U, [L]) \leq \mathrm{rk}(L_U)$ for any Jordan-invariant decomposition $[L]$ of L, seeing as

$$L_U = \bigoplus_{i \in \mathcal{I}} (\mathrm{id} - U)\, L_i.$$

Therefore, we have that

$$\mathrm{rk}\,(L_U) = \dim_a (U, [L]) \leq \dim_a (U) \leq \mathrm{rk}\,(L_U)\,,$$

forcing all inequalities to be equalities. The proof of Theorem 4.9.1 then gives a factorization in terms of $2\,\mathrm{rk}\,(L_U) - \dim_a (U, [L]) = \mathrm{rk}\,(L_U)$ symmetries of L.

Recall that L sits inside the E-Hermitian space $V = E \otimes_{\mathcal{O}_E} L$ and that U is an element of $\mathrm{U}(V)$. Since E is the fraction field of $\mathcal{O}_E$, we have that $\dim_E (V_U) = \mathrm{rk}\,(L_U)$. By [Gro02], we have that any minimal factorization of U as a product of symmetries of V has length $\dim_E (V_U)$. Seeing as a factorization of U as a product of symmetries of L is also a factorization of U as a product of symmetries of V, we have that the decomposition in terms of symmetries of L that we previously obtained is of minimal length $\mathrm{rk}\,(L_U) = \dim_E (V_U)$. ∎

Let us discuss the proof of the Corollary, and investigate what we would need to do to obtain a generalization. Suppose first that $U \in \mathrm{U}(L)$ is strictly diagonalizable, but that we relax the hypothesis that each element x_j of the orthogonal basis

of eigenvectors satisfies $\nu(\langle x_j, x_j - Ux_j\rangle) = \nu(\langle x_j, x_j\rangle)$. Indeed, suppose that x_j satisfies $i = \langle x_j, x_j\rangle < \langle x_j, x_j - Ux_j\rangle$. Then, the anisotropic dimension $\dim_a(U, L_i)$ of U with respect to L_i is strictly less than $\mathrm{rk}((\mathrm{id} - U) L_i)$. From this, we cannot conclude that the bound given by $\dim_a(U, [L]) \leq \dim_a(U)$ is met. Moreover, the lower bound for the length of a factorization over the field is no longer tight, seeing as $\dim_E(V_U) \neq \dim_a(U, [L])$. Therefore, the proof above fails in this case. Nonetheless, it seems reasonable to attempt to find a proof for the following result, as obtaining an orthogonal eigenbasis for a lattice is a very strong condition. In particular, it is easy to count how many isometries appear in the algorithm of Theorem 4.9.1 when we apply it to U and its orthogonal eigenbasis. The eigendecomposition of L induces an eigendecomposition of each modular component appearing in a Jordan-invariant decomposition of L, so that the anisotropic dimension of U is determined by computation on the eigenbasis. In other words, it remains to prove that the anisotropic dimension determines how many isometries appear in a minimal factorization of U, which would require a new technique.

Conjecture 4.10.2. *Any minimal factorization of a strictly diagonalizable isometry U of L has length $2\,\mathrm{rk}(L_U) - \dim_a(U)$. In other words, the bound of Theorem 4.9.1 is met.*

The next possible relaxation consists of admitting an arbitrary Jordan-semisimple isometry of L. This is analogous to introducing the Jordan canonical form of a linear operator as a generalization of diagonalization. When the isometry is Jordan-semisimple, its anisotropic dimension behaves well with respect to the orthogonal components, in the sense that $\sum_{i\in\mathcal{I}} \dim_a(U, L_i) \leq \mathrm{rk}(L_U)$. It then seems plausible to write the following.

Conjecture 4.10.3. *Any minimal factorization of a Jordan-semisimple isometry U of L has length $2\,\mathrm{rk}(L_U) - \dim_a(U)$.*

When we no longer assume that U is Jordan-semisimple, we have the following obstacle. The anisotropic dimension of a Jordan-semisimple isometry of L (with respect to a Jordan-invariant decomposition $[L] = \bigoplus_{i\in\mathcal{I}} L_i$ is defined as

$$\dim_a(U, [L]) = \sum_{i\in\mathcal{I}} \dim_a(U, L_i),$$

where $\dim_a(U, L_i)$ is the rank of $\langle -, - \rangle_U$ restricted to the image of $\mathrm{id} - U$ on L_i. In the cases we have examined, this number is exactly the number of steps of the algorithm of Theorem 4.9.1 which yields a single symmetry, effectively counting the total number of symmetries which appear. Now, when we relax the assumption that each L_i is U-invariant, we no longer have a direct sum decomposition of L_U mirroring the modularity of a Jordan-invariant decomposition of L. So, if we attempt to define anisotropic dimension as the sum of the ranks of $\langle -, - \rangle_U$ on each $\mathrm{id} - U|_{L_i}$, we possibly obtain a quantity strictly greater than $\mathrm{rk}(L_U)$. For example, we cannot easily compare the quantity $\nu(\langle x, x - Ux \rangle)$ with $\nu(\langle x, x \rangle)$ for a given preimage x for $x - Ux$, as there is no canonical choice for $\nu(\langle x, x \rangle)$. The next step would be find a way to consider this modified definition of rank of $\langle -, - \rangle_U$, up to well-chosen elements of the kernel of $\mathrm{id} - U$, in a way that this quantity is equal to the number of inductive steps of the algorithm of Theorem 4.9.1 which yield exactly one symmetry. An invariant of this nature could then potentially be used to determine the length of a minimal factorization of an isometry $U \in \mathrm{U}(L)$, but care needs to be taken to come up with a satisfactory definition.

Bibliography

[AM69] Michael Atiyah and Ian G. Macdonald. *Introduction to Commutative Algebra*. Addison-Wesley Series in Mathematics, Addison-Wesley, Reading, Massachusetts, 1969.

[Ash02] Robert Ash. *Abstract Algebra: The Basic Graduate Year*. Online edition, 2002.

[BH23] Simon Brandhorst and Tommy Hofmann. *Finite subgroups of automorphisms of K3 surfaces*. Forum of Math. Sigma, **e54**, 2023.

[BHM22] Simon Brandhorst, Tommy Hofmann and Sven Manthe. *Generation of Local Unitary Groups*. Journal of Algebra, **609**, pp.484-513, 2022.

[Cas86] J.W.S. Cassels. *Local Fields*. London Mathematical Society Student Texts, Volume 3, Cambridge University Press, Cambridge, United Kingdom, 1986.

[DF04] David S. Dummit and Richard M. Foote. *Abstract Algebra*. 3^{rd} Edition, John Wiley and Sons, Inc., Hoboken, New Jersey, 2004.

[Eis95] David Eisenbud. *Commutative Algebra With a View Toward Algebraic Geometry*. Graduate Texts in Mathematics, Volume 150, Springer New York, NY, 1995.

[Fo17] Timothy J. Ford. *Separable Algebras*. Graduate Studies in Mathematics, Volume 183, of the American Mathematical Society, Providence, Rhode Island, 2017.

[Fu89] William Fulton. *Algebraic Curves*. Advanced Book Classics, Addison-Wesley Publishing Company, Redwood City, 1989.

[Gro02] Larry C. Grove. *Classical Groups and Geometric Algebra*. Graduate Studies in Mathematics, Volume 39, of the American Mathematical Society, Providence, Rhode Island, 2002.

[Hay68] Keizo Hayakawa. *Generation of local integral unitary groups over an unramified dyadic local field.* J. Fac. Sci. Univ. Tokyo, **15**(1), pp.1-11, 1968.

[Jac62] Ronald Jacobowitz. *Hermitian forms over Local Fields*. American Journal of Mathematics, **84**(3), pp. 441-465, 1962.

[Knu91] Max-Albert Knus. *Quadratic and Hermitian Forms over Rings*. Grundlehren der mathematischen Wissenschaften, Volume 294, Springer-Berlin, Heidelberg, 1991.

[Kob84] Neal Koblitz. *p-adic Numbers, p-adic Analysis, and Zeta-Functions.* Graduate Texts in Mathematics, Volume 58, Springer New York, NY, 1984.

[Lam05] T.Y. Lam. *Introduction to Quadratic Forms over Fields*. Graduate Studies in Mathematics, Volume 67, of the American Mathematical Society, Providence, Rhode Island, 2005.

[Lor07] Falko Lorenz. *Volume II: Fields with Structure, Algebras and Advanced Topics*. Springer New York, New York, New York, 2007.

[Mil20] J.S. Milne. *Algebraic Number Theory*. Online edition, 2020.

[Nag62] Nagata, M. *Local rings*. Interscience Tracts in Pure and Applied Mathematics, Volume 13, of Interscience Publishers, New York, 1962.

[Neu99] Jürgen Neukirch. *Algebraic Number Theory*. Grundlehren der mathematischen Wissenschaften, Volume 322, Springer Berlin, Heidelberg, 1999.

[OMe57] O.T. O'Meara. *Integral Equivalence of Quadratic Forms in Ramified Local Fields*. American Journal of Mathematics, Vol. 79, No. 1, pp.157-186, 1957.

[OMe73] O.T. O'Meara. *Introduction to Quadratic Forms*. Grundlehren der mathematischen Wissenschaften, Volume 117, Springer Berlin, Heidelberg, 1973.

[Sc50] Peter Scherk. *On the Decomposition of Orthogonalities into Symmetries*. Proceedings of the American Mathematical Society, 1(4), 481–491, 1950.

[Se79] Jean-Pierre Serre. *Local fields*. Graduate Texts in Mathematics, Volume 67, Springer New York, NY, 1984.

[Voi21] John Voight.*Quaternion Algebras*. Graduate Texts in Mathematics, Volume 288, Springer Cham, 2021.

[Wa59] Gordon E. Wall. *The structure of a unitary factor group*. Publications mathématiques de l'I.H.É.S., **1**, pp.7-23, 1959.

[Wi19] Sarah J. Witherspoon. *Hochschild Cohomology for Algebras*. Graduate Studies in Mathematics, Volume 204, of the American Mathematical Society, Providence, Rhode Island, 2019.