

Algorithms for finite rings

Iuliana Ciocanea Teodorescu

▶ **To cite this version:**

Iuliana Ciocanea Teodorescu. Algorithms for finite rings. General Mathematics [math.GM]. Université de Bordeaux; Universiteit Leiden (Leyde, Pays-Bas), 2016. English. NNT : 2016BORD0121 . tel-01378003v2

HAL Id: tel-01378003

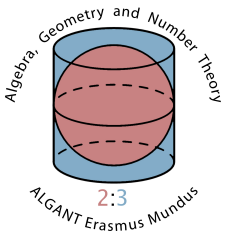
https://theses.hal.science/tel-01378003v2

Submitted on 14 Oct 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.





université de BORDEAUX

THÈSE EN COTUTELLE PRÉSENTÉE

POUR OBTENIR LE GRADE DE

DOCTEUR DE

L'UNIVERSITÉ DE BORDEAUX ET L'UNIVERSITÉ DE LEIDEN

ÉCOLE DOCTORALE DE MATHÉMATIQUES ET INFORMATIQUE
ÉCOLE DOCTORALE DES SCIENCES

SPECIALITÉ : Mathématiques Pures

par Iuliana CIOCĂNEA-TEODORESCU

Algorithmes pour les anneaux finis

Sous la direction de Karim BELABAS
et de Hendrik W. LENSTRA

Soutenue le 22 juin 2016.

Membres du jury:

LENSTRA, Hendrik W.	Professeur	Universiteit Leiden	Directeur
BELABAS, Karim	Professeur	Université de Bordeaux	Directeur
KRICK, Teresa	Professeur	Universidad de Buenos Aires	Rapporteur
TAELMAN, Lenny	Professeur	Universiteit van Amsterdam	Rapporteur
BIESEL, Owen	Docteur	Universiteit Leiden	Examinateur
DE SMIT, Bart	Professeur	Universiteit Leiden	Examinateur
VAN DER KALLEN, Wilberd	Docteur	Universiteit Utrecht	Examinateur

Titre: Algorithmes pour les anneaux finis

Résumé: Cette thèse s'attache à décrire des algorithmes qui répondent à des questions provenant de la théorie des anneaux et des modules. Nous restreindrons essentiellement notre étude à des algorithmes déterministes, en temps polynomial, ainsi qu'aux anneaux et modules finis.

Le premier des principaux résultats de cette thèse concerne le problème de l'isomorphisme entre modules : nous décrivons deux algorithmes distincts qui, étant donné un anneau fini R et deux R-modules M et N finis, déterminent si M et N sont isomorphes. S'ils le sont, les deux algorithmes exhibent un tel isomorphisme.

De plus, nous montrons comment calculer un ensemble de générateurs de taille minimale pour un module donné, et comment construire des couvertures projectives et des enveloppes injectives. Nous décrivons ensuite des tests mettant en évidence le caractère simple, projectif ou injectif d'un module, ainsi qu'un test constructif de l'existence d'un homomorphisme de modules surjectif entre deux modules finis, l'un d'entre eux étant projectif. Par contraste, nous montrons le résultat négatif suivant : le problème consistant à tester l'existence d'un homomorphisme de modules injectif entre deux modules, l'un des deux étant projectif, est NP-complet.

La dernière partie de cette thèse concerne le problème de l'approximation du radical de Jacobson d'un anneau fini. Il s'agit de déterminer un idéal bilatère nilpotent tel que l'anneau quotient correspondant soit "presque" semi-simple. La notion de "semi-simplicité approchée" que nous utilisons est la *séparabilité*.

Mots clés: algorithmes déterministes, anneaux finis, modules finis, isomorphisme, radical de Jacobson, semi-simplicité, séparabilité.

Title: Alorithms for finite rings

Abstract: In this thesis we are interested in describing algorithms that answer questions arising in ring and module theory. Our focus is on deterministic polynomial-time algorithms and rings and modules that are finite.

The first main result of this thesis concerns the module isomorphism problem: we describe two distinct algorithms that, given a finite ring R and two finite R-modules M and N, determine whether M and N are isomorphic. If they are, the algorithms exhibit such an isomorphism.

In addition, we show how to compute a set of generators of minimal cardinality for a given module, and how to construct projective covers and injective hulls. We also describe tests for module simplicity, projectivity, and injectivity, and constructive tests for existence of surjective module homomorphisms between two finite modules, one of which is projective. As a negative result, we show that the problem of testing for existence of injective module homomorphisms between two finite modules, one of which is projective, is NP-complete.

The last part of the thesis is concerned with finding a good working approximation of the Jacobson radical of a finite ring, that is, a two-sided nilpotent ideal such that the corresponding quotient ring is "almost" semisimple. The notion we use to approximate semisimplicity is that of *separability*.

Keywords: deterministic algorithms, finite rings, finite modules, isomorphism, Jacobson radical, semisimplicity, separability.

L'Institut de Mathmatiques de Bordeaux

351, cours de la Libration - F 33 405 Talence, FRANCE

ALGORITHMS FOR FINITE RINGS

Proefschrift
ter verkrijging van
de graad van Doctor aan de Universiteit Leiden
op gezag van Rector Magnificus prof. mr. C.J.J.M. Stolker,
volgens besluit van het College voor Promoties
te verdedigen op woensdag 22 juni 2016
klokke 11:15 uur

door

Iuliana Ciocănea-Teodorescu
geboren te Boekarest, Roemenië
in 1990

Promotores: Prof. dr. Hendrik W. Lenstra (Universiteit Leiden)

Prof. dr. Karim Belabas (Université de Bordeaux)

Samenstelling van de promotiecommissie:

Dr. Owen Biesel (Universiteit Leiden)

Prof. dr. Bart de Smit (Universiteit Leiden)

Prof. dr. Teresa Krick (Universidad de Buenos Aires)

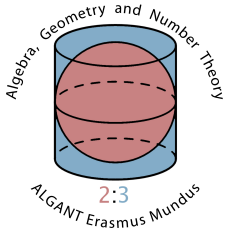
Prof. dr. Lenny Taelman (Universiteit van Amsterdam)

Dr. Wilberd van der Kallen (Universiteit Utrecht)

Prof. dr. Aad van der Vaart (Universiteit Leiden)

This work was funded by Algant-Doc Erasmus Mundus and was carried out at Universiteit Leiden and l'Université de Bordeaux.





université de BORDEAUX

"Once [the reader] explicitly gives up all practical claims, he will realize that he can occupy himself with algorithms without having to fear the bad dreams caused by the messy details and dirty tricks that stand between an elegant algorithmic idea and its practical implementation. He will find himself in the platonic paradise of pure mathematics, where a conceptual and concise version of an algorithm is valued more highly than an ad hoc device that speeds it up by a factor of ten and where words have precise meanings that do not change with the changing world. (...) And in his innermost self he will know that in the end his own work will turn out to have the widest application range, exactly because it was not done with any specific application in mind."

H.W. Lenstra. Algorithms in Algebraic Number Theory (1992). *BAMS*, 26: 211–244

"If $\mathrm{P} = \mathrm{NP}$*, then the world would be a profoundly different place than we usually assume it to be. There would be no special value in creative leaps, no fundamental gap between solving a problem and recognizing the solution once it's found. Everyone who could appreciate a symphony would be Mozart; everyone who could follow a step-by-step argument would be Gauss (...)."*

Scott Aaronson. Personal blog: www.scottaaronson.com/blog/ (2006)

I died for beauty, but was scarce
Adjusted in the tomb,
When one who died for truth was lain
In an adjoining room.

Emily Dickinson. Fr 448, J 449 (1890)

Contents

Introduction **i**

List of symbols **v**

1 Background **1**
1.1 Algorithms and complexity 1
1.2 Basic ring theory 4
1.3 Basic module theory 5
1.4 More ring theory 7
1.5 Idempotents 9
1.6 More module theory 10
1.7 Quasi-Frobenius rings 15
1.8 Frobenius algebras and symmetric algebras 16
1.9 Duality 17

2 Linear algebra over $\mathbb{Z}$: basic algorithms for finite abelian groups **19**
2.1 Lattices 20
2.2 Hermite and Smith normal forms 22
2.3 Representing objects and basic constructions 26
2.4 Homomorphism groups and tensor products 33
2.5 Splitting exact sequences 34
2.6 Torsion subgroups, exponents, orders, cyclic decompositions 35
2.7 Homomorphism groups and tensor products reconsidered 39
2.8 Projective $\mathbb{Z}/m\mathbb{Z}$-modules 40

3 Linear algebra over $\mathbb{Z}$: basic algorithms for finite rings **45**
3.1 Representing objects and basic constructions 46
3.2 Computations with ideals 48
3.3 Computing the centre and the prime subring of a finite ring 49
3.4 Computing the Jacobson radical 50
3.5 Other known algorithms and open questions 50

4 The module isomorphism problem **53**
4.1 Introduction . . . 53
4.2 Context . . . 55
4.3 MIP via non-nilpotent endomorphisms . . . 56
4.4 MIP via an approximation of the Jacobson radical . . . 59
4.5 Remark on implementation and performance . . . 63

5 A miscellaneous collection of algorithms **65**
5.1 Testing if a ring is a field . . . 65
5.2 Testing if a ring is simple . . . 66
5.3 Testing if a module is simple . . . 67
5.4 Testing if a module is projective . . . 67
5.5 Constructing projective covers . . . 68
5.6 Constructing injective hulls . . . 69
5.7 Testing if a module is injective . . . 70
5.8 Testing if a ring is quasi-Frobenius . . . 70
5.9 Constructive tests for existence of injective and surjective module homomorphisms . . . 70

6 Approximating the Jacobson radical of a finite ring **75**
6.1 Introduction . . . 75
6.2 Separability . . . 76
6.3 An approximation of the Jacobson radical . . . 96
6.4 Computing the generalised prime subring . . . 111

Bibliography **115**

Index **123**

Abstract **125**

Résumé **126**

Samenvatting **127**

Acknowledgements **129**

CV **130**

Introduction

Throughout this text, rings are assumed to contain a unit element, but are not necessarily commutative. Modules are always left-unital, unless otherwise specified.

The main goal of this PhD thesis is to develop a toolbox for working with finite rings and finite modules within algorithms. The motivation to study problems concerning finite rings and finite modules is twofold. The first reason is a theoretical one and stems from the fundamental nature of the problems that arise. Since we are mostly interested in viewing algorithms as mathematical objects in their own right, the focus will be on deterministic polynomial-time algorithms. The second reason to study these problems refers to the necessity of having as many algorithms as possible available in computer algebra systems to deal with finite rings.

The first chapter of this thesis contains the necessary background theory on algorithms, complexity, rings and modules. Chapters 2 and 3 contain a series of basic algorithms for finitely generated abelian groups and finite rings. These will be used implicitly and extensively in the rest of the algorithms described.

The first algorithmic problem we tackle is the module isomorphism problem. The *module isomorphism problem* can be formulated as follows: design a deterministic algorithm that, given a ring R and two left R-modules M and N, decides in polynomial time whether they are isomorphic, and if yes, exhibits an isomorphism.

Isomorphism problems are some of the most natural algorithmic questions. Given two objects of the same nature, we would like to be able to tell if they are isomorphic, and if so, we would ideally also want to produce an isomorphism. Objects for which isomorphism problems have been extensively studied include graphs, groups and rings. The easy formulation of these problems and their fundamental nature does not however entail that they have a trivial solution. In fact, for many problems of this type, no deterministic polynomial-time algorithms are known ([11, 52, 53]).

Two intermediate results, valuable in themselves, are proved in Chapter 4:

Theorem. *There exists a deterministic polynomial-time algorithm that, given a finite ring R and two finite R-modules M and N, computes a maximum length R-module C that is isomorphic to a direct summand both of M and of N. Moreover, the algorithm computes direct complements of C both in M and in N, together with the corresponding isomorphisms.*

Theorem. *There exists a deterministic polynomial-time algorithm that, given a finite ring R and a finite R-module M, computes a set of generators for M of minimum cardinality.*

Both of these theorems can be used to provide a solution for the module isomorphism problem.

Theorem. *There exists a deterministic polynomial-time algorithm that, given a finite ring R and two finite R-modules M and N, decides whether M and N are isomorphic, and if they are, exhibits an isomorphism.*

Chapter 5 contains a collection of deterministic polynomial-time algorithms for testing properties of rings and modules.

Theorem. *There exists a deterministic polynomial-time algorithm that, given a finite ring R and a finite R-module M, tests whether M is*

(i) *projective,*
(ii) *injective,*
(iii) *simple.*

Theorem. *There exists a deterministic polynomial-time algorithm that, given a finite ring R, tests whether R is*

(i) *simple,*
(ii) *quasi-Frobenius.*

Moreover,

Theorem. *There exists a deterministic polynomial-time algorithm that, given a finite ring R and a finite R-module M, constructs a projective cover and an injective hull of M.*

We also discuss the algorithmic problem of constructively testing for existence of injective and surjective homomorphisms between two finite length modules over a ring R, i.e. the problem of testing for existence and finding such homomorphisms when they do exist. If R is a finite-dimensional algebra over a field, this problem can be cast in the context of matrix completion, and has been shown to be NP-hard. We consider the case where R is a finite ring and one of the modules is either projective or injective over R.

Theorem. *There exists a deterministic polynomial-time algorithm that, given a finite ring R and two finite R-modules M and N, one of which is R-projective, constructively tests for existence of a surjective R-module homomorphism $M \twoheadrightarrow N$.*

Dually:

Theorem. *There exists a deterministic polynomial-time algorithm that, given a finite ring R and two finite R-modules M and N, one of which is R-injective, constructively tests for existence of an injective R-module homomorphism $M \hookrightarrow N$.*

For the remaining cases we obtain negative results:

Theorem. *The problem of deciding existence of an injective module homomorphism between two modules over a finite ring, one of which is projective over that ring, is* NP*-complete.*

A very important class of rings is that of *semisimple rings*. Let R be a ring and M an R-module. Then M is said to be *semisimple* if every R-submodule of M has a direct complement in M. A ring R is said to be *semisimple* if the left-regular (or equivalently right-regular) module is semisimple. Semisimple rings have a lot of structure: everything breaks down in an orderly fashion. Moreover, the Wedderburn theorem gives a complete classification of such rings as finite products of matrix rings over division rings.

The notion of semisimplicity is inextricably linked to that of the *Jacobson radical* of a ring, defined as the intersection of all maximal left ideals. The Jacobson radical of a ring R is a two-sided ideal, and we denote it by $\mathrm{J}(R)$. The rings R and $R/\mathrm{J}(R)$ have the same simple left modules, which suggests that a study of $R/\mathrm{J}(R)$ will reveal much of the structure of R. Moreover, if R is left-artinian, then $\mathrm{J}(R)$ is a nilpotent ideal of R and R is semisimple if and only if $\mathrm{J}(R) = 0$.

When trying to answer questions about left-artinian rings and modules over them, it is often convenient to reduce the problem at hand to the semisimple case, where structures are much more manageable, and then "lift". This places the computation of the Jacobson radical at the heart of many problems. While it can be done deterministically in polynomial time for matrix algebras over a field [15, 18, 27, 75], we cannot expect to have a deterministic polynomial-time algorithm for the general case, since the problem ultimately reduces to finding the squarefree part of an integer (consider the ring $\mathbb{Z}/n\mathbb{Z}$, for some $n \in \mathbb{Z}_{>0}$). In Chapter 6, we attempt to deterministically construct approximations of the Jacobson radical of a finite ring that are "satisfactory" for many practical purposes, that is, two-sided nilpotent ideals such that when we quotient the ring by them, we are left with something that is "almost" semisimple.

The notion used to approximate semisimplicity is that of separability. Given a commutative ring R, an R-algebra S is said to be *separable* over R if S is projective as an $S \otimes_R S^{\mathrm{o}}$-module, where S^{o} denotes the opposite ring of S. A ring is said to be separable if it is separable as a $\mathbb{Z}$-algebra.

Definition. *Let A be a finite ring and $\mathfrak{j}_A \subset A$ an ideal. We say $\mathfrak{j}_A$ is an* approximation of the Jacobson radical *of A if*

(A1) *$\mathfrak{j}_A$ is a two-sided nilpotent ideal of A,*
(A2) *$A/\mathfrak{j}_A$ is finite separable,*
(A3) *$A/\mathfrak{j}_A$ is projective as a module over its prime subring.*

The resulting ring, $A/\mathfrak{j}_A$, has many good properties, e.g. it has "many" projective modules (via *projectivity lift*), it is quasi-Frobenius, it is isomorphic to its opposite as rings. Moreover, finite separable rings can be classified as finite products of matrix rings over certain commutative rings. We show that approximations of Jacobson radicals can be efficiently computed.

Theorem. *There exists a deterministic polynomial-time algorithm that, given a finite ring A, computes an approximation of the Jacobson radical of A.*

We are interested in deterministic polynomial-time algorithms that produce approximations of the Jacobson radical of a finite ring and have the additional property that, when run on two isomorphic rings, they output isomorphic approximations of their Jacobson radicals, even when the ring isomorphism is unknown.

In fact, we exhibit not one, but two algorithms as described by the above theorem. If we denote by $\mathcal{F}$ the class of finite rings, then the two families of ideals $(\mathfrak{j}_A)_{A\in\mathcal{F}}$ and $(\mathfrak{j}'_A)_{A\in\mathcal{F}}$, produced by the two algorithms are functorial under isomorphisms, i.e. if $\phi : A \to B$ is an isomorphism of finite rings, then $\phi(\mathfrak{j}_A) = \mathfrak{j}_B$ and $\phi(\mathfrak{j}'_A) = \mathfrak{j}'_B$.

List of symbols

$\mathbb{F}_q$	finite field with q elements
$\mathbb{Z}$	ring of integers
$\mathbb{Z}_{>0}$	set of positive integers
$\mathbb{Z}_{\geq 0}$	set of non-negative integers
$\lvert S\rvert$	cardinality of a set S
$\mathcal{M}_n(R)$	ring of $n \times n$ matrices with entries in the ring R
$\mathcal{M}_{n\times m}(R)$	set of $n \times m$ matrices with entries in the ring R
R^{e}	the enveloping algebra of R (page 76)
R^{o}	the opposite ring of the ring R
$R[G]$	the group ring of G over R
$\mathrm{Max}(R)$	set of maximal ideals of a commutative ring R
$\mathrm{Spec}(R)$	set of prime ideals of a commutative ring R
$\mathrm{J}(R)$	Jacobson radical of a ring R
$\mathrm{Z}(R)$	centre of a ring R
$\mathrm{char}(R)$	characteristic of a ring R
$\mathrm{rad}(n)$	the product of all primes dividing an integer n
$\hookrightarrow$	injective map
$\twoheadrightarrow$	surjective map
$M \otimes_R N$	tensor product of M with N over R
$\mathrm{End}_R(M)$	ring of R-endomorphisms of M
$\mathrm{Hom}_R(M, N)$	group of R-homomorphisms from M to N
$\mathfrak{M}_R^{\mathrm{fg}}$	category of finitely generated right R-modules
$\mathfrak{M}_R$	category of right R-modules

${}^{\mathrm{fg}}_{R}\mathfrak{M}$	category of finitely generated left R-modules
${}_{R}\mathfrak{M}$	category of left R-modules
${}_{R}M_{S}$	R-S-bimodule M

Chapter 1

Background

This chapter introduces the terminology that will be used throughout the rest of the text. The first section contains a brief discussion about algorithms and complexity, followed by a list of examples of basic algorithmic questions (primality testing, integer factorisation, coprime factorisation). The remaining sections review basic facts of ring and module theory. We will focus on those results that are specific to noncommutative ring theory.

The main references for this chapter are: [66, 73], for the section concerning algorithms, and [56, 57, 58, 60], for the rest.

1.1 Algorithms and complexity

For an entirely formal discussion of algorithms and complexity, one needs to enter the realm of theoretical computer science jargon. Fortunately, however, this can be avoided, since it so happens that the intuitive notions we have of algorithms, "hardness" of a computational problem, "efficiency" etc., are enough for a meaningful discussion, and complexity theory appears to be "robust" enough to allow us to work with them.

Formally, an algorithm is a *Turing machine*. Intuitively, an algorithm is a sequence of steps that takes as *input* a finite sequence of nonnegative integers and produces an *output* in the form of another finite sequence of nonnegative integers. An integer is represented inside an algorithm by a string of *bits*, and a step in the algorithm is then a *bit operation*. It is also useful to have a notion of the "size" of an input. If $n \in \mathbb{Z}_{\geq 0}$, then the *length* of n is taken to be $\text{length}(n) := \log_2(n+2)$, reflecting the number of bits required to write n down in binary. The length of a negative integer m is $1 + \text{length}(|m|)$ and the length of an input is the sum of the lengths of the integers that compose it.

We would like to study the number of steps needed for an algorithm to perform a certain task. The *running time* represents the number of steps required to produce

an output. An algorithm is said to be *polynomial-time* if its running time is bounded above by a polynomial expression in the length of the input. The running time of an algorithm is often referred to as the *complexity of the algorithm*. In our case, this is the *bit-complexity*, as opposed to e.g. the *arithmetic complexity*, where a step is taken to be an arithmetic operation.

Naturally, we are interested in more than just performing arithmetic in $\mathbb{Z}$. However, virtually any mathematical object of interest can be *encoded* as a sequence of nonnegative integers. For the objects we are interested in, we will see exactly how to do this in the following two chapters.

Throughout this text, we will be exclusively interested in *deterministic polynomial-time algorithms*, i.e. algorithms in the running of which no random bit is generated. While allowing for probabilistic algorithms (e.g. *Las Vegas* or *Monte Carlo* algorithms) leads in practice to increased efficiency, these algorithms reveal less about the intrinsic difficulty of the problem at hand and are thus of less theoretical interest. We shall not think about them.

Furthermore, we will be content with being able to declare a certain algorithm as running in polynomial time, without computing exact exponents. The main reason for this is that we have not conceived the algorithms presented in this thesis with the intention of also implementing them. Therefore, there are countless improvements and randomised variations possible, which we have chosen not to explore in detail. Computing running times of an algorithm that is deliberately non-optimal seems futile.

Algorithms are often thought of as auxiliary objects, whose main reason for existence is to facilitate experimentation within computer algebra systems, with the purpose of confirming or invalidating hypotheses formulated in a more theoretical setting, providing examples or guiding the mathematician's intuition. In these cases, one is rarely interested in the "intrinsic" difficulty of a problem. Instead, one usually focuses one's attention to a very particular instance of a problem and only desires that the algorithm used to solve it output a result in a "reasonable" amount of time.

Under this paradigm, our preference for deterministic polynomial-time algorithms seems at least odd and perhaps even outdated. However, the viewpoint that we adopt in this thesis is that algorithms are mathematical objects *per se*, worthy of independent study. The fact that a problem can be solved deterministically in polynomial time says that the problem is not intrinsically difficult or mysterious.

1.1.1 Complexity classes

After fixing the model of computation, we may wish to classify problems based on the rate at which they use up a certain resource, e.g. time. This gives rise to *complexity classes*.

Within complexity classes, we can order the problems according to their difficulty by using *reductions*. A reduction from a problem Q to a problem P is an intermediate algorithm that, given a solution to a problem P, produces a solution to another

problem Q. We say Q reduces to P. This formulation suggests that problem P is "at least as hard" as Q. Intuitively, a reduction has to be an "easy" computation. We will mainly be interested in reductions that are deterministic polynomial-time algorithms.

The problems that are maximal elements with respect to the partial ordering induced by reductions are said to be *complete* for that complexity class. These problems capture the difficulty of the entire class. Moreover, the existence of a "natural" complete problem in a complexity class guarantees that the class is not "artificial".

The most important complexity classes are listed below, together with informal descriptions:

1. P: consists of problems that can be solved by a deterministic polynomial-time algorithm;

2. NP: consists of problems whose solutions can be verified deterministically in polynomial time;

3. NP-hard: a problem A is NP-hard if every problem B in NP can be reduced to A;

4. NP-complete: consists of problems that are both in NP and NP-hard.

Clearly P $\subseteq$ NP. The question whether the reverse inclusion holds is at this time one of the most important open problems in theoretical computer science.

If P $\neq$ NP, then there exist problems that are in NP, but are neither NP-complete, nor in P (see [73], Theorem 14.1). These are called NP-*intermediate* problems. However, no "natural" NP-intermediate problems are known.

1.1.2 Integer factorisation, coprime factorisation and primality testing

Perhaps the simplest question one might ask oneself is, if given a positive integer, whether one can find a factorisation into primes. Despite its fundamental nature, the problem of integer factorisation is notoriously difficult, which has made it the heart of many algorithms used in cryptography. It is easy to see that integer factorisation lies in the complexity class NP. However, no deterministic polynomial-time algorithm for it is known. It is also not thought to be NP-complete, and is hence considered to be a candidate for the NP-intermediate class. There is an extensive literature devoted to a large variety of algorithms for integer factorisation (see e.g. [13, 62]).

A similar and related problem is that of finding square divisors of a given integer, for which there is also no known deterministic polynomial-time algorithm (see [59] or [12], Section 7.1).

Factoring into primes is out of our reach. However, given a set of integers, we can simultaneously factor them into "coprime" factors.

Definition 1.1.1 ([8], Section 4,7)**.** *Let S be a finite set of positive integers. A* coprime base *for S is a set of positive integers B such that:*

(i) $1 \notin B$,
(ii) *elements of* B *are pairwise coprime,*
(iii) *each element of* S *can be written as a product of powers of elements of* B.

Theorem 1.1.2 ([8], Algorithm 18.1)**.** (Coprime Base Algorithm) *There exists a deterministic polynomial-time algorithm that takes as input a finite set of positive integers* S *and outputs a coprime base* B *for* S*, and a factorisation of each element of* S *into products of powers of elements of* B.

Furthermore, primality testing has been shown to be in P.

Theorem 1.1.3 ([1])**.** *There exists a deterministic polynomial-time algorithm that, given* $n \in \mathbb{Z}_{>1}$*, determines if* n *is prime.*

1.2 Basic ring theory

Definition 1.2.1. *A* ring *is a triple* $(R, +, \cdot)$*, where* R *is a set and* $+, \cdot : R \times R \to R$ *are binary operations such that:*

(R1) $(R, +)$ *is an abelian group,*
(R2) $(R, \cdot)$ *is a monoid, i.e. the operation* $\cdot$ *is associative and has an identity element,*
(R3) *for all* $x, y, z \in R$*, we have* $x \cdot (y + z) = x \cdot y + x \cdot z$ *and* $(x + y) \cdot z = x \cdot z + y \cdot z$.

We say $(R, +, \cdot)$ *is a commutative ring if, in addition* $(R, +, \cdot)$ *satisfies*

(R4) *for all* $x, y \in R$*, we have* $x \cdot y = y \cdot x$.

We denote the identity element of $(R, +)$ *by* 0_R*, and the identity element of* $(R, \cdot)$ *by* 1_R*. A subring of* $(R, +, \cdot)$ *is a subset* $S \subset R$ *such that* $(S, +, \cdot)$ *is itself a ring and* $1_R \in S$.

Definition 1.2.2. *Let* R *be a ring.*

(i) *We define the* centre *of* R *to be*

$$\mathrm{Z}(R) := \{r \in R \mid \forall s \in R : rs = sr\}.$$

(ii) *We define the* characteristic *of* R *to be the integer* $n \in \mathbb{Z}_{\geq 0}$ *such that* $\ker(\mathbb{Z} \to R^+,\ 1 \mapsto 1_R) = n\mathbb{Z}$.

Note 1.2.3. Let R be a finite ring and let R^+ denote the underlying abelian group of R. Then

$$\mathrm{char}(R) = \exp(R^+),$$

where $\exp(R^+)$ is the exponent of the abelian group R^+, i.e. the smallest positive integer m such that for all $r \in R^+$, the composition of r with itself m times equals the identity element.

Definition 1.2.4. *Let* $(R, +, \cdot)$ *be a ring. A* left ideal *of* R *is subset* $I \subset R$ *such that*

(I1) $(I,+)$ *is an abelian subgroup of* $(R,+)$,
(I2) *for all* $r \in R$ *and* $i \in I$, *we have* $ri \in I$.

Analogously, we can define right ideals. *An ideal is said to be* two-sided, *if it is both right and left.*

Definition 1.2.5. *Let* R *be a ring and* $I \subseteq R$ *a one-sided (or two-sided) ideal of* R. *Then*

(i) I *is said to be* nil *if every element of* I *is nilpotent.*
(ii) I *is said to be* nilpotent *if there exists* $n \in \mathbb{Z}_{>0}$ *such that* $I^n = 0$.

Definition 1.2.6. *A ring* R *is said to be* simple *if* R *is nonzero and the only two-sided ideals of* R *are* 0 *and* R.

Definition 1.2.7. *Let* $(R, +_R, \cdot_R)$, $(S, +_S, \cdot_S)$ *be two rings. A* ring homomorphism $F : R \to S$ *is a homomorphism of the underlying abelian groups, such that* $F(1_R) = 1_S$ *and for all* $r_1, r_2 \in R$, *we have* $F(r_1 \cdot_R r_2) = F(r_1) \cdot_S F(r_2)$. *A bijective ring homomorphism is called a* ring isomorphism.

Definition 1.2.8. *Let* R *be a ring. We define the* prime subring *of* R *to be the image of the ring homomorphism* $\mathbb{Z} \to R$, *given by* $1 \mapsto 1_R$.

Definition 1.2.9. *An* algebra *is a pair of rings,* k *and* R, *with* k *commutative, together with a ring homomorphism* $\varphi : k \to R$ *such that* $\mathrm{im}(\varphi) \subseteq \mathrm{Z}(R)$. *We then say that* R *is an* algebra over k.

Theorem 1.2.10 ([27], Theorem 1.1)**.** *Let* R *be a finite-dimensional algebra over a field* $\mathbb{F}$ *and let* $n := \dim_{\mathbb{F}}(R)$. *Then* R *is isomorphic to a subalgebra of* $\mathcal{M}_n(\mathbb{F})$.

Theorem 1.2.11 ([57], Theorem 3.1)**.** *Let* R *be a ring,* $n \in \mathbb{Z}_{>0}$ *and* $S = \mathcal{M}_n(R)$. *Then*

(i) *If* I *is a two-sided ideal of* R, *then* $\mathcal{M}_n(I)$ *is a two-sided ideal of* S.
(ii) *Every two-sided ideal of* S *is of the form* $\mathcal{M}_n(I)$, *for some two-sided ideal* I *of* R.

1.3 Basic module theory

Definition 1.3.1. *Let* R *be a ring. A* left R-module *is an abelian group* $(M,+)$, *together with an action* $R \times M \to M$ *such that:*

(M1) *for all* $r, s \in R$ *and* $x \in M$, *we have* $r(sx) = (rs)x$,
(M2) *for all* $r, s \in R$ *and* $x \in M$, *we have* $(r+s)x = rx + sx$,
(M3) *for all* $r \in R$ *and* $x, y \in M$, *we have* $r(x+y) = rx + ry$,
(M4) *for all* $x \in M$, *we have* $1_R x = x$.

Analogously, we can define right R-modules. *A* submodule *of a left* R*-module* M *is an abelian subgroup* $N \subset M$ *such that* $RN \subseteq N$.

Note 1.3.2. By a *module*, we will always mean a left module.

Definition 1.3.3. *Let R, S be two rings. An R-S-bimodule is an abelian group $(M, +)$ such that*

(B1) *M is a left R-module,*
(B2) *M is a right S-module,*
(B3) *for all $x \in M$, $r \in R$ and $s \in S$, we have $(rm)s = r(ms)$.*

We often write ${}_RM_S$ for an R-S-bimodule M.

Definition 1.3.4. *Let R be a ring. Then the* left-regular *R-module, ${}_RR$, is the abelian group $(R, +)$, together with an action $R \times R \to R$ given by left-multiplication. We can similarly define the* right-regular *R-module, R_R.*

Definition 1.3.5. *Let R be a ring. We say an R-module M is* free *if $M \cong \bigoplus_{i \in I} R_i =: R^{(I)}$ as R-modules, where I is an arbitrary indexing set and $R_i \cong R$ for all $i \in I$.*

Definition 1.3.6. *Let R be a ring. We say that R has left IBN (Invariant Basis Number) if for all $n, m \in \mathbb{Z}_{>0}$, whenever ${}_RR^n \cong {}_RR^m$, we have that $n = m$.*

Note 1.3.7 ([56], Corollary 1.2)**.** Let R be a ring. If ${}_RR^{(I)} \cong {}_RR^{(J)}$, where R is nonzero and I is infinite, then $|I| = |J|$.

Definition 1.3.8. *Let R be a ring with left IBN and let $M \cong R^{(I)}$ be a free R-module, for some indexing set I. The* rank *of M over R, which we denote by $\mathrm{rk}_R(M)$ is the cardinality of I.*

Example 1.3.9 ([56], Example 1.6)**.** The following rings have left IBN: division rings, local rings, nonzero commutative rings, nonzero left-artinian rings.

Definition 1.3.10. *Let R be a ring and M, N two R-modules. A* module homomorphism *$f : M \to N$ is a homomorphism of the underlying abelian groups, such that for all $r \in R$, we have $f(rm) = rf(m)$.*

Definition 1.3.11. *Let R be a ring and M an R-module. Then*

(i) *M is* simple *if $M \neq 0$ and its only submodules are 0 and M.*
(ii) *M is* indecomposable *if $M \neq 0$ and M cannot be written as the direct sum of two nontrivial, proper submodules.*
(iii) *M is* semisimple *if for any submodule $N \leq M$, there exists $C \leq M$ such that $M = N \oplus C$.*
(iv) *M is* artinian *if every descending chain of submodules of M stabilizes.*
(v) *M is* noetherian *if every ascending chain of submodules of M stabilizes.*
(vi) *M is* finitely generated *over R if there exists a finite set $X \subset M$ such that $M = \sum_{x \in X} Rx$.*
(vii) *M has* finite length *if M has a finite composition series, i.e. there exists $t \in \mathbb{Z}_{\geq 0}$ and a sequence $(N_i)_{i=0}^t$ of submodules of M such that $M = N_t > N_{t-1} > \ldots > N_1 > N_0 = 0$ and for all $0 \leq i \leq t-1$, we have that N_{i+1}/N_i is simple.*

Proposition 1.3.12 ([57], Theorem 19.16)**.** (Fitting's Lemma) *Let R be a ring, M a finite-length R-module and $f \in \mathrm{End}_R(M)$. Then there exists $n \in \mathbb{Z}_{>0}$ such that*

$$M = \ker(f^n) \oplus \mathrm{im}(f^n).$$

Theorem 1.3.13 ([57], Corollary 19.22)**.** (Krull-Remak-Schmidt Theorem) *Let R be a ring and M an R-module of finite length. Then there exist $n \in \mathbb{Z}_{>0}$ and indecomposable submodules $M_i \leq M$ such that*

$$M = \bigoplus_{i=1}^{n} M_i.$$

Moreover, n is uniquely determined, and the sequence $(M_i)_{i=1}^n$ is uniquely determined up to isomorphism, and up to a permutation.

Proposition 1.3.14. *Let R be a ring and $I \subset R$ a two-sided ideal. Let M be an abelian group. Then M is an R/I-module if and only if M is an R-module that is annihilated by I.*

Proof. Suppose M is an R-module that is annihilated by I. Then we can define an R/I-module structure on M, given by $R/I \times M \to M$, $(r+I)m \mapsto rm$. Conversely, if M is an R/I-module, then M is an R-module via $R \times M \to M$, $rm \mapsto \overline{r}m$, where $^- : R \to R/I$. Clearly M is then annihilated by I. □

1.4 More ring theory

1.4.1 Menagerie of rings I

Definition 1.4.1. *Let R be a ring. Then*

(i) *R is a* division ring *if $R \neq 0$ and for all $0 \neq r \in R$, there exists $s \in R$ such that $rs = sr = 1_R$.*
(ii) *R is* Dedekind-finite *if every element of R that is left-invertible is also right-invertible.*
(iii) *R is* left-artinian (resp. right-artinian) *if ${}_RR$ (resp. R_R) is artinian.*
(iv) *R is* left-noetherian (resp. right-noetherian) *if ${}_RR$ (resp. R_R) is noetherian.*

Proposition 1.4.2 ([57], Theorem 3.3)**.** *Let D be a division ring and let $R = \mathcal{M}_n(D)$, for some $n \in \mathbb{Z}_{>0}$. Then, up to isomorphism, R has a unique simple left module V, and $V \cong D^n$ as R-modules.*

1.4.2 Semisimple rings

One of the most important class of rings is that of *semisimple rings*.

Theorem 1.4.3 ([57], Theorems 2.5, 2.8, Corollary 3.7)**.** *Let R be a ring. Then the following are equivalent:*

(i) *The left-regular module, ${}_RR$, is semisimple.*
(ii) *All left R-modules are semisimple.*
(iii) *All left R-modules are projective.*
(iv) *All left R-modules are injective.*

Replacing "left" with "right" gives further equivalent conditions.

Definition 1.4.4. *Let R be a ring. If R satisfies any of the conditions of Theorem 1.4.3, then R is said to be a* semisimple ring.

Theorem 1.4.5 ([57], Theorem 3.5)**.** (Wedderburn's Theorem) *Let R be a ring. Then R is semisimple if and only if*

$$R \cong \prod_{i=1}^{t} \mathcal{M}_{n_i}(D_i),$$

where $t \in \mathbb{Z}_{\geq 0}$, $n_i \in \mathbb{Z}_{>0}$ and the D_i are division rings.

Note 1.4.6. Let R be a semisimple ring. Then the isomorphism classes of simple R-modules form a finite set. Moreover, the proof of Theorem 1.4.5 shows that

$$R \cong \prod_{S \text{ simple}} \mathrm{End}_{\mathrm{End}_R(S)}(S),$$

where the product ranges over the isomorphism classes of simple R-modules.

1.4.3 The Jacobson radical

The notion of semisimplicity is inextricably linked to that of the Jacobson radical.

Definition 1.4.7. *Let R be a ring. The* Jacobson radical *is defined as*

$$\mathrm{J}(R) := \bigcap_{\substack{I \subset R \\ I \text{ max left ideal}}} I.$$

Theorem 1.4.8 ([57], Corollary 4.2)**.** *Let R be a ring. Then*

$$\mathrm{J}(R) = \bigcap_{\substack{M \\ M \text{ simple } R\text{-module}}} \mathrm{ann}_R(M)$$

Theorem 1.4.9 ([57], Lemma 4.11, Theorems 4.12,4.14)**.** *Let R be a ring and $\mathrm{J}(R)$ its Jacobson radical. Then*

(i) *$\mathrm{J}(R)$ is a two-sided ideal of R.*
(ii) *If $I \subset R$ is a nil one-sided ideal, then $I \subseteq \mathrm{J}(R)$.*
(iii) *If R is left-artinian, then $\mathrm{J}(R)$ is the largest nilpotent left (resp. right) ideal of R.*
(iv) *R is semisimple if and only if R is left-artinian and $\mathrm{J}(R) = 0$.*

Theorem 1.4.10 ([18], Section 2)**.** *Let R be a finite-dimensional algebra of matrices over a field $\mathbb{F}$, where* $\mathrm{char}(\mathbb{F}) = 0$*. Then*

$$\mathrm{J}(R) = \{r \in R \mid \mathrm{Tr}(rs) = 0 \text{ for all } s \in R\}. \tag{1.1}$$

Proposition 1.4.11 ([57], Exercise 4.12B)**.** *For any collection of rings $\{A_i\}_{i\in I}$ we have $\mathrm{J}(\prod_i A_i) = \prod_i \mathrm{J}(A_i)$.*

Proposition 1.4.12 ([57], Example 21.14)**.** *Let R be a ring and $n \in \mathbb{Z}_{>0}$. Then $\mathrm{J}(\mathcal{M}_n(R)) = \mathcal{M}_n(\mathrm{J}(R))$.*

Proposition 1.4.13. *Let R be a ring, $I \subseteq R$ a two-sided nilpotent ideal and M an R-module. Then M is an R/I-modules, and M is simple over R/I if and only if it is simple over R.*

Proof. This is an easy corollary of Proposition 1.3.14. □

1.4.4 Menagerie of rings II

Definition 1.4.14. *Let R be a ring. Then*

(i) *R is* semilocal *if $R/\mathrm{J}(R)$ is semisimple.*
(ii) *R is* semiprimary *if $\mathrm{J}(R)$ is nilpotent and $R/\mathrm{J}(R)$ is semisimple.*
(iii) *R is* local *if $R/\mathrm{J}(R)$ is a division ring.*

Theorem 1.4.15 ([57], Theorem 19.1)**.** *Let R be a ring. Then R is local if and only if R has a unique maximal left (equiv. right) ideal.*

1.5 Idempotents

Definition 1.5.1. *Let R be a ring. An element $e \in R$ is an* idempotent *if $e^2 = e$. Two idempotents e_1 and e_2 are said to be* orthogonal *if $e_1e_2 = e_2e_1 = 0$.*

Definition 1.5.2. *Let R be a ring and $e \in R$ an idempotent. Then*

(i) *e is* central *if $e \in \mathrm{Z}(R)$.*
(ii) *e is* primitive *if $e \neq 0$ and it cannot be written as the sum of two nonzero orthogonal idempotents.*
(iii) *e is* centrally primitive *if $e \in \mathrm{Z}(R)$, $e \neq 0$ and e cannot be written as the sum of two nonzero orthogonal central idempotents.*

Definition 1.5.3. *A ring R is said to be* connected *if $R \neq 0$ and the only central idempotents in R are 0 and 1.*

Theorem 1.5.4. *Let R be a ring, and M an R-module.*

(i) *Let N, P be R-modules. Then $M = N \oplus P$ if and only if there exists an idempotent $e \in \mathrm{End}_R(M)$ such that $N = e(M)$ and $P = (1-e)(M)$.*

(ii) *Let A, B be R-modules. Then $R = A \oplus B$ if and only if there exists an idempotent $e \in R$ such that $A = Re$ and $B = R(1-e)$.*

(iii) ([77], Proposition 1.1.14) *Let R_1, R_2 be two-sided ideals of R. Then $R = R_1 \times R_2$ if and only if there exist central orthogonal idempotents e_1, e_2 such that $e_1 + e_2 = 1$, with $R_i = Re_i$, for $i = 1, 2$.*

Let R be a ring and suppose that $1 \in R$ can be written as a finite sum of orthogonal centrally primitive idempotents. Then such a decomposition $1 = e_1 + \ldots + e_n$ is unique up to permutation of the summands, and R can be written as a finite product of connected rings. Moreover, we have

$$R = Re_1 \oplus \ldots \oplus Re_n.$$

We call this a *block decomposition* of R.

Theorem 1.5.5 ([57], Proposition 22.2)**.** *Let R be a left-noetherian ring. Then R has a block decomposition.*

Proposition 1.5.6. *Let R be a ring. If R has a block decomposition $R = Re_1 + \ldots + Re_n$, where $\{e_i\}_{i=1}^n$ is a set of orthogonal centrally primitive idempotents of sum 1, then $\mathrm{Z}(R)$ has block decomposition $\mathrm{Z}(R) = \mathrm{Z}(R)e_1 + \ldots + \mathrm{Z}(R)e_n$.*

Theorem 1.5.7 ([57], Corollary 19.19)**.** *A nonzero left-artinian ring R is local if and only if R has no nontrivial idempotents.*

Proposition 1.5.8. *Let R be a left-artinian ring with Jacobson radical $\mathrm{J}(R)$. Then the natural projection $p : R \to R/\mathrm{J}(R)$ induces a surjective map on the set of idempotents.*

Proof. Let $E \in R$ be an idempotent. Then certainly $p(E)$ is an idempotent in $R/\mathrm{J}(R)$. Suppose $e \in R/\mathrm{J}(R)$ is an idempotent, i.e. $e^2 - e \in \mathrm{J}(R)$. What we want to find is an element satisfying $x^2 - x = 0$ in R, which is mapped to e. Consider the polynomial $F(x) = 3x^2 - 2x^3$. Let $e_1 := F(e)$. Then

$$e_1^2 - e_1 = (3e^2 - 2e^3)^2 - (3e^2 - 2e^3) = (4e^2 - 4e - 3)(e^2 - e)^2 \in \mathrm{J}(R)^2,$$

so $e_1^2 - e_1 \in \mathrm{J}(R)^2$. Moreover, $e_1 = e - (2e-1)(e^2 - e)$, so $e_1 \equiv e \mod \mathrm{J}(R)$.

We define $e_i := F(e_{i-1})$. By induction, we have $e_i^2 - e_i \in \mathrm{J}(R)^{2^i}$ and $e_i \equiv e \mod \mathrm{J}(R)$. Since R is left-artinian, $\mathrm{J}(R)$ is nilpotent, so there exists $n \in \mathbb{Z}_{\geq 0}$ such that $e_n^2 - e_n \in \mathrm{J}(R)^n = 0$. Then $E = e_n$ is the element we were after. □

Remark 1.5.9. The key to the above proof is that $e^2 - e$ is nilpotent. Hence we can use the same lifting technique against any nil ideal of R.

1.6 More module theory

1.6.1 Schur's Lemma, Converse Schur Lemma

Proposition 1.6.1 ([57], Lemma 3.6)**.** (Schur's Lemma) *Let R be a ring and M a simple module. Then $\mathrm{End}_R(M)$ is a division ring.*

Note 1.6.2. The converse is not necessarily true. To see this, let F be a field and consider the ring

$$R = \begin{pmatrix} F & F \\ 0 & F \end{pmatrix}$$

and the R-module

$$M = R \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 0 & F \\ 0 & F \end{pmatrix}.$$

Then $\mathrm{End}_R(M) \cong F$, but M is not simple.

Definition 1.6.3. *Let R be a ring. We say ${}_R\mathfrak{M}$, the category of R-modules, satisfies the* converse of Schur's Lemma *if every R-module whose endomorphism ring is a division ring, is in fact simple.*

Theorem 1.6.4 ([71], Theorem 1.6)**.** (Converse Schur) *Let R be a semiprimary ring. Then the category of R-modules, ${}_R\mathfrak{M}$, satisfies the converse of Schur's Lemma if and only if R is a finite direct product of full matrix rings over local rings.*

1.6.2 Nakayama's Lemma

Theorem 1.6.5 ([57], Lemma 4.22)**.** (Nakayama's Lemma) *Let R be a ring and $J \subseteq R$ a left ideal of R. Then the following are equivalent:*

(i) $J \subseteq \mathrm{J}(R)$.
(ii) *For any finitely generated left R-module M,*

$$J \cdot M = M \Rightarrow M = 0.$$

(iii) *For any left R-modules $N \leq M$ such that M/N is finitely generated,*

$$N + J \cdot M = M \Rightarrow N = M.$$

1.6.3 Projective and injective modules

Definition 1.6.6. *Let R be a ring and P an R-module. Then P is said to be* projective *if for any surjective R-module homomorphism $g : B \twoheadrightarrow C$ and any R-module homomorphism $f : P \to C$, there exists an R-module homomorphism $h : P \to B$ such that $f = gh$:*

$$\begin{array}{ccccc} & & & & P \\ & & & \swarrow^{h} & \downarrow f \\ B & & \xrightarrow{g} & & C \longrightarrow 0. \end{array}$$

Theorem 1.6.7 ([56], §2A)**.** *Let R be a ring and P an R-module. Then the following are equivalent:*

(i) *P is projective.*
(ii) *P is a direct summand of a free R-module.*

(iii) *Every surjective R-module homomorphism $M \twoheadrightarrow P$ splits.*
(iv) *The functor $\mathrm{Hom}_R(P, -)$ is exact on ${}_R\mathfrak{M}$.*

Finitely generated projective modules over $\mathbb{Z}$ and $\mathbb{Z}/n\mathbb{Z}$, for $n \in \mathbb{Z}_{>0}$, are easy to describe.

Proposition 1.6.8. (i) *A $\mathbb{Z}$-module is finitely generated projective if and only if it is free of finite rank.*
(ii) *Let p be a prime and let $e \in \mathbb{Z}_{>0}$. A $\mathbb{Z}/p^e\mathbb{Z}$-module is finitely generated projective if and only if it is free of finite rank.*
(iii) *Let $n \in \mathbb{Z}_{>0}$. A $\mathbb{Z}/n\mathbb{Z}$-module is finitely generated projective if and only if it is a direct sum of copies of modules of the form $\mathbb{Z}/m\mathbb{Z}$, with $m \mid n$ such that $\gcd(\frac{n}{m}, m) = 1$.*

Proof. Part (i) is a consequence of $\mathbb{Z}$ being a principal ideal domain. Part (ii) holds since $\mathbb{Z}/p^e\mathbb{Z}$ is a local ring.

For part (iii), note that $\mathbb{Z}/m\mathbb{Z}$ is a $\mathbb{Z}/n\mathbb{Z}$-module if and only if $m \mid n$. It is now enough to show that if $m \mid n$, then $\mathbb{Z}/m\mathbb{Z}$ is $\mathbb{Z}/n\mathbb{Z}$-projective if and only if $\gcd(\frac{n}{m}, m) = 1$. Suppose $n = \prod_{i \in I} p_i^{a_i}$, where I is a finite indexing set and all p_i are distinct primes. Then $\gcd(\frac{n}{m}, m) = 1$ if and only if $m = \prod_{j \in J} p_j^{a_j}$, for some subset $J \subseteq I$. But this happens if and only if $\mathbb{Z}/m\mathbb{Z} = \bigoplus_{j \in J} \mathbb{Z}/p_j^{a_j}\mathbb{Z}$, which is a direct summand of $\mathbb{Z}/n\mathbb{Z}$. □

Proposition 1.6.9 ([20], Proposition 1.4)**.** *Let k be a commutative ring and let R be a k-algebra such that R is projective as a k-module. Let M be a projective R-module. Then M is projective over k.*

Definition 1.6.10. *Let R be a ring and I an R-module. Then I is said to be* injective *if for any injective R-module homomorphism $g : A \hookrightarrow B$ and any R-module homomorphism $f : A \to I$, there exists an R-module homomorphism $h : B \to I$ such that $f = hg$:*

$$\begin{array}{ccccc} & & I & & \\ & & {\scriptstyle f}\uparrow & \nwarrow^{h} & \\ 0 & \longrightarrow & A & \xrightarrow[g]{} & B. \end{array}$$

Definition 1.6.11. *Let R be a ring. If R is injective as a left-regular (resp. right-regular) module, we say that R is* left (resp. right) self-injective.

Theorem 1.6.12 ([56], §3A; [76], Proposition 3.42)**.** *Let R be a ring and I an R-module. Then the following are equivalent:*

(i) *I is injective.*
(ii) *Every injective R-module homomorphism $I \hookrightarrow M$ splits.*
(iii) (Baer's Test) *For all left ideals $K \subset R$, any R-homomorphism $K \to I$ can be extended to a map $R \to I$.*
(iv) *Every short exact sequence $0 \to I \to M \to N \to 0$, where M is an R-module and N is a cyclic R-module, splits.*
(v) *The functor $\mathrm{Hom}_R(-, I)$ is exact on ${}_R\mathfrak{M}$.*

1.6.4 Flat and finitely presented modules

Definition 1.6.13. *Let R be a ring and M an R-module. We say M is* flat *over R if the functor $- \otimes_R M$ is exact.*

Proposition 1.6.14 ([56], Proposition 4.3; [57], Theorem 23.20)**.** *Over a left-artinian ring, the notions of projective modules and flat modules coincide.*

Definition 1.6.15. *Let R be a ring and M an R-module. We say M is* finitely presented *over R if there is an exact sequence $R^m \to R^n \to M \to 0$, for some $m, n \in \mathbb{Z}_{\geq 0}$.*

Proposition 1.6.16 ([56], Proposition 4.29)**.** *A ring R is left-noetherian if and only if every finitely generated R-module is finitely presented.*

1.6.5 Rank of a projective module

In this section, suppose R is a commutative ring. Denote by $\mathrm{Spec}(R)$ the set of prime ideals of R and by $\mathrm{Max}(R)$ the set of maximal ideals of R. Let M be an R-module and $\mathfrak{p} \in \mathrm{Spec}(R)$. Then we denote by $M_{\mathfrak{p}}$ the localisation of M at $R \backslash \mathfrak{p}$.

Proposition 1.6.17 ([58], Corollary 3.4)**.** *Let M be a finitely presented R-module. Then the following are equivalent:*

(i) *M is projective over R,*
(ii) *for all $\mathfrak{m} \in \mathrm{Max}(R)$, we have that $M_{\mathfrak{m}}$ is projective over $R_{\mathfrak{m}}$,*
(iii) *for all $\mathfrak{p} \in \mathrm{Spec}(R)$, we have that $M_{\mathfrak{p}}$ is free over $R_{\mathfrak{p}}$.*

Let P be a projective R-module. Consider the function

$$\mathrm{rk}_R(P) : \mathrm{Spec}(R) \to \mathbb{Z}, \quad \mathfrak{p} \mapsto \mathrm{rk}_{R_{\mathfrak{p}}}(P_{\mathfrak{p}}).$$

Definition 1.6.18. *Let P be a projective R-module. If $\mathrm{rk}_R(P)$ is a constant function, then we say P has* constant rank*.*

Proposition 1.6.19 ([58], Corollary 3.6)**.** *If R is connected, then every projective R-module has constant rank.*

1.6.6 Hom & $\otimes$

Let R, S, T be rings, let M be an R-S-bimodule, N an R-T-bimodule and P an S-T-bimodule. Then

(i) $\mathrm{Hom}_R({}_RM_S, {}_RN_T)$ is an S-T-bimodule, where for all $s \in S$, $t \in T$, $m \in M$ and $f \in \mathrm{Hom}_R(M, N)$, we have $s \cdot f(m) = f(ms)$ and $(f \cdot t)(m) = f(m)t$.
(ii) $\mathrm{Hom}_T({}_RN_T, {}_SP_T)$ is an S-R-bimodule, where for all $s \in S$, $r \in R$, $n \in N$ and $g \in \mathrm{Hom}_T(N, P)$, we have that $s \cdot g(n) = sg(n)$ and $(g \cdot r)(n) = g(rn)$.
(iii) ${}_RM_S \otimes_S {}_SP_T$ is an R-T-bimodule, where for all $r \in R$, $t \in T$, $m \in M$ and $n \in N$, we have $r \cdot (m \otimes n) = rm \otimes n$ and $(m \otimes n) \cdot t = m \otimes nt$.

Proposition 1.6.20. *Let R, S be two rings, let $\alpha : R \to S$ be a ring homomorphism and M an S-R-bimodule. Then*

$$\mathrm{Hom}_S({}_SS_R, {}_SM_R) \cong {}_RM_R,$$

as R-R-bimodules.

Proposition 1.6.21 ([79], Proposition 18.44)**.** *Let R, S, T be rings, let M be an R-S-bimodule, N an S-T-bimodule and P an R-module. Then*

$$\mathrm{Hom}_R(M \otimes_S N, P) \cong \mathrm{Hom}_S(N, \mathrm{Hom}_R(M, P)),$$

as T-modules.

Proposition 1.6.22 ([58], Chapter I, Example 2.2(4), Proposition 2.13)**.** *Let R, R' be commutative rings, $\alpha : R \to R'$ a ring homomorphism and P, Q two finitely generated projective R-modules. Then*

$$\mathrm{Hom}_R(P, Q) \otimes_R R' \cong \mathrm{Hom}_{R'}(P \otimes_R R', Q \otimes_R R'),$$
$$(P \otimes_R Q) \otimes_R R' \cong (P \otimes_R R') \otimes_{R'} (Q \otimes_R R'),$$

as R'-modules.

1.6.7 Projective covers and injective hulls

Definition 1.6.23. *Let M be an R-module. A* superfluous submodule *of M is an R-module $S \subseteq M$ such that*

$$\forall N \leq M : (S + N = M \Rightarrow N = M).$$

If S is a superfluous submodule of M, we write $S \subseteq_s M$.

Definition 1.6.24. *Let M be an R-module. A* projective cover *of M is a pair (P, ϕ), where P is a projective R-module, $\phi : P \twoheadrightarrow M$ is an epimorphism, and $\ker(\phi) \subseteq_s P$.*

Theorem 1.6.25 ([57], Proposition 24.10, Example 24.11(3), Theorem 24.18)**.** *Let R be a ring.*

(i) *If R is left-artinian, then any R-module has a projective cover.*
(ii) *Let M be an R-module. Suppose (P, ϕ) and (P', ϕ') are two projective covers of M. Then there exists an isomorphism $\alpha : P' \to P$ such that $\phi' = \phi\alpha$.*
(iii) *Let $M_1, \ldots, M_n$ be R-modules. Suppose (P_i, ϕ_i) is a projective cover of M_i, for all $1 \leq i \leq n$. Then $(\bigoplus_{i=1}^n P_i, \bigoplus_{i=1}^n \phi_i)$ is a projective cover of $\bigoplus_{i=1}^n M_i$.*

Definition 1.6.26. *Let M be an R-module. An* essential extension *of M is an R-module $E \supseteq M$ such that*

$$\forall F \leq E : (F \cap M = 0 \Rightarrow F = 0)$$

If E is an essential extension of M, we write $M \subseteq_e E$.

Theorem 1.6.27 ([57], Theorem 3.30)**.** *Let R be a ring and $M \subseteq I$ two R-modules. Then the following are equivalent:*

(i) *I is maximal essential over M, i.e. $I \supseteq_e M$ and no module properly containing I can be an essential extension of M.*
(ii) *I is injective, and is essential over M.*
(iii) *I is minimal injective over M, i.e. I is injective and if I' is an injective module such that $M \subseteq I' \subseteq I$, then $I = I'$.*

Definition 1.6.28. *Let M be an R-module. An* injective hull *of M is an R-module $I \supseteq M$ satisfying one of the conditions of Theorem 1.6.27.*

Theorem 1.6.29 ([57], Lemma 3.29, Corollary 3.32, Example 3.38)**.** *Let R be a ring.*

(i) *Every R-module has an injective hull.*
(ii) *Let M be an R-module. Suppose I and I' are two injective hulls of M. Then there exists an isomorphism $I \to I'$ which is the identity on M.*
(iii) *Let $M_1, \ldots, M_n$ be R-modules. Suppose I_j is an injective hull of M_j, for all $1 \leq j \leq n$. Then $\bigoplus_{j=1}^n I_j$ is an injective hull of $\bigoplus_{j=1}^n M_j$.*

Theorem 1.6.30 ([56], Lemma 3.28, Theorem 3.30)**.** *Let R be a ring and M an R-module. Let I be an injective hull of M. Then M is injective if and only if $M = I$.*

1.7 Quasi-Frobenius rings

Theorem 1.7.1 ([56], Theorems 15.1, 15.9, Remark 15.10)**.** *Let R be a ring. Then the following are equivalent:*

(i) *R is left-noetherian and left self-injective.*
(ii) *R is right-noetherian and left self-injective.*
(iii) *R is left-noetherian and right self-injective.*
(iv) *R is right-noetherian and right self-injective.*
(v) *all projective R-modules are injective.*
(vi) *all injective R-module are projective.*

Definition 1.7.2. *Let R be a ring. If R satisfies any of the conditions of Theorem 1.7.1, then R is said to be a* quasi-Frobenius ring.

Example 1.7.3. The following rings are quasi-Frobenius:

(i) fields,
(ii) $\mathbb{Z}/n\mathbb{Z}$, for $n \in \mathbb{Z}_{>0}$,
(iii) semisimple rings,
(iv) $\mathcal{M}_n(R)$, for R a quasi-Frobenius ring and $n \in \mathbb{Z}_{\geq 0}$,
(v) the group ring $R[G]$, for R a quasi-Frobenius ring and G a finite group,
(vi) Galois rings (see Note 6.2.59).

1.8 Frobenius algebras and symmetric algebras

Let k be a commutative ring and A a k-algebra that is finitely generated projective as a module over k. The k-dual, $\mathrm{Hom}_k(A,k)$, is an A-A-bimodule. The left module structure is given by

$$a \cdot f = (x \mapsto f(xa)),$$

and the right module structure is given by

$$f \cdot a = (x \mapsto f(ax)),$$

where $a \in A$ and $f \in \mathrm{Hom}_k(A,k)$. These two actions are compatible: for any $a, a', x \in A$, we have $((a \cdot f) \cdot a')(x) = f(a'xa) = (a \cdot (f \cdot a'))(x)$.

Comparing the A-A-bimodule structures of A and $\mathrm{Hom}_k(A,k)$ leads to the following two notions.

Definition 1.8.1. *Let k be a commutative ring and A a k-algebra that is finitely generated projective as a module over k. If $A \cong \mathrm{Hom}_k(A,k)$ as left A-modules, then we say A is a* Frobenius algebra. *If $A \cong \mathrm{Hom}_k(A,k)$ as A-A-bimodules, then we say A is a* symmetric algebra.

Theorem 1.8.2 ([56], Theorems 16.54)**.** *Let k be a commutative ring and A a k-algebra that is finitely generated projective as a module over k. Then A is a symmetric algebra over k if and only if there exists a k-bilinear map $B : A \times A \to k$ such that*

(i) *B is symmetric, i.e. for all $x, y \in A$, we have $B(x,y) = B(y,x)$,*
(ii) *B is nonsingular, i.e. the map $A \to \mathrm{Hom}_k(A,k)$, given by $x \mapsto (y \mapsto B(x,y))$ is a k-module isomorphism,*
(iii) *B is associative, i.e. for all $x, y, z \in A$, we have $B(xy,z) = B(x,yz)$,*

Example 1.8.3 ([56], 16.56-59)**.** (Symmetric algebras)

1. Let k be a field and G a finite group. Then the group ring $A = k[G]$ is a symmetric k-algebra. To see this, consider the map $B : A \times A \to k$ given by $B(\sum_{g\in G} a_g g, \sum_{h\in G} b_h h) = \sum_{g\in G} a_g b_{g^{-1}}$, where for all $g \in G$, we have $a_g, b_g \in k$.
2. Let k be a field and $A = \mathcal{M}_n(k)$, for some $n \in \mathbb{Z}_{>0}$. Then A is a symmetric k-algebra. To see this, consider the map $B : A \times A \to k$, given by $B(X,Y) = \mathrm{tr}(XY)$, where tr denotes the usual trace map.
3. Let k be a field. Then any finite-dimensional semisimple k-algebra is symmetric.

1.8.1 Generators and progenerators

Definition 1.8.4. *Let R be a ring and M and R-module. The* trace ideal *of M over R is defined to be*

$$\mathfrak{T}_R(M) := \sum_{f \in \mathrm{Hom}_R(M,R)} \mathrm{im}(f).$$

Note 1.8.5. It is easy to check that $\mathfrak{T}_R(M)$ is a two-sided ideal of R.

Definition 1.8.6. *Let R be a ring. An R-module M is an* R-generator *if $\mathfrak{T}_R(M) = R$. If, in addition, M is finitely generated and projective, then it is said to be an* R-progenerator.

Note 1.8.7. Over a commutative ring R, any faithful finitely generated projective module is a progenerator. The converse also holds.

1.9 Duality

Let R be a finite ring. Denote by ${}^{\mathrm{fg}}_{R}\mathfrak{M}$ and $\mathfrak{M}^{\mathrm{fg}}_{R}$ the categories of finitely generated left, respectively right, R-modules.

Definition 1.9.1. *Let R be a finite ring and denote by ${}^{\mathrm{fg}}_{R}\mathfrak{M}$ and $\mathfrak{M}^{\mathrm{fg}}_{R}$ the categories of finitely generated left and right R-modules, respectively. We define the* character functors

$$\widehat{\ }\colon {}^{\mathrm{fg}}_{R}\mathfrak{M} \rightleftharpoons \mathfrak{M}^{\mathrm{fg}}_{R}, \quad M \mapsto \widehat{M} := \mathrm{Hom}_{\mathbb{Z}}(M, \mathbb{Q}/\mathbb{Z}).$$

The module $\widehat{M}$ is called the character module *of M.*

Theorem 1.9.2 ([56], §19C,D)**.** *Let R be a finite ring. Consider the contravariant functors*

$$F : {}^{\mathrm{fg}}_{R}\mathfrak{M} \longrightarrow \mathfrak{M}^{\mathrm{fg}}_{R} \quad \textit{and} \quad G : \mathfrak{M}^{\mathrm{fg}}_{R} \longrightarrow {}^{\mathrm{fg}}_{R}\mathfrak{M}, \tag{1.2}$$

defined by taking character modules. Then $G \circ F$ and $F \circ G$ are naturally equivalent to the identity functors, i.e. F and G define a duality between ${}^{\mathrm{fg}}_{R}\mathfrak{M}$ and $\mathfrak{M}^{\mathrm{fg}}_{R}$.

Chapter 2

Linear algebra over $\mathbb{Z}$: basic algorithms for finite abelian groups

When working algorithmically with finite-dimensional algebras over a field, we rely on the vector space structure for our computations (most importantly for solving systems of linear equations). However, when presented with an arbitrary finite ring, we would like to be able to handle the situation regardless of whether it contains a field or not. In the absence of an underlying field, it is the additive group structure of the ring in question that we wish to exploit.

This chapter lays the foundation of everything that succeeds it. At the end of it, we will have built a toolbox for working with finite abelian groups within algorithms. This will allow our later algorithms to have a natural proof-like flow. We will not have to think about the bit operations that go on behind the scenes, and we will talk of algebraic structures, rather than of the strings of integers representing them.

Our algorithms are purposely conceptual. In this way, we aim to concentrate on the structural properties of our objects, rather than rely on seemingly random matrix manipulations that end up giving the "right" result.

We will represent finitely generated abelian groups via generators and relations. Correspondingly, we show how to represent group homomorphisms, subgroups and quotients of groups. Building on this, we describe deterministic polynomial-time algorithms that accomplish the following tasks in the abelian case:

1. test if a group homomorphism is injective,
2. test if a group homomorphism is surjective,
3. decide if two group homomorphisms are equal,
4. compute subgroups generated by a given finite set of elements in a group,

5. compute the quotient of a group by a subgroup,
6. compute kernels, images and cokernels of group homomorphisms,
7. compute direct sums of groups,
8. compute homomorphism groups and tensor products,
9. split exact sequences,
10. compute the order of a finite group,
11. compute the torsion subgroup of a finitely generated group,
12. compute the order of a given group element,
13. compute the exponent of a finite group,
14. write a finitely generated group as a direct sum of cyclic groups.

The last of these is particularly important, as it will allow us to assume in later chapters that a finite abelian group is given by specifying the sizes of its cyclic direct summands.

Working with finitely generated abelian groups in the representation we have chosen ultimately reduces to carrying out integer matrix computations. The way to keep the entries of these matrices under control is either to employ modular techniques, or to give the group a lattice structure and use basis reduction algorithms.

2.1 Lattices

The main references for this section are [67, 68].

Definition 2.1.1. *A* lattice *is an additive subgroup $L \subseteq \mathbb{R}^n$, where $n \in \mathbb{Z}_{\geq 0}$, for which there exists $\epsilon \in \mathbb{R}_{>0}$ such that for all $x \in L$, $x \neq 0$, we have $\langle x, x\rangle \geq \epsilon$, where $\langle \cdot, \cdot \rangle$ denotes the standard inner product on $\mathbb{R}^n$. A* sublattice *of L is a subgroup of L.*

Proposition 2.1.2 ([68], Section 2)**.** *A subset $L \subset \mathbb{R}^n$ is a lattice if and only if there exists a set $B \subset \mathbb{R}^n$ of $\mathbb{R}$-linearly independent vectors such that*

$$L = \sum_{b \in B} \mathbb{Z} b.$$

A set B as in Proposition 2.1.2 is said to be a *basis* of L, and the cardinality of B is the *rank* of L. Suppose $B = \{b_1, \ldots, b_m\}$, for some $m \in \mathbb{Z}_{>0}$ and let A be the matrix whose i^{th} column is given by b_i. The *determinant* of L is

$$\det(L) := \det(\langle b_i, b_j \rangle)_{1 \leq i,j \leq m}^{1/2} = |\det(A)|.$$

It can be shown that the rank and determinant of a lattice are well-defined.

Definition 2.1.3. *Two lattices L and L' are said to be* isomorphic *if there exists a bijective $\mathbb{Z}$-linear transformation $\tau : L \to L'$ such that for all $x, y \in L$, we have $\langle x, y\rangle = \langle \tau(x), \tau(y)\rangle$. If such a transformation exists, we write $L \cong L'$.*

Since most real numbers cannot be represented inside algorithms using a finite number of bits, we will only consider lattices whose vectors are rational numbers. In this case, we represent a lattice by giving a matrix $A \in \mathcal{M}_{n\times m}(\mathbb{Q})$ of rank m. Then L is taken to be the lattice with basis given by the m columns of A, and we write $L = \mathcal{L}(A)$.

An important notion in the theory of lattices is that of a *reduced basis*. A precise definition can be found in [68], Section 10. Intuitively, reduced bases can be thought of as consisting of "short" vectors that are "nearly orthogonal". To the notion of a reduced basis we associate a parameter $c > 4/3$. Roughly speaking, c is a qualitative measure of the reduction – the smaller the value of c, the better the reduction. When no such parameter is specified, it is typically taken to be 2.

An algorithm that, given a lattice, produces a reduced basis thereof is called a *lattice basis reduction algorithm*. An example of such an algorithm, that is deterministic and runs in polynomial time, is the LLL algorithm ([63]).

Definition 2.1.4. *Let L be a lattice of rank n in $\mathbb{R}^n$. The* dual lattice *of L is given by*

$$L^\star = \{x \in \mathbb{R}^n \mid \langle x, L\rangle \subset \mathbb{Z}\},$$

where $\langle\cdot,\cdot\rangle$ is the standard inner product.

Note 2.1.5.

(i) The dual lattice is a lattice.
(ii) $\mathrm{rank}(L^\star) = \mathrm{rank}(L)$ and $\det(L^\star) = \det(L)^{-1}$.
(iii) $L^{\star\star} = L$.
(iv) If L has basis given by the columns of a matrix A, then $L^\star$ has basis given by the columns of the inverse of the transpose of A.

2.1.1 Kernels, images and systems of linear equations over $\mathbb{Z}$

One of the basic tools that we will use is the efficient computability of kernels and images.

Theorem 2.1.6 ([68], Section 14)**.** *There exists a deterministic polynomial-time algorithm that, given a triple (m, n, f), with $n, m \in \mathbb{Z}_{\geq 0}$ and $f \in \mathcal{M}_{n\times m}(\mathbb{Z})$ a matrix representing a group homomorphism $f : \mathbb{Z}^m \to \mathbb{Z}^n$, computes $k := \mathrm{rank}(f)$ and a basis $b_1, \ldots, b_m$ for $\mathbb{Z}^m$ such that $b_1, \ldots, b_{m-k}$ is a basis for* $\ker f$ *and $f(b_{m-k+1}), \ldots, f(b_m)$ is a basis for* $\mathrm{im}\, f$.

This algorithm can then be used to solve systems of linear equations over $\mathbb{Z}$.

Theorem 2.1.7 ([68], Section 14)**.** *There exists a deterministic polynomial-time algorithm that, given a triple (m, n, f), with $n, m \in \mathbb{Z}_{\geq 0}$ and $f \in \mathcal{M}_{n\times m}(\mathbb{Z})$, together with a vector $b \in \mathbb{Z}^n$, computes the set of solutions of the equation $fx = b$, or determines that there is no solution.*

2.1.2 Intersection, sum, inclusion and equality of lattices

A subgroup $H \subseteq \mathbb{Z}^n$ is given to an algorithm by specifying a sequence of elements of $\mathbb{Z}^n$ that is a basis of H over $\mathbb{Z}$. Note that by Theorem 2.1.6, we can recover a basis of H from any generating set.

Proposition 2.1.8. *There exists a deterministic polynomial-time algorithm that, given $n \in \mathbb{Z}_{>0}$ and two subgroups $H_1, H_2 \subseteq \mathbb{Z}^n$, computes $H_1 \cap H_2$ and $H_1 + H_2$, together with the inclusion maps $H_1 \cap H_2 \to H_i$ and $H_i \to H_1 + H_2$, for $i = 1, 2$.*

Proof. Consider the group $H_1 \oplus H_2$, i.e. the group with elements of the form (h_1, h_2), where $h_1 \in H_1$ and $h_2 \in H_2$, together with componentwise addition. Let $\phi : H_1 \oplus H_2 \to \mathbb{Z}^n$ be the map given by $(h_1, h_2) \mapsto h_1 - h_2$. Then $\ker(\phi) = H_1 \cap H_2$ and $\mathrm{im}(\phi) = H_1 + H_2$, and both can be efficiently computed by Theorem 2.1.6. This produces bases of $H_1 \cap H_2$ and $H_1 + H_2$ in terms of the standard basis of $\mathbb{Z}^n$.

Now $H_1 \cap H_2$ is equal to the image of the projection $\ker(\phi) \to H_1$. This gives a basis for $H_1 \cap H_2$ in terms of the basis of H_1. Similarly for H_2. Further, $H_1 = H_1 \cap (H_1 + H_2)$, which gives a basis for H_1 in terms of the basis of $H_1 + H_2$.

□

As a consequence of this, we are able to determine inclusion and equality of two subgroups of $\mathbb{Z}^n$.

Corollary 2.1.9. *There exists a deterministic polynomial-time algorithm such that, given $n \in \mathbb{Z}_{>0}$ and two subgroups $H_1, H_2 \subseteq \mathbb{Z}^n$, determines whether $H_1 \subseteq H_2$.*

Proof. Note that $H_1 \subseteq H_2$ if and only if $H_1 \cap H_2 = H_1$. Since $H_1 \cap H_2 \subseteq H_1$, testing equality is equivalent to testing whether the determinants of the two lattices, $H_1 \cap H_2$ and H_1, are equal. Computing determinants of lattices reduces to computing determinants of integer matrices, which can be done in polynomial time. □

Corollary 2.1.10. *There exists a deterministic polynomial-time algorithm such that, given $n \in \mathbb{Z}_{>0}$ and two subgroups $H_1, H_2 \subseteq \mathbb{Z}^n$, determines whether $H_1 = H_2$.*

2.2 Hermite and Smith normal forms

This section draws on Section 2.4 of [19].

There are two canonical forms of a matrix A that are of interest: the *Hermite normal form* and the *Smith normal form*. These can be obtained by applying row and column operations to A.

Definition 2.2.1. *Let $m, n \in \mathbb{Z}_{>0}$ and $A \in \mathcal{M}_{n \times m}(\mathbb{Z})$. A* column operation *on A is one of the following:*

(i) *interchanging two columns of A,*
(ii) *multiplying one column of A by -1,*
(iii) *adding a nonzero multiple of a column of A to another column.*

Note 2.2.2. (i) Each column operation corresponds to postmultiplying A with an appropriate invertible matrix over $\mathbb{Z}$.
(ii) If A' is a matrix obtained from A via a sequence of column operations, then there exists an invertible matrix V such that $A' = AV$. Conversely, if two matrices differ by a postmultiplied invertible matrix, then one can be obtained from the other by a series of column operations.
(iii) Applying column operations to a square matrix does not change the absolute value of its determinant.

Note 2.2.3. We can similarly define *row operations*. These correspond to premultiplying by a certain invertible matrix over $\mathbb{Z}$.

It is easy to see that performing column operations on a matrix does not change the lattice the columns generate.

Proposition 2.2.4. *Let $A, B \in \mathcal{M}_{n\times m}(\mathbb{Z})$. Then the lattice generated by the columns of A is equal to the lattice generated by the columns of B if and only if there exists $V \in \mathrm{GL}_m(\mathbb{Z})$ such that $AV = B$.*

Note 2.2.5. Let F be a free $\mathbb{Z}$-module of finite rank. In choosing to represent a subgroup $H \hookrightarrow F$ via a matrix $A \in \mathcal{M}_{n\times m}(\mathbb{Z})$, we are making a choice of basis of H and of F. Applying column operations to A corresponds to a change of basis of H, while keeping the basis for F fixed. Applying row operations corresponds to a change of basis of F, while keeping the basis for H fixed.

We are now ready to introduce the Hermite normal form, which is useful for representing subgroups of $\mathbb{Z}^n$ in a canonical way.

Definition 2.2.6. *Let $A = (a_{i,j}) \in \mathcal{M}_{n\times m}(\mathbb{Z})$, for some $m, n \in \mathbb{Z}_{>0}$. Then A is said to be in* Hermite normal form (HNF) *if there exists $0 \leq k \leq m$ such that the last $m - k$ columns are zero and for each $1 \leq j \leq k$, there exists an entry $a_{i_j,j} > 0$ such that*

(i) *For all $i' < i_j$, we have $a_{i',j} = 0$.*
(ii) *For all $j' < j$, we have $a_{i_j,j} > a_{i_j,j'} \geq 0$.*
(iii) *For all $j' < j$, we have $i_{j'} < i_j$.*

Note 2.2.7. The nonzero entry $a_{i_j,j}$ is called the *leading coefficient* of the j^{th} column. Informally, a matrix is in Hermite normal form if all its zero columns lie on the right, the leading coefficients of all nonzero columns are strictly positive and have nonnegative and strictly smaller entries to their left, and occur strictly below the position of the leading coefficient of the previous column, if this exists.

Note 2.2.8. We have seen that applying column operations to a matrix does not change the lattice it generates. Thus, finding the Hermite normal form of a matrix corresponds to finding a basis of the associated lattice, such that the basis vectors can be ordered in such a way that they have an increasing number of leading zero entries.

Proposition 2.2.9. (i) *Each integer matrix can transformed into a matrix in Hermite normal form by a sequence of column operations.*

(ii) ([12], Section 5.3) *Each integer matrix has a unique Hermite normal form, i.e. if $A = (a_{i,j}) \in \mathcal{M}_{n\times m}(\mathbb{Z})$, then there exists a matrix $V \in \mathrm{GL}_m(\mathbb{Z})$ such that AV is in Hermite normal form. If there is another $V' \in \mathrm{GL}_m(\mathbb{Z})$ such that AV' is in Hermite normal form, then $AV = AV'$.*

Sketch of proof of (i). Let $A = (a_{i,j}) \in \mathcal{M}_{n\times m}(\mathbb{Z})$. Suppose that A has block form

$$A = \begin{bmatrix} A_0 & \mathbf{0} \end{bmatrix},$$

where A_0 is an $n \times m'$ matrix, for some $m' \leq m$, with no nonzero columns. Otherwise interchange columns to arrive at this form. Let $a_{01}, \ldots, a_{0m'}$ be the entries in the first nonzero row of A_0. Then at least one a_{0i} must be nonzero and we can ensure that they are all nonnegative (by applying suitable column operations), so using the extended Euclid algorithm (see [33]), we can compute $g := \gcd(\{a_{0i}\}_i)$. This reduces to applying a sequence of column operations at the end of which the first nonzero row of A will be $[g \;\; 0 \;\; 0 \ldots 0]$. The matrix now has form

$$\begin{bmatrix} \mathbf{0} & \mathbf{0} & \mathbf{0} \\ g & 0 & 0 \\ \star & A_1 & \mathbf{0} \end{bmatrix}.$$

We proceed by computing the greatest common divisor of the entries in the first nonzero row of A_1. We perform a couple of extra column operations to ensure that the entry to the left of the leading entry of the second column is strictly smaller than the leading entry. We continue in this way until we attain the Hermite normal form. □

Note 2.2.10. The matrix V in Proposition 2.2.9, part (ii), need not be unique.

Note 2.2.11. Since the Hermite normal form is unique, we see that the leading entry of the first column is the greatest common divisor of the entries in the first nonzero row of the original matrix.

There are many algorithms available in the literature that compute the Hermite normal form of a given integer matrix. The main difficulty in achieving polynomial time is to keep the entries of the intermediate matrices small. The straightforward column-operation-algorithm presented in Proposition 2.2.9 suffers from coefficient blow-up. This can be avoided by using modular techniques, for example by working modulo an integer d, where d is chosen to be the determinant of a full rank submatrix of A. Another way to circumvent coefficient blow-up is to employ lattice basis reduction techniques.

For detailed accounts of deterministic polynomial-time algorithms for computing the Hermite normal form of an integer matrix, together with a transformation matrix V, see [16, 30, 33, 85]. We will only record their existence.

Theorem 2.2.12 ([12], Proposition 5.4)**.** *There exists a deterministic polynomial-time algorithm that, given a matrix $A = (a_{i,j}) \in \mathcal{M}_{n\times m}(\mathbb{Z})$, computes a matrix $V \in \mathrm{GL}_m(\mathbb{Z})$ such that AV is in Hermite normal form.*

Note 2.2.13. Theorem 2.2.12 can also be obtained from Theorem 2.1.6 if we equip $\mathcal{L}(A)$ with a suitable "length function" $q : L \to \mathbb{R}$ (see [68], Section 14).

The second canonical form we wish to examine is the Smith normal form.

Definition 2.2.14. *Let $A = (a_{i,j}) \in \mathcal{M}_{n\times m}(\mathbb{Z})$, for some $m, n \in \mathbb{Z}_{\geq 0}$. Then A is said to be in* Smith normal form (SNF) *if there exists $0 \leq k \leq \min\{n, m\}$ such that the last $n-k$ rows and the last $m-k$ columns are zero and the matrix $(a_{i,j})_{i=1,j=1}^{k,k}$ is diagonal, with $a_{i,i} > 0$ for all $1 \leq i \leq k$, and $a_{i,i} \mid a_{i+1,i+1}$ for all $1 \leq i < k$.*

Note 2.2.15. The nonzero entries are called *elementary divisors* of A.

The following is a standard result.

Proposition 2.2.16. *Every integer matrix has a unique Smith normal form, i.e. if $A = (a_{i,j}) \in \mathcal{M}_{n\times m}(\mathbb{Z})$, then there exist matrices $U \in \mathrm{GL}_n(\mathbb{Z})$ and $V \in \mathrm{GL}_m(\mathbb{Z})$ such that UAV is in Smith normal form. If there exist other $U' \in \mathrm{GL}_n(\mathbb{Z})$ and $V' \in \mathrm{GL}_m(\mathbb{Z})$ such that $U'AV'$ is in Smith normal form, then $U'AV' = UAV$.*

Theorem 2.2.17 ([86], Section 8.2)**.** *There exists a deterministic polynomial-time algorithm that, given a matrix $A = (a_{i,j}) \in \mathcal{M}_{n\times m}(\mathbb{Z})$, computes the Smith normal form of A, together with transformation matrices $U \in \mathrm{GL}_n(\mathbb{Z})$ and $V \in \mathrm{GL}_m(\mathbb{Z})$.*

The relevance of the Hermite and Smith normal forms to the study of finitely generated abelian groups lies in the following theorem.

Theorem 2.2.18. *Let $n \in \mathbb{Z}_{>0}$. Suppose $H \subseteq \mathbb{Z}^n$ is a subgroup.*

(i) *There exists a unique full column rank matrix A in Hermite normal form such that H is generated over $\mathbb{Z}$ by the columns of A.*
(ii) *If H has rank n, then there exists a unique square matrix A in Smith normal form such that*

$$\mathbb{Z}^n/H \cong \bigoplus_{i=1}^{n} \mathbb{Z}/d_i\mathbb{Z}, \tag{2.1}$$

as $\mathbb{Z}$-modules, where $d_1, \ldots, d_n$ are the elementary divisors of A.

Proof. Part (i) is a consequence of Propositions 2.2.4 and 2.2.9. For a proof of part (ii), see Theorem 2.4.13 of [19]. □

2.3 Representing objects and basic constructions

Given that we wish to bound the running time of an algorithm in terms of the length of the input, it is of crucial importance to make clear how we represent objects inside algorithms. Different representations may lead to essentially different computational tasks, with different complexities.

There are several ways to represent groups inside an algorithm. These include giving a finite presentation, giving the group as a permutation group of a finite set or a matrix group over a ring, black-box representations, or giving the group as automorphisms of certain objects (e.g. graphs, field extensions). For more details on group representations and the algorithmic problems they give rise to, see [35, 38].

For finitely generated abelian groups, matters simplify greatly, since these are nothing else than $\mathbb{Z}$-modules and thus can be represented by matrices.

2.3.1 Representing finite and finitely generated abelian groups

The proof of the following result can be found in any introductory algebra textbook.

Theorem 2.3.1. *Let G be a finitely generated abelian group. Then:*

(i) *There exists $k \in \mathbb{Z}_{\geq 0}$ such that $G \cong \mathbb{Z}^k \oplus H$, where H is a finite abelian group.*
(ii) *If G is finite, then there exist $n \in \mathbb{Z}_{>0}$ and a subgroup $L \subseteq \mathbb{Z}^n$ of rank n, such that $G \cong \mathbb{Z}^n/L$.*
(iii) (Fundamental Theorem of Finite Abelian Groups) *If G is finite, then there exists a unique $t \in \mathbb{Z}_{\geq 0}$ and a unique sequence of integers $d_1, \ldots, d_t \in \mathbb{Z}_{>1}$ such that $d_1 \mid d_2 \mid \ldots \mid d_t$ and*

$$G \cong \bigoplus_{i=1}^{t} \mathbb{Z}/d_i\mathbb{Z}.$$

To represent a group G, we give the algorithm a set of generators and relations. More precisely, suppose G has generators $x_1, \ldots, x_n$ and relations in G are of the form $\sum_{i=1}^{n} a_{ij}x_i$ for $1 \leq j \leq m$, where $a_{ij} \in \mathbb{Z}$ for all i, j. Then the matrix $f = (a_{ij}) \in \mathcal{M}_{n\times m}(\mathbb{Z})$ is said to be a *presentation matrix* of G. Consider the exact sequence

$$\mathbb{Z}^m \xrightarrow{f} \mathbb{Z}^n \to \operatorname{coker}(f) \to 0,$$

Then $\operatorname{coker}(f) \cong G$ and an element $g \in \operatorname{coker}(f)$ corresponds, in a non-unique way, to a vector in $\mathbb{Z}^n$ mapping to g, i.e. it is specified as a $\mathbb{Z}$-linear combination of generators.

Definition 2.3.2. *Let G be a finitely generated abelian group. An* exact-sequence representation *of G consists of a triple (m, n, f), where $m, n \in \mathbb{Z}_{>0}$ and $f \in \operatorname{Hom}(\mathbb{Z}^m, \mathbb{Z}^n)$ are such that $\operatorname{coker}(f) \cong G$.*

Proposition 2.3.3. *Let G be a finitely generated abelian group and let $f \in \mathcal{M}_{n\times m}(\mathbb{Z})$ be a presentation matrix of G. If $n = m$, then G is finite if and only if $\det(f) \neq 0$. In this case, $|G| = |\det(f)|$.*

Proof. This follows from Theorem 2.2.18, Theorem 2.3.1. □

From now on, given a presentation matrix f for a finitely generated abelian group G, we will identify G with $\operatorname{coker}(f)$.

2.3.2 Group homomorphisms

Let G_1, G_2 be two finitely generated abelian groups. Suppose G_1 and G_2 are represented by triples (m_1, n_1, f_1) and (m_2, n_2, f_2) respectively. Then we have two exact sequences:

$$\begin{array}{ccccccc} \mathbb{Z}^{m_1} & \xrightarrow{f_1} & \mathbb{Z}^{n_1} & \xrightarrow{\pi_1} & G_1 & \longrightarrow & 0 \\ & & \downarrow & & \downarrow & & \\ \mathbb{Z}^{m_2} & \xrightarrow{f_2} & \mathbb{Z}^{n_2} & \xrightarrow{\pi_2} & G_2 & \longrightarrow & 0. \end{array}$$

Any group homomorphism $G_1 \to G_2$ is induced by a map $\mathbb{Z}^{n_1} \to \mathbb{Z}^{n_2}$. This is the same as saying that a group homomorphism is determined by the images of the generators. However, not all assignments of generators correspond to well-defined group homomorphisms. That is to say, not every map $g : \mathbb{Z}^{n_1} \to \mathbb{Z}^{n_2}$ gives rise to a group homomorphism $\overline{g} : G_1 \to G_2$.

Proposition 2.3.4. *There exists a deterministic polynomial-time algorithm that, given two finitely generated abelian groups G_1 and G_2, represented by triples (m_1, n_1, f_1) and (m_2, n_2, f_2) respectively, and a map $g : \mathbb{Z}^{n_1} \to \mathbb{Z}^{n_2}$, decides whether or not g induces a group homomorphism $\overline{g} : G_1 \to G_2$.*

Proof. The map induced by g takes an element of G_1, lifts it to $\mathbb{Z}^{n_1}$ and then maps it to G_2 under $\pi_2 \circ g$. For this to be a well-defined group homomorphism, we must ensure that it is independent of the lift to $\mathbb{Z}^{n_1}$ we choose. For this, we require that $\operatorname{im}(f_1) \subseteq \ker(\pi_2 \circ g)$, or equivalently, $\operatorname{im}(g \circ f_1) \subseteq \operatorname{im}(f_2)$.

By Theorem 2.1.6 and Corollary 2.1.9, we can compute kernels and images, and test for inclusion, so, given a map $g : \mathbb{Z}^{n_1} \to \mathbb{Z}^{n_2}$, we can test if g induces a group homomorphism $G_1 \to G_2$ by checking whether

$$\operatorname{im}(g \circ f_1) \subseteq \operatorname{im}(f_2).$$

□

Note 2.3.5. The condition $\operatorname{im}(g \circ f_1) \subseteq \operatorname{im}(f_2)$ is equivalent to the existence of a map $h : \mathbb{Z}^{m_1} \to \mathbb{Z}^{m_2}$ such that $g \circ f_1 = f_2 \circ h$.

Once we have ensured that we have a well-defined group homomorphism, we may test whether it is injective or surjective.

Proposition 2.3.6. *There exists a deterministic polynomial-time algorithm that, given two finitely generated abelian groups G_1 and G_2, represented by triples (m_1, n_1, f_1) and (m_2, n_2, f_2) respectively, and a map $g : \mathbb{Z}^{n_1} \to \mathbb{Z}^{n_2}$ inducing a group homomorphism $\overline{g} : G_1 \to G_2$, decides whether $\overline{g}$ is injective.*

Proof. For $\overline{g}$ to be injective, we require that for all $x \in \mathbb{Z}^{n_1}$, if $g(x) \in \mathrm{im}(f_2)$, then $x \in \mathrm{im}(f_1)$, or equivalently, that

$$\mathrm{im}(f_1) \supseteq g^{-1}(\mathrm{im}(f_2)).$$

To express this condition in terms of the maps that are part of the input, i.e. in terms of f_1, f_2 and g, consider the map

$$-f_2 + g : \mathbb{Z}^{m_2} \oplus \mathbb{Z}^{n_1} \to \mathbb{Z}^{n_2}, \quad (x, y) \mapsto -f_2(x) + g(y). \tag{2.2}$$

Let $p_2 : \mathbb{Z}^{m_2} \oplus \mathbb{Z}^{n_1} \to \mathbb{Z}^{n_1}$ be the projection map to the second component. It is easy to see that

$$g^{-1}(\mathrm{im}(f_2)) = p_2(\ker(-f_2 + g)).$$

Thus, the condition for injectivity becomes

$$p_2(\ker(-f_2 + g)) \subseteq \mathrm{im}\, f_1,$$

which can be tested deterministically in polynomial time using Theorem 2.1.6 and Corollary 2.1.9. □

Proposition 2.3.7. *There exists a deterministic polynomial-time algorithm that, given two finitely generated abelian groups G_1 and G_2, represented by triples (m_1, n_1, f_1) and (m_2, n_2, f_2) respectively, and a map $g : \mathbb{Z}^{n_1} \to \mathbb{Z}^{n_2}$ inducing a group homomorphism $\overline{g} : G_1 \to G_2$, decides whether $\overline{g}$ is surjective.*

Proof. For $\overline{g} : G_1 \to G_2$ to be surjective, we require that for all $\overline{x} \in G_2 = \mathbb{Z}^{n_2}/\mathrm{im}(f_2)$, there exist $y \in \mathbb{Z}^{n_1}$ such that $\pi_2 \circ g(y) = \overline{x}$. This is equivalent to requiring that

$$\forall x \in \mathbb{Z}^{n_2}, \quad \exists y \in \mathbb{Z}^{n_1} \text{ such that } g(y) - x \in \mathrm{im}(f_2),$$

which is further equivalent to the map $-f_2 + g$ defined in (2.2), being surjective. We can check this using Theorem 2.1.6. □

Proposition 2.3.8. *There exists a deterministic polynomial-time algorithm that, given two finitely generated abelian groups G_1 and G_2, represented by triples (m_1, n_1, f_1) and (m_2, n_2, f_2) respectively, and two maps $g, h : \mathbb{Z}^{n_1} \to \mathbb{Z}^{n_2}$ inducing group homomorphisms $\overline{g}, \overline{h} : G_1 \to G_2$, decides whether $\overline{g} = \overline{h}$.*

Proof. The problem of deciding if two group homomorphisms are equal is equivalent to deciding if a given group homomorphism is the zero homomorphism. Let $\overline{g} : G_1 \to G_2$ be a group homomorphism induced by the map $g : \mathbb{Z}^{n_1} \to \mathbb{Z}^{n_2}$. Then $\overline{g} \equiv 0$ if and only if $\mathrm{im}(g) \subseteq \mathrm{im}(f_2)$. □

2.3.3 Subgroups and quotients

Representing subgroups

Let G be a finitely generated abelian group. To represent a subgroup H of G, we need to represent it as a group and additionally, to produce an embedding $H \hookrightarrow G$ that tells us how H sits inside G, i.e. we need to give the algorithm a triple (r, s, h) and a map of free $\mathbb{Z}$-modules g, that induces an injective map $H \hookrightarrow G$.

Example 2.3.9. It is important to specify the injection: if $G = \mathbb{Z}/2\mathbb{Z} \oplus \mathbb{Z}/2\mathbb{Z}$, then the subgroup $\mathbb{Z}/2\mathbb{Z}$ sits inside G in three different ways.

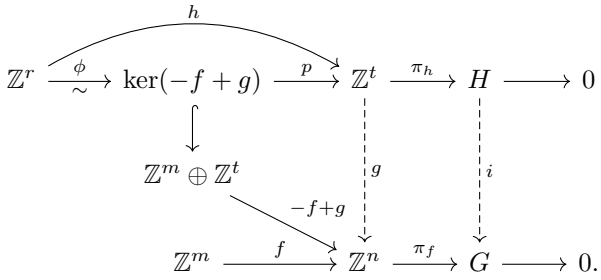
Proposition 2.3.10. *There exists a deterministic polynomial-time algorithm that, given a finitely generated abelian group G and a finite set $\overline{S} \subset G$, computes the subgroup $H \leq G$ generated by $\overline{S}$.*

Proof. The desired output of the algorithm is a triple (r, s, h) representing H as a group, together with an injective group homomorphism $i : H \hookrightarrow G$. Suppose the group G is represented by triple (m, n, f).

Let $S = \{s_1, \ldots, s_t\} \subset \mathbb{Z}^n$ be a set of representatives of the elements of $\overline{S}$ in G. Consider the map $g : \mathbb{Z}^t \to \mathbb{Z}^n$ given by $e_i \mapsto s_i$, where for $1 \leq i \leq t$, the vector e_i is the i^{th} canonical basis element of $\mathbb{Z}^t$:

$$\begin{array}{ccccccc} & & & & \mathbb{Z}^t & & \\ & & & & \downarrow{\scriptstyle g} & & \\ \mathbb{Z}^m & \xrightarrow{f} & \mathbb{Z}^n & \xrightarrow{\pi_f} & G & \longrightarrow & 0. \end{array} \tag{2.3}$$

Consider the map $(-f + g) : \mathbb{Z}^m \oplus \mathbb{Z}^t \to \mathbb{Z}^n$, given by $(x, y) \mapsto -f(x) + g(y)$. Compute $\ker(-f + g) = \{(x, y) \in \mathbb{Z}^m \oplus \mathbb{Z}^t \mid f(x) = g(y)\}$ and project it to $\mathbb{Z}^t$ via a map p. Put $r := \operatorname{rank}(\ker(-f + g))$ and let ϕ be the matrix whose columns are the basis vectors of $\ker(-f + g)$. The following diagram illustrates what was described above:



Put $h := p \circ \phi$. Then (r, t, h) represents H. Note that t is not necessarily the minimum number of generators of H.

The map g induces a map $i : H \to G$ in the following way: an element of $H = \mathbb{Z}^t / \operatorname{im}(h)$ is lifted to $\mathbb{Z}^t$ and then is mapped to G. We claim that $i : H \to G$ is an injective group homomorphism.

First we show that i is independent of the chosen lift to $\mathbb{Z}^t$. Let $\overline{x} \in H$ and lift it to $x+k \in \mathbb{Z}^t$, for some $k \in \mathrm{im}(h)$. Then $g(x+k) = g(x)+g(k)$. But, since $k \in \mathrm{im}(h)$, it follows that k is the image under p of some $(y,k) \in \ker(-f+g)$, with $y \in \mathbb{Z}^m$. Then $g(k) = f(y)$. Hence $\pi_f g(k) = 0$ and so $\overline{g}$ does not depend on the lift.

For injectivity, suppose $i(\overline{y}) = 0_G$, for some $\overline{y} \in H$, where 0_G denotes the zero element in G. Let y be a lift of $\overline{y}$ to $\mathbb{Z}^t$. Then $\pi_f(g(y)) = 0_G$, so $g(y) = f(x)$, for some $x \in \mathbb{Z}^t$. But then $(x,y) \in \ker(-f+g)$ and so $y \in \mathrm{im}(p)$, i.e. $\overline{y} = 0_H$.

Moreover, by construction, $\mathrm{im}(i) = \langle\{\overline{s_i}\}\rangle = H$, where $\overline{s_i} := \pi_f(s_i)$. □

Example 2.3.11. Let G be any group. If $S = \emptyset$, then $t = 0$, and $H = 0$, as expected.

Example 2.3.12. Let $G = \mathbb{Z}/2\mathbb{Z} \oplus \mathbb{Z}/4\mathbb{Z} \oplus \mathbb{Z}/4\mathbb{Z}$.

(i) Let $S = \{(1,1,2)\}$. Then $H := \langle S\rangle \cong \mathbb{Z}/4\mathbb{Z}$. The map $g - f : \mathbb{Z}^3 \oplus \mathbb{Z} \to \mathbb{Z}^3$ is given by $(x_1,x_2,x_3,y) \mapsto (y-2x_1, y-4x_2, 2y-4x_3)$, so $\ker(g-f) = (2,1,2,4)\mathbb{Z}$. Hence $\mathrm{im}(\pi) = 4\mathbb{Z}$ and the subgroup H is given by $(1,1,4\cdot\mathrm{id})$.

(ii) Let $S = \{(1,1,2),(0,2,0)\}$. Again, $H := \langle S\rangle \cong \mathbb{Z}/4\mathbb{Z}$. The map $g-f : \mathbb{Z}^3 \oplus \mathbb{Z}^2 \to \mathbb{Z}^3$ is given by $(x_1,x_2,x_3,y_1,y_2) \mapsto (2y_1 - 4x_1, y_1 + 2y_2 - 4x_2, 2y_1 - 4x_3)$, so $\ker(g-f) = (1,0,1,2,1)\mathbb{Z} \oplus (0,1,0,0,2)\mathbb{Z}$. Hence $\mathrm{im}(\pi) = (2,1)\mathbb{Z} \oplus (0,2)\mathbb{Z}$ and the subgroup H is given by $(2,2,h)$, where h is the map represented by the matrix

$$\begin{pmatrix} 2 & 1 \\ 0 & 2 \end{pmatrix}.$$

It can be checked that the Smith normal form of this matrix has diagonal $(4,1)$, so $(2,2,h)$ indeed represents the cyclic group of order 4.

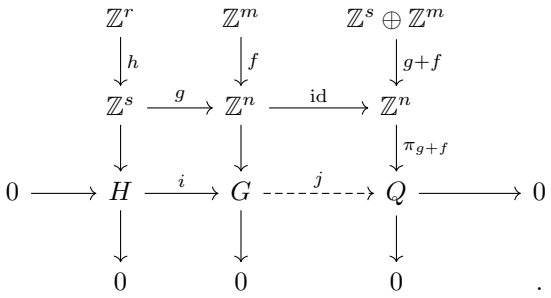
Representing quotients

Let G be a finitely generated abelian group and let $H \leq G$ be a subgroup. To represent the quotient G/H, we need to represent it as a group, which we denote by Q, and additionally, to produce a surjection $j : G \twoheadrightarrow Q$. Then $Q \cong G/H$.

Proposition 2.3.13. *There exists a deterministic polynomial-time algorithm that, given a finitely generated abelian group G and a subgroup $H \leq G$, computes the quotient group G/H.*

Proof. Suppose the group G is represented by triple (m,n,f) and H is represented by triple (r,s,h), together with a map $g : \mathbb{Z}^s \to \mathbb{Z}^n$ inducing an injection $i : H \to G$, as in the previous subsection. Let $g+f : \mathbb{Z}^s \oplus \mathbb{Z}^m \to \mathbb{Z}^n$ be the map given by

$(x, y) \mapsto g(x) + f(y)$. Consider the following diagram

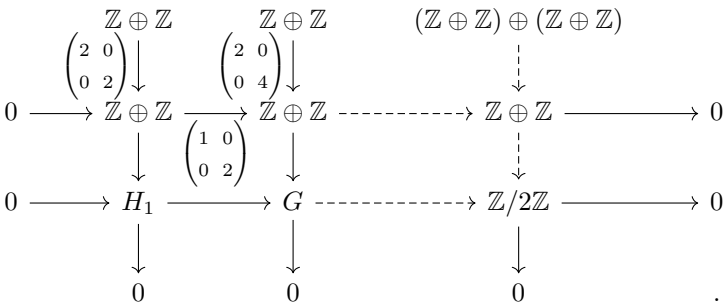


Then G/H is represented by triple $(s + m, n, g + f)$.

The map $g + f$ induces a map $j : G \to Q$ in the following way: an element of $G = \mathbb{Z}^n/\operatorname{im}(f)$ is lifted to $\mathbb{Z}^n$ and then mapped to Q via $\pi_{g+f} \circ \mathrm{id}$. We claim that this map is a surjective group homomorphism.

To see that it is well-defined, let $\overline{x} \in G$ and a consider a lift of $\overline{x}$ to $\mathbb{Z}^n$ given by $x + f(y)$, for some $y \in \mathbb{Z}^m$. Since $f(y) \in \operatorname{im}(g + f)$ by construction, it is sent to zero under π_{g+f}. Hence j is independent of the choice of lift. Surjectivity follows by construction. □

Example 2.3.14. Let $G = \mathbb{Z}/2\mathbb{Z} \oplus \mathbb{Z}/4\mathbb{Z}$. Let $H_1 = \mathbb{Z}/2\mathbb{Z} \oplus \mathbb{Z}/2\mathbb{Z}$. Then the quotient G/H_1 is given by data

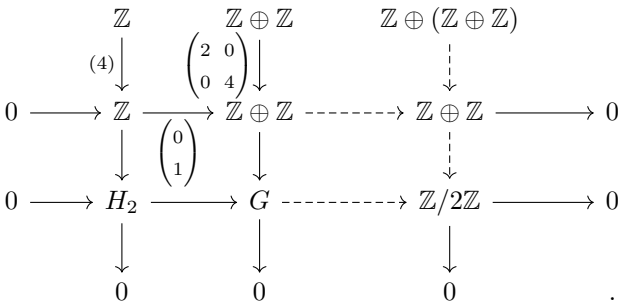


The presentation matrix for G/H_1 is

$$f_{G/H_1} = \begin{pmatrix} 1 & 0 & 2 & 0 \\ 0 & 2 & 0 & 4 \end{pmatrix},$$

whose Smith normal form has nonzero entries $(2\ \ 1)$, and so $G/H_1 \cong \mathbb{Z}/2\mathbb{Z}$.

If we let $H_2 = \mathbb{Z}/4\mathbb{Z}$, the quotient G/H_2 is given by data



Then the presentation matrix for G/H_2 is

$$f_{G/H_2} = \begin{pmatrix} 0 & 2 & 0 \\ 1 & 0 & 4 \end{pmatrix},$$

whose Smith normal form again has nonzero entries $(2\ 1)$, and so $G/H_2 \cong \mathbb{Z}/2\mathbb{Z}$.

Proposition 2.3.15. *There exists a deterministic polynomial-time algorithm that, given two finitely generated abelian groups G_1 and G_2, represented by triples (m_1, n_1, f_1) and (m_2, n_2, f_2) respectively, and a map $g : \mathbb{Z}^{n_1} \to \mathbb{Z}^{n_2}$, computes*

(i) $\ker(\overline{g})$, *together with the corresponding injective group homomorphism* $\ker(\overline{g}) \to G_1$,
(ii) $\operatorname{im}(\overline{g})$, *together with the corresponding injective group homomorphism* $\operatorname{im}(\overline{g}) \to G_2$,
(iii) $\operatorname{coker}(\overline{g})$, *together with the corresponding surjective group homomorphism* $G_2 \to \operatorname{coker}(\overline{g})$,

where $\overline{g} : G_1 \to G_2$ is the group homomorphism induced by g.

Proof. Kernels can be computed as:

$$\begin{aligned} \ker(G_1 \to G_2) &= \ker\left(\mathbb{Z}^{n_1}/\operatorname{im}(f_1) \longrightarrow \mathbb{Z}^{n_2}/\operatorname{im}(f_2)\right) \\ &= \{x \in \mathbb{Z}^{n_1} \mid \pi_2 \circ g(x) = 0\}/\operatorname{im}(f_1) \\ &= \ker(\pi_2 \circ g)/\operatorname{im}(f_1) \\ &= p_2(\ker(-f_2 + g))/\operatorname{im}(f_1). \end{aligned}$$

Images can be computed as:

$$\begin{aligned} \operatorname{im}(G_1 \to G_2) &= \operatorname{im}(\pi_2 \circ g) \\ &= \operatorname{im}(g)/(\operatorname{im}(g) \cap \operatorname{im}(f_2)) \\ &= \operatorname{im}(-f_2 + g)/\operatorname{im}(f_2). \end{aligned}$$

Cokernels can be computed as:

$$\begin{aligned}\operatorname{coker}(G_1 \to G_2) &= G_2/\operatorname{im}(G_1 \to G_2)\\ &= (\mathbb{Z}^{n_2}/\operatorname{im}(f_2))\,/\,(\operatorname{im}(-f_2+g)/\operatorname{im}(f_2))\\ &= \operatorname{coker}(-f_2+g).\end{aligned}$$

All these computations can be carried out in polynomial time by Theorem 2.1.6 and Proposition 2.1.8. The maps $\ker(\overline{g}) \to G_1$, $\operatorname{im}(\overline{g}) \to G_2$ and $G_2 \to \operatorname{coker}(\overline{g})$ can be obtained from Propositions 2.3.10 and 2.3.13. □

2.3.4 Direct sums of groups

Proposition 2.3.16. *There exists a deterministic polynomial-time algorithm that, given two finitely generated abelian group G_1, G_2, computes the direct sum $G_1 \oplus G_2$.*

Proof. Suppose G_1 and G_2 are represented via triples (m_1, n_1, f_1) and (m_2, n_2, f_2) respectively. Then the direct sum $G_1 \oplus G_2$ is represented by triple (m_1+m_2, n_1+n_2, F), where F is an $(m_1+m_2) \times (n_1+n_2)$ integer matrix with block form

$$F = \begin{pmatrix} f_1 & 0 \\ 0 & f_2 \end{pmatrix}.$$

□

2.4 Homomorphism groups and tensor products

In the case of free abelian groups of finite rank, the tensor product is easy to construct: if $\{x_i\}_{i=1}^n$ is a basis for $\mathbb{Z}^n$ and $\{y_i\}_{i=1}^m$ is a basis for $\mathbb{Z}^m$, then the set $\{x_i \otimes y_j \mid 1 \le i,j \le n\}$ is a basis for $\mathbb{Z}^n \otimes \mathbb{Z}^m$. Constructing homomorphism groups of free abelian groups of finite rank is also easy, since $\operatorname{Hom}(\mathbb{Z}^m, \mathbb{Z}^n) = \mathcal{M}_{n\times m}(\mathbb{Z})$.

Suppose now that G_1, G_2 are two finitely generated abelian groups with presentations

$$\mathbb{Z}^{m_1} \xrightarrow{f_1} \mathbb{Z}^{n_1} \longrightarrow G_1 \longrightarrow 0 \tag{2.4}$$

and

$$\mathbb{Z}^{m_2} \xrightarrow{f_2} \mathbb{Z}^{n_2} \longrightarrow G_2 \longrightarrow 0 \tag{2.5}$$

respectively.

Proposition 2.4.1. *There exists a deterministic polynomial-time algorithm that, given two finitely generated abelian groups G_1, G_2, computes*

(i) *the tensor product, $G_1 \otimes G_2$, together with the corresponding bilinear map $G_1 \times G_2 \to G_1 \otimes G_2$,*
(ii) *the homomorphism group, $\operatorname{Hom}(G_1, G_2)$, together with the corresponding bilinear map $\operatorname{Hom}(G_1, G_2) \times G_1 \to G_2$.*

Proof. Suppose G_1, G_2 have presentations (2.4) and (2.5), respectively.

Recall that tensoring is right-exact, so tensoring (2.4) with G_2 and tensoring (2.5) with $\mathbb{Z}^{m_1}$ and $\mathbb{Z}^{n_1}$ gives

$$\begin{array}{ccccccc}
\mathbb{Z}^{m_1}\otimes\mathbb{Z}^{m_2} & & \mathbb{Z}^{n_1}\otimes\mathbb{Z}^{m_2} & & & & \\
\downarrow & & \downarrow & & & & \\
\mathbb{Z}^{m_1}\otimes\mathbb{Z}^{n_2} & & \mathbb{Z}^{n_1}\otimes\mathbb{Z}^{n_2} & & & & \\
\downarrow & & \downarrow & & & & \\
\mathbb{Z}^{m_1}\otimes G_2 & \longrightarrow & \mathbb{Z}^{n_1}\otimes G_2 & \longrightarrow & G_1\otimes G_2 & \longrightarrow & 0. \\
\downarrow & & \downarrow & & & & \\
0 & & 0 & & & &
\end{array}$$

By construction of quotient groups (Theorem 2.3.13), it follows that $G_1 \otimes G_2$ has presentation given by

$$(\mathbb{Z}^{m_1}\otimes\mathbb{Z}^{n_2}) \oplus (\mathbb{Z}^{n_1}\otimes\mathbb{Z}^{m_2}) \longrightarrow \mathbb{Z}^{n_1}\otimes\mathbb{Z}^{n_2} \longrightarrow G_1\otimes G_2 \longrightarrow 0,$$

where the first map is given by $(f_1 \otimes \mathrm{id}) + (\mathrm{id} \otimes f_2)$.

Applying $\mathrm{Hom}(-, G_2)$ to (2.4) gives

$$0 \longrightarrow \mathrm{Hom}(G_1, G_2) \longrightarrow \mathrm{Hom}(\mathbb{Z}^{n_1}, G_2) \longrightarrow \mathrm{Hom}(\mathbb{Z}^{m_1}, G_2).$$

Hence

$$\mathrm{Hom}(G_1, G_2) = \ker\left(\mathrm{Hom}(\mathbb{Z}^{n_1}, G_2) \xrightarrow{\circ f_1} \mathrm{Hom}(\mathbb{Z}^{m_1}, G_2), h \mapsto h \circ f_1\right),$$

which can be computed using Proposition 2.3.15. Now, for $k \in \mathbb{Z}_{\geq 0}$, we have that $\mathrm{Hom}(\mathbb{Z}^k, -)$ is an exact functor, so

$$\mathrm{Hom}(\mathbb{Z}^k, G_2) = \mathrm{coker}\left(\mathrm{Hom}(\mathbb{Z}^k, \mathbb{Z}^{m_2}) \xrightarrow{f_2\circ} \mathrm{Hom}(\mathbb{Z}^k, \mathbb{Z}^{n_2}), h \mapsto f_2 \circ h\right),$$

which we compute for k equal to n_1 and m_1.

The bilinear maps $G_1 \times G_2 \to G_1 \otimes G_2$ and $\mathrm{Hom}(G_1, G_2) \times G_1 \to G_2$ are represented by listing the images of all pairs of generators, and can be readily obtained from the above construction. □

2.5 Splitting exact sequences

Consider an exact sequence of finitely generated abelian groups

$$0 \longrightarrow H \xrightarrow{i} G \xrightarrow{p} K \longrightarrow 0. \tag{2.6}$$

Proposition 2.5.1. *There exists a deterministic polynomial-time algorithm that, given an exact sequence as in (2.6), decides if it is split, and if it is split, produces a right-inverse of p and a left-inverse of i.*

Proof. We consider the map

$$\psi : \mathrm{Hom}(K, G) \to \mathrm{Hom}(K, K), \quad k \mapsto p \circ k.$$

Deciding if the sequence is split reduces to deciding if id_K is in the image of ψ. This can be done by solving a system of linear equations over $\mathbb{Z}$ (Theorem 2.1.7). If the system has a solution, solving it will also produces a right inverse of p, i.e. an element $s \in \mathrm{Hom}(K, G)$ such that $\psi(s) = \mathrm{id}_K$. Similarly, we construct a left-inverse of i. □

Note 2.5.2. Suppose that the exact sequence (2.6) is right-split, i.e. there exists a group homomorphism $s : K \to G$ such that $p \circ s = \mathrm{id}_K$. Moreover, suppose that we have found such an s. Then we know that $G = i(H) \oplus s(K)$, and we can construct images of group homomorphism and direct sums of groups. The isomorphism $H \oplus K \xrightarrow{\sim} G$ is given by $(h, k) \mapsto i(h) + s(k)$, with inverse $G \xrightarrow{\sim} H \oplus K$ given by $g \mapsto (t(h), p(k))$, where t is a left splitting of i, that is, a group homomorphism such that $t \circ i = \mathrm{id}_H$.

2.6 Torsion subgroups, exponents, orders, cyclic decompositions

2.6.1 Computing the order of a finite abelian group

Recall that a finitely generated abelian group G represented by triple (m, n, f) is finite if and only if the basis of $\mathrm{im}\, f$ given by the algorithm in Theorem 2.1.6 has n elements.

Proposition 2.6.1. *There exists a deterministic polynomial-time algorithm that, given a finitely generated abelian group G, computes the order of G.*

Proof. To test if G is finite, we run the algorithm of Theorem 2.1.6 and check if the basis of $\mathrm{im}\, f$ has n elements. If so, we compute the determinant of the matrix whose columns are given by these n basis elements, which is then equal to $|G|$. Otherwise we conclude that G has infinite order. □

2.6.2 Computing the torsion subgroup of a finitely generated abelian group

Suppose we are given a finitely generated abelian group G. We have seen that we can determine if G is finite or not. If we find that it is not finite, we would like to find its torsion subgroup. To do this, we introduce a construction described in [68], at the end of Section 14.

Theorem 2.6.2. *There exists a deterministic polynomial-time algorithm that, given a finitely generated abelian group G, determines its torsion subgroup T, and produces an isomorphism $G \cong T \oplus G/T$.*

Proof. Suppose G is represented by triple (m, n, f). Let $L := \mathbb{Z}^n$ and $H = \mathrm{im}(f) \subseteq \mathbb{Z}^n$. Suppose $H \cong \mathbb{Z}^k$ and let $F \in \mathcal{M}_{n\times k}(\mathbb{Z})$ be a matrix whose columns are a $\mathbb{Z}$-basis of H. Consider the map

$$\phi : L \to \mathrm{Hom}(H, \mathbb{Z}), \quad \phi(x)(y) = \langle x, y\rangle,$$

for $x \in L$, $y \in H$ and $\langle\cdot,\cdot\rangle$ the standard inner product. Then ϕ is represented by the matrix $F^{\mathrm{t}} \in \mathcal{M}_{k\times n}(\mathbb{Z})$. This is because $\langle x, y\rangle = y^{\mathrm{t}}F^{\mathrm{t}}x$, for any $x \in L, y \in H$. We denote the kernel of this map by

$$H^{\perp} := \{x \in L \mid \langle x, H\rangle = 0\}.$$

Since $H^{\perp}$ is the kernel of a map between two free $\mathbb{Z}$-modules, Theorem 2.1.6 produces for us a basis $b_1 \ldots, b_n$ of L, where $b_1, \ldots, b_{n-k}$ is a $\mathbb{Z}$-basis of $H^{\perp}$.

We repeat the process above by considering the map

$$\phi' : L \to \mathrm{Hom}(H^{\perp}, \mathbb{Z}),$$

where the matrix representing ϕ' is given by the transpose of the matrix whose columns are $b_1, \ldots, b_{n-k}$. We denote the kernel of this map by

$$H^{\perp\perp} := \{a \in L \mid \langle a, H^{\perp}\rangle = 0\} = (\mathbb{Q} \cdot H) \cap L. \tag{2.7}$$

Thus,

$$L/H^{\perp\perp} \cong (L/H) \Big/ (L/H)_{\mathrm{tor}},$$

where $(L/H)_{\mathrm{tor}}$ denotes the torsion subgroup of L/H. Our goal becomes to split the exact sequence

$$0 \longrightarrow H^{\perp\perp}/H \longrightarrow L/H \longrightarrow L/H^{\perp\perp} \longrightarrow 0, \tag{2.8}$$

which we do using Proposition 2.5.1. □

2.6.3 Computing the order of a group element

Theorem 2.6.3. *There exists a deterministic polynomial-time algorithm that, given a finitely generated abelian group G and an element $g \in G$, determines the order of g.*

Proof. Consider the map $\psi : \mathbb{Z} \to G$, given by $1 \mapsto g$. If ψ is injective, which we can test, then g has infinite order. Otherwise, $\ker(\psi)$ is of the form $l\mathbb{Z}$, giving the order of g in G as equal to l. □

Note 2.6.4. This result depends heavily on the way we are representing G. Suppose $G = (\mathbb{Z}/n\mathbb{Z})^\times$, for some $n \in \mathbb{Z}_{>2}$. Suppose we choose to represent G by only giving the integer n. Then we can certainly carry out computations in G, but given the element $2 \in G$, we cannot in general efficiently compute its order without knowing the factorisation of n, and even then it is hard.

We can use this tool to decide if two elements are equal, by simply determining if their difference is equal to the zero element, i.e. if it has order 1.

Corollary 2.6.5. *There exists a deterministic polynomial-time algorithm that, given a finitely generated abelian group G and an element $g \in G$, determines if $g = 0_G$.*

Corollary 2.6.6. *There exists a deterministic polynomial-time algorithm that, given a finitely generated abelian group G and two elements $g, h \in G$, determines if $g = h$.*

2.6.4 Computing the exponent of a group

The exponent of a group G is computable as the generator over $\mathbb{Z}$ of

$$\ker\left(\mathbb{Z} \to \mathrm{Hom}(G, G),\ n \mapsto (x \mapsto nx)\right).$$

However, it is also useful to be able to exhibit an element of G of order equal to the exponent of G. To do this, we begin with a couple of preliminary results.

Lemma 2.6.7. *There exists a deterministic polynomial-time algorithm that, given a finite abelian group G and two elements $x, y \in G$ of orders n and m respectively, outputs an element z, expressed in terms of x and y, of order equal to* $\mathrm{lcm}(n, m)$.

Proof. Using Theorem 2.6.3, compute n and m, the orders of x and y, respectively. Apply the Coprime Base Algorithm (Theorem 1.1.2) to the set $\{n, m\}$ to obtain a set $\mathcal{P}$ of coprime divisors of nm and a factorisation $n = \prod_{p\in\mathcal{P}} p^{n_p}$ and $m = \prod_{p\in\mathcal{P}} p^{m_p}$, where $n_p, m_p \in \mathbb{Z}_{\geq 0}$. Define

$$n' = \prod_{\substack{p\in\mathcal{P}\\ n_p > m_p}} p^{n_p} \quad \text{and} \quad m' = \prod_{\substack{p\in\mathcal{P}\\ n_p \leq m_p}} p^{m_p}$$

Let $x' = \frac{n}{n'}x$ and $y' = \frac{m}{m'}y$. Since n' and m' are coprime, the order of $z := x' + y'$ is equal to $\mathrm{lcm}(n', m') = \mathrm{lcm}(n, m)$. □

Lemma 2.6.8. *Let $L = \mathbb{Z}^n$ and $H \subseteq L$ a subgroup. Suppose $B = \{b_1, \ldots, b_n\}$ is a basis of L. Then the exponent of L/H is equal to the lowest common multiple of the orders of $b_1, \ldots, b_n$ in L/H.*

Proof. Let $e_1, \ldots, e_n$ be the respective orders of $b_1 + H, \ldots, b_n + H$ in L/H. Let $l = \mathrm{lcm}_i(\{e_i\})$. By Lemma 2.6.7, there exists an element of L/H of order equal to l, so $\exp(L/H) \geq l$. Moreover, since B is a basis of L, every $x \in L/H$ is of the form $\sum_i \alpha_i b_i + H$, for some $\alpha_i \in \mathbb{Z}$. Then $lx \in H$, so x has order dividing l. Hence $\exp(L/H) = l$. □

This now enables us to compute the exponent of any given finite abelian group.

Theorem 2.6.9. *There exists a deterministic polynomial-time algorithm such that, given a finite abelian group G, computes the exponent of G and produces an element $g \in G$ of order equal to the exponent.*

Proof. We begin by determining the order of all basis vectors of L in L/H using Theorem 2.6.3. This gives a sequence of integers whose lowest common multiple is the exponent of L/H. Applying Proposition 2.6.7 repeatedly, we produce an element with the required property. □

2.6.5 Writing a finitely generated abelian group as a direct sum of cyclic groups

Let G be a finite abelian group of exponent n. Then we have seen that there exists $g \in G$ such that the order of g is n. Moreover, for any such g, the cyclic subgroup generated by it is a direct summand of G. This suggests a method of decomposing a finite abelian group into a direct sum of cyclic subgroups by computing the exponent of the group, producing an element of that order, quotienting out by the subgroup it generates and repeating the process for the remaining part.

Theorem 2.6.10. *There exists a deterministic polynomial-time algorithm that, given a finite abelian group G, produces a direct-sum-decomposition of G into cyclic subgroups.*

Proof. Suppose G is represented by triple (m, n, f), and put $L := \mathbb{Z}^n$ and $H := \mathrm{im}(f)$. Then $G = L/H$. By Theorem 2.6.9, we can compute e_1, the exponent of G, and produce an element $x_1 \in G$ of order e_1. The subgroup generated by x_1 is now a direct summand of G. We would now like to determine a subgroup $G_2 \leq G$ such that $G \cong G_2 \oplus \mathbb{Z}/e_1\mathbb{Z}$.

To do this, apply Proposition 2.5.1 to the exact sequence

$$0 \to \mathbb{Z}/e_1\mathbb{Z} \to G \to G/\mathbb{Z}x_1 \to 0.$$

Now replace G by G_2 and repeat. In the end we will have produced a positive integer $t \in \mathbb{Z}_{>0}$, a sequence of integers $e_1, \ldots, e_t \in \mathbb{Z}_{>0}$, a sequence of subgroups $G_1, \ldots, G_t \leq G$ such that the exponent of G_i is e_i, and a sequence of elements $x_1, \ldots, x_t$ such that $x_i \in G_i$ and the order of x_i is equal to e_i. Moreover, we have an isomorphism

$$G \xleftarrow[\sim]{\psi} \bigoplus_{i=1}^{t} \mathbb{Z}/e_i\mathbb{Z},$$

where $\psi : 1_{\mathbb{Z}/e_i\mathbb{Z}} \mapsto x_i$. The algorithm requires $t \leq \log_2 |G|$ iterations. □

Note 2.6.11. Another way to obtain a decomposition of G into cyclic subgroups is to apply Theorem 2.2.17 to the presentation matrix of G.

2.7 Homomorphism groups and tensor products reconsidered

We have already seen how to compute tensor products and homomorphism groups, without knowing a cyclic direct sum decomposition of the groups involved. The disadvantage of that approach is that the number of generators produced by the algorithm can get unnecessarily large. In this section we construct tensor products and homomorphism groups of finite abelian groups by making use of the fact that we can compute a cyclic direct sum decomposition.

Suppose

$$G_1 \cong \bigoplus_{i=1}^{t_1} \mathbb{Z}/c_i\mathbb{Z} \quad \text{and} \quad G_2 \cong \bigoplus_{j=1}^{t_2} \mathbb{Z}/d_i\mathbb{Z}, \tag{2.9}$$

for some $t_1, t_2 \in \mathbb{Z}_{>0}$ and $c_i, d_j \in \mathbb{Z}_{>0}$, for all $1 \leq i \leq t_1$ and $1 \leq j \leq t_2$.

Let $n, m \in \mathbb{Z}_{>0}$ and $d := \gcd(n, m)$. Then we have an isomorphism

$$\mathbb{Z}/n\mathbb{Z} \otimes \mathbb{Z}/m\mathbb{Z} \xrightarrow{\sim} \mathbb{Z}/d\mathbb{Z}, \quad x \otimes y \mapsto xy.$$

Similarly, every group homomorphism $\phi : \mathbb{Z}/n\mathbb{Z} \to \mathbb{Z}/m\mathbb{Z}$ is determined by the image of 1 in $\mathbb{Z}/m\mathbb{Z}$, which must be a multiple of $\frac{m}{d}$. This is because

$$n\phi(1) = \phi(n) \equiv 0 \mod m,$$

and so

$$\frac{n}{d}\phi(1) \equiv 0 \mod \frac{m}{d}.$$

Since n/d and m/d are coprime, it must be the case that $\phi(1) \equiv 0 \mod \frac{m}{d}$. Hence we have an isomorphism

$$\mathrm{Hom}(\mathbb{Z}/n\mathbb{Z}, \mathbb{Z}/m\mathbb{Z}) \xrightarrow{\sim} \mathbb{Z}/\gcd(n, m)\mathbb{Z}, \quad \phi \mapsto \left(\frac{m}{d}\right)^{-1} \phi(1).$$

Proposition 2.7.1. *There exists a deterministic polynomial-time algorithm that, given two finite abelian groups G_1, G_2 via direct-sum representations as in (2.9), computes*

(i) *the tensor product, $G_1 \otimes G_2$, together with the corresponding bilinear map $G_1 \times G_2 \to G_1 \otimes G_2$,*
(ii) *the homomorphism group, $\mathrm{Hom}(G_1, G_2)$, together with the corresponding bilinear map $\mathrm{Hom}(G_1, G_2) \times G_1 \to G_2$.*

Proof. Note that one can always obtain a direct-sum representation using Theorem

2.2.17. We then have that

$$\begin{aligned} G_1 \otimes G_2 &\cong \left(\bigoplus_{i=1}^{t_1} \mathbb{Z}/c_i\mathbb{Z}\right) \otimes \left(\bigoplus_{j=1}^{t_2} \mathbb{Z}/d_i\mathbb{Z}\right) \\ &\cong \bigoplus_{i,j=1}^{t_1,t_2} (\mathbb{Z}/c_i\mathbb{Z} \otimes \mathbb{Z}/d_j\mathbb{Z}) \\ &\cong \bigoplus_{i,j=1}^{t_1,t_2} \mathbb{Z}/\gcd(c_i, d_j)\mathbb{Z}, \end{aligned}$$

and all isomorphisms occurring are known and computable.

Similarly,

$$\begin{aligned} \mathrm{Hom}(G_1, G_2) &\cong \mathrm{Hom}\left(\bigoplus_{i=1}^{t_1} \mathbb{Z}/c_i\mathbb{Z}, \bigoplus_{j=1}^{t_2} \mathbb{Z}/d_i\mathbb{Z}\right) \\ &\cong \bigoplus_{i,j=1}^{t_1,t_2} \mathrm{Hom}(\mathbb{Z}/c_i\mathbb{Z}, \mathbb{Z}/d_j\mathbb{Z}) \\ &\cong \bigoplus_{i,j=1}^{t_1,t_2} \mathbb{Z}/\gcd(c_i, d_j)\mathbb{Z}, \end{aligned}$$

and all isomorphisms occurring are known and computable. □

2.8 Projective $\mathbb{Z}/m\mathbb{Z}$-modules

We have seen in Proposition 1.6.8, what the projective modules over $\mathbb{Z}/m\mathbb{Z}$ are, for $m \in \mathbb{Z}_{>0}$.

Suppose now we are given a finite abelian group A_1 and we would like to find the largest integer $m \mid \exp(A_1)$ such that A_1/mA_1 is projective as a module over $\mathbb{Z}/m\mathbb{Z}$.

Proposition 2.8.1. *There exists a deterministic polynomial-time algorithm that, given a finite abelian group A_1 computes the largest integer $m \in \mathbb{Z}_{>0}$ such that $m \mid \exp(A_1)$ and A_1/mA_1 is projective as a $\mathbb{Z}/m\mathbb{Z}$-module.*

Proof. Suppose

$$A_1 \cong \bigoplus_{\substack{p \in \mathcal{P} \\ a \in \mathbb{Z}_{>0}}} (\mathbb{Z}/p^a\mathbb{Z})^{n_{a,p}},$$

where $\mathcal{P}$ is a set of pairwise coprime integers greater than 1 and $n_{a,p} \in \mathbb{Z}_{\geq 0}$. Note that this is a situation we can reduce to if we first compute a decomposition of A_1 as a direct sum of cyclic subgroups using Theorem 2.6.10 and then apply the Coprime

Base Algorithm (Theorem 1.1.2) to the set of sizes of these cyclic components. For each $p \in \mathcal{P}$, set

$$\alpha_p = \begin{cases} \min\{a \mid n_{a,p} \neq 0\}, & \text{if } \exists a \text{ s. t. } n_{a,p} \neq 0 \\ 0, & \text{otherwise .} \end{cases}$$

Let

$$m = \prod_{p \in \mathcal{P}} p^{\alpha_p}.$$

Then

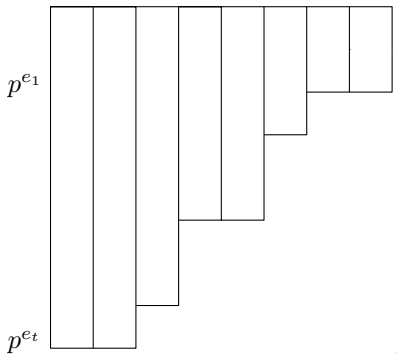
$$A_1/mA_1 \cong \bigoplus_{p \in \mathcal{P}} (\mathbb{Z}/p^{\alpha_p}\mathbb{Z})^{\sum_a n_{a,p}},$$

so is projective as a $\mathbb{Z}/m\mathbb{Z}$-module.

This value of m is the largest possible. To see this, it is enough to consider the case when $\mathcal{P} = \{p\}$ for some $p \in \mathbb{Z}_{>1}$. Suppose $m = p^{\beta_p}$, where β_p is a positive integer larger than α_p. Then A_1/mA_1 will have $\mathbb{Z}/p^{\alpha_p}\mathbb{Z}$ as a direct summand and thus cannot be projective as a $\mathbb{Z}/m\mathbb{Z}$-module.

□

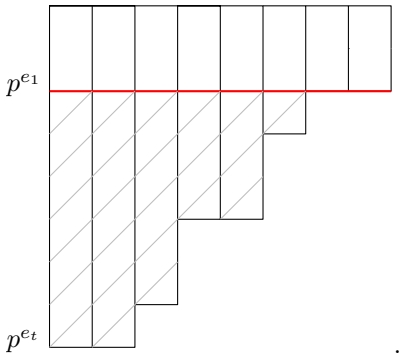
To see maximality of m graphically, let p be a prime and consider the p-group $A_1 = \bigoplus_{i=1}^{t} \mathbb{Z}/p^{e_i}\mathbb{Z}$, where $e_i \leq e_{i+1}$ and $e_1 > 0$. We represent A_1 by the following diagram:



. (2.10)

The number of vertical boxes is equal to the number of cyclic direct summands of A_1 and the height of each such box proportional to the length of the cyclic group it represents.

To make this into a free module, we need to "cut out a rectangle", so we need to cut along the smallest invariant, p^{e_1}, or along any other $p^{e'}$, for $e' \leq e_1$:



The remaining part is isomorphic to $(\mathbb{Z}/p^{e_1}\mathbb{Z})^t$ as an abelian group.

Suppose A_2 is a finite abelian group and we want to find the least integer $m' \mid \exp(A_2)$ such that $A_2/A_2[m']$ is projective over $\mathbb{Z}/\frac{\exp(A_2)}{m'}\mathbb{Z}$, where $A_2[m'] = \ker(A_2 \to A_2,\ x \mapsto m'x)$.

Proposition 2.8.2. *There exists a deterministic polynomial-time algorithm that, given a finite abelian group A_2 of exponent m, for some $m \in \mathbb{Z}_{>0}$, computes the smallest $m' \mid m$ such that $A_2/A_2[m']$ is projective over $\mathbb{Z}/\frac{m}{m'}\mathbb{Z}$.*

Proof. Suppose

$$A_2 \cong \bigoplus_{\substack{q \in \mathcal{Q} \\ b \in \mathbb{Z}_{>0}}} \left(\mathbb{Z}/q^b\mathbb{Z}\right)^{n_{b,q}}, \tag{2.11}$$

where $\mathcal{Q}$ is a set of coprime integers greater than 1 and $n_{b,q} \in \mathbb{Z}_{\geq 0}$. For each $q \in \mathcal{Q}$, let

$$\mu(q) = \begin{cases} \max\{b \mid n_{b,q} \neq 0\}, & \text{if } \exists b \text{ s.t. } n_{b,q} \neq 0 \\ 0, & \text{otherwise,} \end{cases}$$

and

$$\alpha(q) = \begin{cases} \max\{\{b \mid n_{b,q} \neq 0\} \backslash \mu(q)\}, & \text{if } \exists b \text{ s.t. } n_{b,q} \neq 0 \text{ and } b \neq \mu(q) \\ 0, & \text{otherwise.} \end{cases}$$

Then $\mu(q)$ is the largest power of q occurring as the exponent of one of the cyclic groups in the direct sum (2.11), and $\alpha(q)$ is the second largest power of q occurring. Set

$$m' = \prod_{q \in \mathcal{Q}} q^{\alpha(q)}.$$

Then

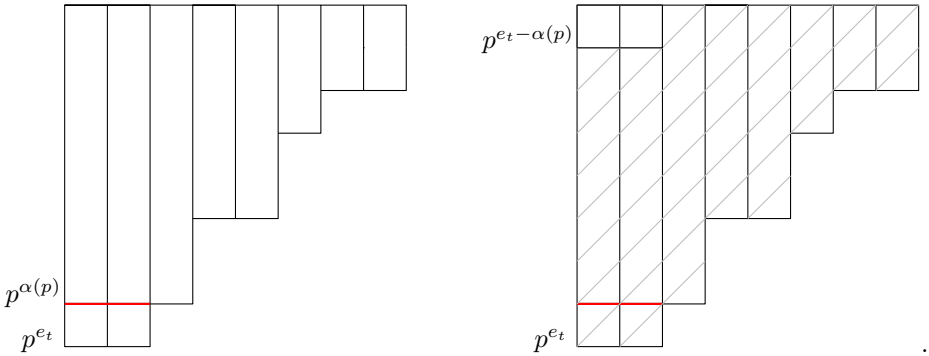
$$A_2/A_2[m'] = \bigoplus_{q \in \mathcal{Q}} \left(\mathbb{Z}/q^{\alpha(q)}\mathbb{Z}\right)^{n_{\mu(q),q}},$$

where $n_{0,q} := 0$.

This value of m' is the smallest possible. To see this, it is enough to consider the case when $\mathcal{Q} = \{p\}$, for some $p \in \mathbb{Z}_{>1}$. Suppose $m' = \beta(p)$, where $\beta(q)$ is a

positive integer smaller than $\alpha(q)$. Then $A_2/A_2[m]$ will have $\mathbb{Z}/p^{\alpha(q)-\beta(q)}\mathbb{Z}$ as a direct summand, and thus cannot be projective as a $\mathbb{Z}/\frac{m}{m'}\mathbb{Z}$-module. □

To see minimality of m' graphically, reduce again to the p-group case. Suppose $A_2 = \bigoplus_{i=1}^{t} \mathbb{Z}/p^{e_i}\mathbb{Z}$, where $e_i \leq e_{i+1}$ and $e_1 > 0$, and suppose that there are k copies of $\mathbb{Z}/p^{e_t}\mathbb{Z}$ in this sum. Then A_2 is represented by a diagram like in (2.10). This time however, to get a rectangle, we do not cut along the smallest invariant, but along the second largest one:



The remaining part is represented by the upper-left rectangle left unhatched, and is isomorphic to $(\mathbb{Z}/p^{e_t-\alpha(p)}\mathbb{Z})^k$.

Chapter 3

Linear algebra over $\mathbb{Z}$: basic algorithms for finite rings

We have seen in the previous chapter how to represent and carry out basic computations with finite abelian groups. This enables us to deal with the underlying additive group of a finite ring, which is crucial for all our algorithms. This chapter is a compendium of basic algorithms to do with finite rings.

We will represent finite rings using "basis representations" and describe algorithms that accomplish the following tasks:

1. compute homomorphism groups between two finite modules over a finite ring,
2. compute the ideal generated by a given set of elements,
3. compute sums, products and intersections of ideals,
4. compute the quotient of a finite ring by a two-sided ideal,
5. compute the characteristic, centre and prime subring of a finite ring.

We will also look at the problem of computing the Jacobson radical of a ring. In the case where the given ring is a finite-dimensional algebra over a field, this can be accomplished deterministically in polynomial time ([27, 75]). However, in general, such a result cannot be expected, in view of our inability to compute the largest square divisor of an integer.

Finally, we will briefly look at some other known algorithms for finite rings, and some open algorithmic questions.

3.1 Representing objects and basic constructions

3.1.1 Representing rings and modules

To represent finite rings inside algorithms, we will use *basis representations*. These are considered to be the "right" representations for complexity considerations, since they are neither too verbose (so as to make all problems quasipolynomial), nor too compact (so as to make all problems NP-hard). For more on the different representations and the complexity of problems on different representations, see [2, 4, 52, 53].

Definition 3.1.1. *Let R be a finite ring. A* basis representation *of R consists of a sequence of integers $d_1, \ldots, d_t \in \mathbb{Z}_{>1}$, for some $t \in \mathbb{Z}_{\geq 0}$ such that*

$$R^+ \cong \bigoplus_{i=1}^{t} (\mathbb{Z}/d_i\mathbb{Z}), \tag{3.1}$$

together with a bilinear map

$$\sigma : R^+ \times R^+ \to R^+ \tag{3.2}$$

$$(e_i, e_j) \mapsto e_i e_j, \tag{3.3}$$

where e_i is a generator of the cyclic subgroup $\mathbb{Z}/d_i\mathbb{Z}$, for $1 \leq i \leq t$, and we express $e_i e_j$ linearly in terms of $\{e_i\}$, i.e. for each $1 \leq i, j \leq t$ we give a sequence $a_{ijk} \in \mathbb{Z}/d_k\mathbb{Z}$, for $1 \leq k \leq t$, such that $e_i e_j = \sum_{k=1}^{t} a_{ijk} e_k$.

Note 3.1.2. We have established in Definition 2.3.2 that the default representation of a finite abelian group is the exact-sequence representation. However, by the results of Section 2.6.5, we may assume that R^+ is in fact given by a direct-sum decomposition into cyclic groups.

Note 3.1.3. The map σ in (3.2) will be referred to as the *multiplication map* of R. Specifying σ amounts to giving t^3 integers a_{ijk}, which are called *structure constants*.

Note 3.1.4. Given $d_1, \ldots, d_t$ and a sequence of t^3 integers, a_{ijk}, we can check in polynomial time whether they define a ring. This amounts to checking a series of equalities and solving systems of linear equations over $\mathbb{Z}$.

Note 3.1.5. The size of a basis representation is equal to

$$\sum_{i=1}^{t} \log_2(d_i) \cdot t^2 = \log_2(|R|) \cdot t^2 \leq \log_2^3(|R|), \tag{3.4}$$

since $t \leq \log_2(|R|)$. Thus, when we say an algorithm with input R runs in polynomial time, we mean that the number of bit operations is bounded above by $(\log_2(2 + |R|))^C$, for some constant C. The 2 is added in order to accommodate the zero ring.

To input a finite module, we give a finite abelian group $(M, +)$ and a bilinear map

$$\alpha : R^+ \times M \to M, \tag{3.5}$$

which describes the action of R on M. For every additive generator of R and M, we express the image in terms of the additive generators of M.

3.1.2 Representing ring and module homomorphisms

Let R_1 and R_2 be two finite rings. A ring homomorphism $\rho : R_1 \to R_2$ is a homomorphism of the underlying abelian groups that sends the unit element of R_1 to the unit element of R_2 and respects the multiplicative structure of the rings, i.e. for all $r, s \in R_1$, we have that

$$\rho(rs) = \rho(r)\rho(s). \tag{3.6}$$

Proposition 3.1.6. *There exists a deterministic polynomial-time algorithm that, given two finite rings R_1 and R_2, and a group homomorphism $\rho : R_1^+ \to R_2^+$, decides whether ρ is a ring homomorphism, and if it is, decides whether it is injective or surjective.*

Proof. As in Section 2.3.2, the map ϕ is given by a matrix which specifies the image of each additive generator of R_1 as a linear combination of additive generators of R_2. Given such a matrix, we can easily check if it induces a homomorphism of rings, by verifying that the induced map preserves multiplication and sends the unit element of one ring to the unit element of the other. This amounts to checking equalities over $\mathbb{Z}$. By Propositions 2.3.6 and 2.3.7, we can also check for injectivity or surjectivity. □

Let R be a ring and M, N two R-modules. A module homomorphism $\phi : M \to N$ is a homomorphism of abelian groups which is R-linear, i.e. for all $r \in R, m \in M$, we have that

$$\phi(rm) = r\phi(m). \tag{3.7}$$

Proposition 3.1.7. *There exists a deterministic polynomial-time algorithm that, given a finite ring R, two R-modules M and N, and a group homomorphism $\phi : M \to N$, decides whether ϕ is an R-module homomorphism, and if it is, decides whether it is injective or surjective.*

Proof. As in Chapter 2, Section 2.3.2, the map ϕ is given by a matrix which specifies the image of each additive generator of M as a linear combination of additive generators of N. Given such a matrix, we can easily check if it induces a module homomorphism by verifying that the induced map preserves scalar multiplication. This amounts to checking equalities over $\mathbb{Z}$. By Propositions 2.3.6 and 2.3.7, we can also check for injectivity or surjectivity. □

3.1.3 Homomorphism group

Proposition 3.1.8. *There exists a deterministic polynomial-time algorithm that, given a finite ring R and two finite R-modules M_1 and M_2, computes the homomorphism group* $\mathrm{Hom}_R(M_1, M_2)$.

Proof. We can certainly compute $\mathrm{Hom}_{\mathbb{Z}}(M_1, M_2)$. Then

$$\mathrm{Hom}_R(M_1, M_2) = \{f \in \mathrm{Hom}_{\mathbb{Z}}(M_1, M_2) \mid f(rx) = rf(x),\ \forall r \in R, \forall x \in M_1\}. \quad (3.8)$$

It is enough to ensure that the relation $f(rx) = rf(x)$ holds for the additive generators of R and M_1. Consider the map

$$\mathrm{Hom}_{\mathbb{Z}}(M_1, M_2) \to \bigoplus_{\substack{r \text{ additive generator of } R \\ x \text{ additive generator of } M_1}} M_2$$
$$f \mapsto (f(rx) - rf(x))_{r,x}.$$

Then $\mathrm{Hom}_R(M_1, M_2)$ is the kernel of this map, which we can compute. □

3.2 Computations with ideals

Let R be a ring. A left ideal I is, in particular, a left R-module, so it is given to the algorithm as an additive subgroup of R, together with a map $R \times I \to I$, as in (3.5). Right ideals and two-sided ideals are given in a similar way.

3.2.1 Computing the ideal generated by a given set of elements

Proposition 3.2.1. *There exists a deterministic polynomial-time algorithm that, given a finite ring R and a set $S \subset R$, computes the ideal generated by S in R.*

Proof. Suppose $S = \{s_1, \ldots, s_u\}$. The left ideal I, generated by S in R has underlying additive group generated by the set $\{e_i s_j \mid 1 \le i \le t, 1 \le j \le k\}$, which can be computed using Proposition 2.3.10. This will also produce an injective group homomorphism $I \hookrightarrow R$ specifying I^+ as a subgroup of R^+. To determine the R-action, we look at the map $\sigma : R^+ \times R^+ \to R^+$ as in (3.2), giving the multiplicative structure of R. Suppose $e_i s_j = \sum_{n=1}^{t} b_{ijn} e_n$. Then

$$e_k e_i s_j = \left(\sum_m a_{kim} e_m\right) s_j = \sum_{m,n} a_{kim} b_{mjn} e_n. \quad (3.9)$$

Now we express this sum as a linear combination of the $e_i s_j$. □

The right and two-sided ideals are dealt with in a similar manner.

3.2.2 Sum, product and intersection of ideals

Proposition 3.2.2. *There exists a deterministic polynomial-time algorithm that, given a finite ring R and two ideals $I, J \subset R$, computes the ideals $I+J, I \cap J$ and IJ.*

Proof. Note that $(I+J)^+ = I^+ + J^+$ and $(I \cap J)^+ = I^+ \cap J^+$, and the action of R is induced by the ring multiplication map σ (as in (3.2)).

Suppose $x_1, \ldots, x_n$ is a set of additive generators of I and $y_1, \ldots, y_m$ is a set of additive generators of J. Recall that $IJ \subseteq J$. We have that

$$(IJ)^+ = \langle \{x_i \cdot y_j\}_{i,j} \rangle_{\mathbb{Z}}, \tag{3.10}$$

where the product $x_i \cdot y_j$ is computed by writing each x_i and y_j in terms of additive generators of R and then using the multiplication map σ to write each $x_i \cdot y_j$ as a linear combination of additive generators of R. The action of R on $(IJ)^+$ is again induced by σ. □

The right and two-sided ideals are dealt with in a similar manner.

3.2.3 Quotient of ring and two-sided ideal

Proposition 3.2.3. *There exists a deterministic polynomial-time algorithm that, given a finite ring R and a two-sided ideal I, computes the quotient ring R/I.*

Proof. Let I be a two-sided ideal of R. Then

$$(R/I)^+ = R^+/I^+, \tag{3.11}$$

which can be computed using Proposition 2.3.13. This will also produce a surjective group homomorphism $R^+ \twoheadrightarrow (R/I)^+$. The multiplication map for R/I is induced by the multiplication map for R. □

3.3 Computing the centre and the prime subring of a finite ring

Given a finite ring, we will often want to view it as an algebra over its centre or its prime subring. We show how to compute these.

Theorem 3.3.1. *There exists a deterministic polynomial-time algorithm that, given a finite ring R, computes its centre, prime subring and characteristic.*

Proof. To compute the centre of a finite ring R, we consider the map

$$\begin{aligned} \psi : R^+ &\longrightarrow \bigoplus_{r \text{ additive generator of } R} R^+ \\ s &\longmapsto (rs - sr)_r. \end{aligned}$$

The centre of R is the kernel of ψ, which we can compute. Its ring structure is induced by σ, the map defining multiplication in R.

To compute the prime subring of a finite ring R, we consider the map

$$\alpha : \mathbb{Z} \to R^+, \quad 1 \mapsto 1_R.$$

The prime subring of R is the image of α, which we can compute. The multiplication map is induced by σ.

The cardinality of the prime subring is the characteristic of R, which we can also compute by taking the lowest common multiple of the sizes of the cyclic direct summands of R^+. □

3.4 Computing the Jacobson radical

When dealing with problems concerning rings, it is often convenient to reduce to the semisimple case. When the ring at hand is left-artinian, this reduces to computing the Jacobson radical (see Theorem 1.4.9, part (iv)). Thus, the natural question to ask is if Jacobson radicals can be efficiently computed. For our purposes, the appropriate version of this question is whether the Jacobson radical of a finite ring can be computed deterministically in polynomial time.

If the given ring R is a finite-dimensional algebra over a "nice" field $\mathbb{F}$, there do exist deterministic polynomial-time algorithms that compute the Jacobson radical of R (see [18, 27, 75]). For $\mathbb{F}$ a field of characteristic 0, this reduces to solving a system of linear equations over $\mathbb{F}$ by Dickson's theorem (see Theorem 1.4.10). If $\mathbb{F}$ is a finite field, then Friedl and Rónyai showed how to recursively construct a sequence of ideals of R, whose last element is equal to $\mathrm{J}(R)$. Using a very similar technique, Cohen, Ivanyos and Wales generalised these results, showing how to compute the Jacobson radical in the case that $\mathbb{F}$ is any field in which one can perform arithmetic and over which one can solve *semilinear equations* of the form $\sum_{i=1}^{k} a_i x_i^p = 0$, where $p = \mathrm{char}(\mathbb{F})$, $k \in \mathbb{Z}_{>0}$ and $a_i \in \mathbb{F}$ for all $1 \leq i \leq k$. Finite fields are examples of such fields.

Theorem 3.4.1 ([18])**.** *There exists a deterministic polynomial-time algorithm that, given a finite-dimensional algebra R over a field $\mathbb{F}$, where $\mathbb{F}$ is a field over which we can perform arithmetic and solve semilinear equations, computes the Jacobson radical of R.*

Note 3.4.2. We cannot in general expect to be able to compute the Jacobson radical for rings not containing a field. To see this, consider rings of the form $\mathbb{Z}/n\mathbb{Z}$, with $n \in \mathbb{Z}_{>0}$, for which the task ultimately reduces to finding square divisors of n. This is not something we know how to do deterministically in polynomial time.

3.5 Other known algorithms and open questions

Rings are ubiquitous. It is thus important to have a wide range of algorithms to deal with finite rings. This list of deterministic polynomial-time algorithms for finite rings

however, is not as long as would be expected for such basic objects. One of the reasons for this is that many problems for finite rings reduce to rings of the type $\mathbb{Z}/n\mathbb{Z}$, and at this stage the fact that we cannot factor n efficiently becomes a serious issue.

The study of algorithmic problems involving automorphisms and isomorphisms of finite rings intensified after the first deterministic polynomial-time primality test was formulated by Agrawal, Saxena and Kayal in terms of automorphisms of a certain finite ring (see [1]). Subsequently, the same authors studied how some of the most important open algorithmic questions, like integer factorisation, polynomial factorisation over finite fields and graph isomorphism can be reduced to ring automorphism questions. The questions for which deterministic polynomial-time algorithms are sought are the following:

1. Ring Isomorphism Problem
   - (i) Decision version: Given two finite rings, decide if they are isomorphic.
   - (ii) Search version: Given two finite rings, find an isomorphism if one exists.
   - (iii) Counting version: Given two finite rings, compute the number of isomorphisms between them.
2. Ring Automorphism Problem
   - (i) Decision version: Given a finite ring, decide if it has a nontrivial ring automorphism.
   - (ii) Search version: Given a finite ring, find a nontrivial automorphism if one exists.
   - (iii) Counting version: Given a finite ring, compute the number of its automorphisms.

As far as deterministic polynomial-time algorithms are concerned, all of the above problems are open, with the exception of the decision version of the ring automorphism problem, which was shown to be in P in [53]. The algorithm given there relies on the classification of rigid rings (i.e. rings with no nontrivial automorphisms).

Theorem 3.5.1 ([53], Theorem 7.1)**.** *There exists a deterministic polynomial-time algorithm that, given a finite ring R, determines whether R has a nontrivial automorphism.*

It is shown in [53] that both integer factorisation and the graph isomorphism problem reduce to the counting version of the ring automorphism problem, which is unlikely to be NP-complete. The decision version of the ring isomorphism problem is shown to be at least as hard as the graph isomorphism problem. Moreover, integer factorisation reduces to the search version of the ring isomorphism problem.

For finite fields, the isomorphism problem can be handled deterministically in polynomial time.

Theorem 3.5.2 ([65], Theorem 1.2)**.** *There exists a deterministic polynomial-time algorithm that, given two finite fields of the same cardinality, exhibits an isomorphism between them.*

Apart from these problems, to which systematic study has been devoted, the list of algorithms for finite rings remains quite a short one. It is one of the goals of this thesis to expand on this list, thus supplementing the toolbox for algorithmically dealing with finite rings.

Chapter 4

The module isomorphism problem

Let R be a finite ring and let M_1, M_2 be two finite left R-modules. We present two distinct deterministic algorithms that decide in polynomial time whether or not M_1 and M_2 are isomorphic, and if they are, exhibit an isomorphism. As by-products, we are able to determine the largest isomorphic common direct summand between two modules and the minimum number of generators of a module. By not requiring R to contain a field, avoiding computation of the Jacobson radical and not distinguishing between large and small characteristic, both algorithms constitute improvements to known results. We have not attempted to implement either of the two algorithms, but we have no reason to believe that they would not perform well in practice.

4.1 Introduction

The *module isomorphism problem* (MIP) can be formulated as follows: design a deterministic algorithm that, given a ring R and two left R-modules M_1 and M_2, decides in polynomial time whether they are isomorphic, and if yes, exhibits an isomorphism.

This problem is as fundamental as it is easily stated and has been studied extensively, due both to its intrinsic theoretical value and to its broad range of applications. As a theoretical question it is one in a long series of isomorphism problems. These are some of the most natural questions that occur in algorithmic contexts: given two objects "of the same nature", one wishes to determine if they are equal, or isomorphic. Examples of these problems include the group isomorphism problem, the graph isomorphism problem and the ring isomorphism (see [52, 53]). From a complexity point of view, these problems have a special status, namely, they are thought to be NP-intermediate, i.e. pertaining to the class consisting of problems that are known

This chapter is an extended version of the paper *The module isomorphism problem for finite rings and related results*, arXiv:1512.08365v1 ([17])

to be in NP, but are not known to be in P, or to be NP-complete. The class of NP-intermediate problems is nonempty if and only if P $\neq$ NP (see [73], Theorem 14.1). However, even under the hypothesis that P $\neq$ NP, no "natural" NP-intermediate problems are known.

The practical value of the module isomorphism problem comes from viewing it in the larger context of algorithms for finite rings. Finite rings are fundamental objects and one wishes to have as many algorithms as possible for handling them at one's disposal, that is to say, in one's computer algebra system toolbox.

A brief overview of the results related to the module isomorphism problem is included in Section 4.2. In particular, polynomial-time algorithms for the module isomorphism problem were given in [10, 15] for the case where R is a finite-dimensional algebra over a field and M_1, M_2 are finite-dimensional modules over that field. For our purposes, R will be a finite ring (not necessarily containing a field) and M_1, M_2 will be finite R-modules. We give an algorithm as described by the following theorem:

Theorem 4.1.1. *There exists a deterministic polynomial-time algorithm that, given a finite ring R and two finite R-modules M_1 and M_2, computes a maximum length R-module C that is isomorphic to a direct summand both of M_1 and of M_2. Moreover, the algorithm computes direct complements of C both in M_1 and in M_2, together with the corresponding isomorphisms.*

We establish the result in Theorem 4.1.1 by a direct generalisation of the methods given in [10], where the rings considered were finitely generated algebras over a field and the modules were finite-dimensional over that field. This approach relies on the ability of finding non-nilpotent elements in non-nilpotent ideals of the endomorphism ring $\mathrm{End}_R(M_1)$ (cf. Proposition 4.3.3). As a direct consequence of Theorem 4.1.1 we have that:

Theorem 4.1.2. *There exists a deterministic polynomial-time algorithm that, given a finite ring R and two finite R-modules M_1 and M_2, decides whether M_1 and M_2 are isomorphic, and if they are, exhibits an isomorphism.*

In Section 3, we show how the module isomorphism problem reduces to determining freeness of rank one of a module. We then give an algorithm that computes the minimum number of generators of a given finite module.

Theorem 4.1.3. *There exists a deterministic polynomial-time algorithm that, given a finite ring R and a finite R-module M, computes a set for generators of M of minimum cardinality.*

In particular, we can determine if a module is cyclic or free (by comparing cardinalities), which gives a second deterministic polynomial-time algorithm for the module isomorphism problem.

It is important to note that the algorithms given here work for any finite ring, do not distinguish between large and small characteristic and avoid computation of the

Jacobson radical, thus constituting an improvement to known results. Moreover, the algorithm in Theorem 4.1.3 is an interesting object *per se*, due to its structure and the techniques it employs. A common approach to this type of problems is to reduce to the semisimple case and then "lift" (e.g. [15, 42]). In our algorithm, we work *as if* the ring were semisimple and we have a list, $S_1, \ldots, S_t$, of candidates for the isomorphism classes of simple modules composing it. During the running of the algorithm, we allow ourselves to be contradicted in our assumption about the simplicity of the S_i, in which case we update our list, quotient the ring by an appropriate two-sided nilpotent ideal and start again. If we are not contradicted, we may still draw conclusions. In this way, there is always a side-exit available and what forces an output in polynomial time is that we cannot take the side-exit too many times.

4.2 Context

All of the results referred to in this section deal with the case where the given ring R is a finite-dimensional algebra over a field $\mathbb{F}$ and the two modules are finite-dimensional over that field.

One of the first attempts at tackling the module isomorphism problem made use of the MeatAxe algorithm due to Parker (see [74]). Holt and Rees produce an efficient randomized algorithm in [36], that in the case where $\mathbb{F}$ is finite, tests whether a module is simple, and if it is not, produces a proper submodule. This algorithm is then used to test for isomorphism between two modules of which one is known to be simple. However, their method fails in the cases where the given module has a very special structure.

The first deterministic polynomial-time algorithm for the module isomorphism problem was presented by Chistov, Ivanyos and Karpinski in [15]. For the purpose of their paper, $\mathbb{F}$ is a finite field or an algebraic number field. Their method reduces the problem to the semisimple case by computing the Jacobson radical of R and then further reduces the problem to testing cyclicity of modules. Because this method requires the computation of the Jacobson radical, it does not lend itself to generalisation in the case of general finite rings (cf. Chapter 3, Section 3.4).

The next solution to the module isomorphism problem was given by Brooksbank and Luks in [10]. Their algorithm is based on finding nilpotent endomorphisms of one of the modules and in fact does more than test for module isomorphism – it produces the largest isomorphic common direct summand between two modules, i.e. the maximal length direct summand of one of the modules, which has an isomorphic copy as a direct summand of the other. In principle, this method works for any underlying field, as long as we can perform arithmetic in it.

Another related result was presented by Ivanyos, Karpinski and Saxena in [42]. Their algorithm computes the minimum number of generators of a given module and is in particular useful for testing module cyclicity. The drawback of this method is that it distinguishes between small and large characteristic of the underlying field.

Our goal is to obtain corresponding algorithms for the case that R is a finite ring

(not necessarily containing a field) and the given modules are also finite.

4.3 MIP via non-nilpotent endomorphisms

In this section, the algorithm of Brooksbank and Luks [10] is generalised to solve the module isomorphism problem for finite rings and finite modules over that ring.

4.3.1 Finding non-nilpotent elements in non-nilpotent ideals

An ideal I of a ring R is said to be *nil* if all its elements are nilpotent. The following is a well known fact.

Proposition 4.3.1. *If R is a left-artinian ring, then a left (or right) ideal I of R is nil if and only if it is nilpotent.*

Proof. Clearly if a left (or right) ideal is nilpotent, then it is nil. For the other direction, note that any nil left (or right) ideal of R is contained in the Jacobson radical $\mathrm{J}(R)$ (see [57], Lemma 4.11). So it is enough to show that $J := \mathrm{J}(R)$ is nilpotent.

Suppose J is not nilpotent. Since R is left-artinian, the sequence of consecutive powers of J must stabilize, i.e. $\exists n \in \mathbb{Z}_{>0}$ such that $J^n = J^{2n} \neq 0$. Then there exists a left ideal L s.t. $J^n L \neq 0$. Suppose L is minimal with this property. Since $J^n L \neq 0$, there exists $a \in L$ such that $J^n a \neq 0$. Now $Ra \subseteq L$ and $J^n a \subseteq L$. Moreover, $0 \neq J^n a = J^n J^n a$ and $0 \neq J^n a \subseteq J^n Ra$, so by minimality of L, we have that $L = Ra = J^n a$. Hence $a = xa$, for some $x \in J^n$. It follows that $a = 0$ (otherwise, $1 - x$ would be both a unit and a left zero-divisor, which is not possible). This gives the desired contradiction. □

Note 4.3.2. The "only if" direction in the statement of the previous proposition is not true if we remove the artinianity condition. Consider the commutative ring $R = \mathbb{Z}[x_1, x_2, \ldots]/(x_1^2, x_2^2, \ldots)$. Then the ideal $I = (x_1, x_2, \ldots)$ is nil, but I is not nilpotent.

If R is finite, the proof of Proposition 4.3.1 can be translated into an algorithm for finding a non-nilpotent element in a non-nilpotent left ideal of R.

Proposition 4.3.3. *There exists a deterministic polynomial-time algorithm that, given a finite ring R and a left ideal I, determines whether or not I is nilpotent and if it is not, produces a non-nilpotent element lying inside it.*

Proof. Suppose R is a k-algebra, where k is a commutative ring (R will certainly be an algebra over its centre or its prime subring) and let I be a non-nilpotent left ideal of R. Since R is finite, we can find $n \in \mathbb{Z}_{>0}$ such that $I^n = I^{2n} \neq 0$. Suppose I^n is generated over k by a set A and over R by a set B, i.e.

$$I^n = \sum_{\alpha \in A} k\alpha = \sum_{\beta \in B} R\beta.$$

Then $0 \neq I^{2n} = \sum_{\substack{\alpha \in A \\ \beta \in B}} R\beta\alpha$, so there exists $b \in B$ and $a \in A$ such that $b \cdot a \neq 0$. Consider the ideals $I^n a \subseteq Ra$, where $I^n a \neq 0$ since it contains ba. If this inclusion is in fact an equality, then we can write $a = xa$, for some $x \in I^n$ and since $a \neq 0$, it must be the case that x is non-nilpotent, as required (otherwise $1 - x$ would be a unit and a left zero divisor at the same time, which is impossible). Suppose now that the inclusion is strict. Then there exists $c \in I^n a$ such that $I^n c \neq 0$ (otherwise $I^n I^n a = I^n a = 0$, which is a contradiction, since $0 \neq ba \in I^n a$). Moreover, we will be able to find such a c among k-generators of $I^n a$, otherwise, for all $\alpha \in A$, we would have that $I^n \alpha a = 0$, and so $0 = \sum_{\alpha \in A} k I^n \alpha a = I^{2n} a = I^n a \neq 0$, which is a contradiction.

We have now produced a smaller ideal $Ra \supsetneq I^n a \supset Rc$ such that $I^n c \neq 0$. We replace a by c and keep going. This process must terminate in polynomial time, since $\text{length}(Ra) \leq \text{length}({}_R R) \leq \log_2(|R|)$, and so the length of the descending chain of ideals obtained, and hence the number of steps performed by the algorithm, is bounded above by $\log_2(|R|)$. We also note that I is nilpotent if and only if $I^{\text{length}({}_R R)} = 0$, so it is possible to test if I is nilpotent in polynomial time. □

4.3.2 Splitters

Let R be an algebra over a commutative artinian ring k such that R is finitely generated as a k-module. Let M_1 and M_2 be two R-modules of finite length over k. Note that $\text{Hom}_R(M_1, M_2)$ is a left k-module via the action

$$\begin{aligned} k \times \text{Hom}_R(M_1, M_2) &\to \text{Hom}_R(M_1, M_2) \\ (r, f) &\mapsto (m \mapsto f(mr)), \end{aligned}$$

and a right $\text{End}_R(M_1)$-module via the action

$$\begin{aligned} \text{Hom}_R(M_1, M_2) \times \text{End}_R(M_1) &\to \text{Hom}_R(M_1, M_2) \\ (f, g) &\mapsto f \circ g. \end{aligned}$$

Moreover, $\text{End}_R(M_1)$ is left (and right) artinian. In particular, all its nil ideals are nilpotent and *vice versa*.

Following [10], we make the following definition:

Definition 4.3.4. *Let $f \in \text{Hom}_R(M_1, M_2)$. A decomposition $M_1 = N_1 \oplus K_1$, for $N_1, K_1 \leq M_1$, is called an* f-decomposition *if $N_1 \neq 0$, $\ker(f) \leq K_1$ and the image of N_1 under f, which we denote by fN_1, is a direct summand of M_2. If M_1 has an f-decomposition, we say that f is a* splitter.

Note that the condition $\ker(f) \leq K_1$ implies that $N_1 \cong fN_1 = N_2$, so f induces an isomorphism $N_1 \cong N_2$.

The following proposition and its proof, together with Proposition 4.3.3, allow us to algorithmically decide if a given homomorphism f is a splitter, and if it is, to produce an f-decomposition.

Proposition 4.3.5. *Let $f \in \mathrm{Hom}_R(M_1, M_2)$. Then f is a splitter if and only if there exists $g \in \mathrm{Hom}_R(M_2, M_1)$ such that gf is not nilpotent.*

Proof. The proof of this proposition is the same as the one given in [10], Lemma 3.3, which treats the case when R is a finitely generated algebra over a field and M_1, M_2 are finite-dimensional over that field. For completeness we include the proof of the "if" statement here.

Suppose $g \in \mathrm{Hom}_R(M_2, M_1)$ is such that gf is not nilpotent. Let $s = gf$ and $t = fg$. Since M_1 and M_2 have finite length over k (so are both artinian and noetherian over R), we can apply Fitting's Lemma (see Proposition 1.3.12) to say that $M_1 = \ker(s^d) \oplus \mathrm{im}(s^d)$, $M_2 = \ker(t^d) \oplus \mathrm{im}(t^d)$, for $d = \max\{\mathrm{length}_R(M_1), \mathrm{length}_R(M_2)\}$ and the restriction of t to $\mathrm{im}(t^d)$ is an automorphism. We have that $s^d M_1 \neq 0$ since s is not nilpotent, $\ker(f) \leq \ker(s^d)$ by definition of s and

$$f(s^d M_1) = t^d(fM_1) \subseteq t^d M_2 = fg(t^d M_2) = fs^d(gM_2) \subseteq f(s^d M_1),$$

so $f\,\mathrm{im}(s^d) = f(s^d M_1) = t^d M_2 = \mathrm{im}(t^d)$, which is what is required for f to be a splitter. □

Suppose now that we are given $f \in \mathrm{Hom}_R(M_1, M_2)$ and we wish to know whether it is a splitter. To do this, we first compute a set C of k-generators of $\mathrm{Hom}_R(M_2, M_1)$. Now consider the left ideal I of $\mathrm{End}_R(M_1)$ generated by the set $Cf = \{cf \mid c \in C\}$. If I is nilpotent (which we can determine by Proposition 4.3.3), then since $\mathrm{Hom}_R(M_2, M_1)f = \mathrm{span}_k(Cf) = I$, we will not be able to find a non-nilpotent element of the form gf, so f cannot be a splitter. Otherwise the algorithm of Proposition 4.3.3 will produce some $g \in \mathrm{Hom}_R(M_2, M_1)$ witnessing that f is a splitter and we can produce an f-decomposition by Proposition 4.3.5.

We now have a way of identifying splitters. But we would not like to have to look for them over all homomorphisms. The following proposition tells us that we can restrict our attention to a considerably smaller set.

Proposition 4.3.6. *If a splitter exists, then there exists a splitter in any set of k-module generators of $\mathrm{Hom}_R(M_1, M_2)$ and in any set of $\mathrm{End}_R(M_1)$-generators of $\mathrm{Hom}_R(M_1, M_2)$.*

Proof. To see that a set of $\mathrm{End}_R(M_1)$-module generators is enough, note that

$$f \text{ is not a splitter} \iff \mathrm{Hom}_R(M_2, M_1)f \subseteq \mathrm{J}(\mathrm{End}_R(M_1)).$$

Let B be a set of $\mathrm{End}_R(M_1)$-module generators of $\mathrm{Hom}_R(M_1, M_2)$. Suppose that B does not contain any splitters. Then

$$\begin{aligned}\mathrm{Hom}_R(M_2, M_1)\,\mathrm{Hom}_R(M_1, M_2) &= \sum_{b \in B} \mathrm{Hom}_R(M_2, M_1) b\,\mathrm{End}_R(M_1) \\ &\subseteq \mathrm{J}(\mathrm{End}_R(M_1)),\end{aligned}$$

and therefore $\mathrm{Hom}_R(M_1, M_2)$ cannot contain a splitter.

Finally, note that any set of left-k-module generators is also a set of right-$\mathrm{End}_R(M_1)$-module generators. □

Putting these results together, we can construct an algorithm satisfying the requirements of Theorem 4.1.1 as follows.

Proof of Theorem 4.1.1: We view R as an algebra over its prime subring k. Let C be an auxiliary variable that at the end of the algorithm will become equal to the desired maximal length R-module that is isomorphic to a direct summand both of M_1 and of M_2. At the beginning of the algorithm, we put C equal to zero. Similarly, we define an auxiliary variable $f \in \mathrm{Hom}_R(M_1, M_2)$ and initialise it at zero.

We compute B, a set of k-generators of $\mathrm{Hom}_R(M_1, M_2)$ (or a set of $\mathrm{End}_R(M_1)$-module generators thereof). For each element of B, we test if it is a splitter: by Proposition 4.3.6, if a splitter exists, we will find one inside B. Finding a splitter $b \in B$ also gives us a homomorphism $c \in \mathrm{Hom}_R(M_2, M_1)$ such that cb is not nilpotent, and a decomposition $M_1 = N_1 \oplus K_1$ and $M_2 = N_2 \oplus K_2$, where $N_1 = \mathrm{im}((cb)^d)$, $K_1 = \ker((cb)^d)$, $N_2 = \mathrm{im}((bc)^d)$, $K_2 = \ker((bc)^d)$, and $d = \max\{\mathrm{length}_R(M_1), \mathrm{length}_R(M_2)\}$ (see Proposition 4.3.5). We make the following replacements: $C := C \oplus N_1$ and $f := f \oplus$ (the restriction of b to N_1), $M_1 := K_1$, $M_2 := K_2$, $B :=$ (set of k-module generators of the new $\mathrm{Hom}_R(M_1, M_2)$), and we repeat the process. Note that we are all the time assuming the Krull-Remak-Schmidt Theorem (see Theorem 1.3.13), which ensures existence and uniqueness up to isomorphism of the direct summands of M_1 and M_2.

The algorithm produces the following data: the R-module C and the R-module homomorphism f such that $f : C \xrightarrow{\sim} f(C)$. Put $D := f(C)$. The algorithm also produces injections of C and D into M_1 and M_2 respectively, that define them as submodules. By splitting these injections (see Proposition 2.5.1), we can recover the direct complements of C and D in M_1 and M_2 respectively, together with the corresponding isomorphisms. □

4.4 MIP via an approximation of the Jacobson radical

In this section, another solution to the module isomorphism problem for finite rings is presented. We start with a problem reduction, showing that it is enough to be able to test if a module is free of rank 1. We then construct an algorithm that, given a finite ring and a finite module over that ring, computes a set of generators of minimum cardinality for that module.

4.4.1 Problem reduction

Let R be a finite ring. We observe that determining whether two finite R-modules M_1 and M_2 are isomorphic reduces to determining if a module is free of rank one. Let $E := \mathrm{End}_R(M_1)$ and $K := \mathrm{Hom}_R(M_2, M_1)$. Note that K is a left E-module.

Proposition 4.4.1. *Let R, M_1, M_2, E, K be as above. The following are equivalent:*

(i) *$M_1 \cong M_2$ as R-modules.*

(ii) *$E \cong K$ as E-modules and the image of 1_E in K under any such isomorphism is an isomorphism between M_1 and M_2.*
(iii) *There exists an E-module isomorphism $\phi : E \to K$ such that $\phi(1_E)$ is an isomorphism between M_1 and M_2.*

Proof. The implications (ii)⇒(iii) and (iii)⇒(i) are immediate. For (i)⇒(ii), suppose $f : M_2 \xrightarrow{\sim} M_1$. Then $E \cong K$, so let $\phi : E \xrightarrow{\sim} K$ be any such isomorphism. Let $\lambda := \phi(1)$. Then there exists a unique $\epsilon \in E$ such that

$$f = \phi(\epsilon) = \phi(\epsilon \cdot 1) = \epsilon\phi(1) = \epsilon\lambda,$$

where the third equality follows by E-linearity of ϕ. Since f is injective, so must λ be. Moreover, since M_1 and M_2 have the same length, λ must also be surjective. □

Hence an algorithm that can establish freeness of rank one of a module provides a solution to the module isomorphism problem in the following way. Given M_1 and M_2, we might first test whether they have the same cardinality, otherwise it is pointless to proceed. If the cardinalities do agree, we compute E and K, and test if $E \cong K$ as E-modules (in the case that $E \cong K$, we are assuming this test will also produce an isomorphism). If this is not the case, we conclude that $M_1 \ncong M_2$. Otherwise, suppose we have found an isomorphism $\psi : E \xrightarrow{\sim} K$. Set $\lambda := \psi(1_E)$. Then by Proposition 4.4.1, we have that $M_1 \cong M_2$ if and only if λ is an isomorphism.

For the remainder of this section we will concentrate on the task of computing the minimum number of generators of a given module. In particular, this will allow us to test for cyclicity and, by comparing cardinalities, for freeness of rank 1.

4.4.2 Computing minimum number of generators

Let R be a left-artinian ring. Suppose, along with R, we are given a collection $S_1, \ldots, S_t$ of nonzero left R-modules of finite length. Ideally, we would like this collection to be a set of representatives for the isomorphism classes of simple R-modules. However, for now, we only require that each simple R-module occurs in at least one of the S_i, i.e. it occurs as a quotient in its composition series. We can take, for example, $t = 1$ and $S_1 = {}_RR$.

Let

$$\mathfrak{a} = \bigcap_{i=1}^{t} \operatorname{ann}_R(S_i),$$

where $\operatorname{ann}_R(S_i) = \ker(R \to \operatorname{End}_{\mathbb{Z}}(S_i))$. Again, ideally we would like $\mathfrak{a}$ to be the Jacobson radical, $\mathrm{J}(R)$. In reality, we only have one inclusion, namely $\mathfrak{a} \subseteq \mathrm{J}(R)$, since an element of $\mathfrak{a}$ will kill all the S_i's and so will kill all submodules, quotients and submodules of quotients of the S_i. Since every simple R-module occurs in at least one of the S_i, we have that $\mathfrak{a}$ kills all simple R-modules, which is equivalent to being inside the Jacobson radical. Furthermore, since R is left-artinian, $\mathrm{J}(R)$ is nilpotent and hence $\mathfrak{a}$ is nilpotent.

We will construct an algorithm that, given R and a collection of S_i, either "improves" the sequence of S_i or computes $\mathfrak{a}$ and an isomorphism of R-modules

$$R/\mathfrak{a} \xrightarrow{\sim} \bigoplus_{i=1}^{t} S_i^{a_i}, \tag{4.1}$$

for suitable $a_i \in \mathbb{Z}_{\geq 0}$. Similarly, we construct an algorithm that, when also given a finitely generated R-module M, either "improves" the sequence of S_i or computes an isomorphism of R-modules

$$M/\mathfrak{a}M \xrightarrow{\sim} \bigoplus_{i=1}^{t} S_i^{c_i},$$

for suitable $c_i \in \mathbb{Z}_{\geq 0}$. Consider the quantity $l(S_i)_{i=1}^t = \sum_{i=1}^t (2\,\mathrm{length}(S_i) - 1) \in \mathbb{Z}_{\geq 0}$. An "improvement" in the sequence is measured by a decrease in $l(S_i)_{i=1}^t$ and occurs either when we remove one of the S_i from the list (either because the list already contains an isomorphic copy of it or because it is not needed) or when we discover a nonzero proper submodule T of one of the S_i which witnesses nonsimplicity of S_i and which we use to replace S_i by S_i/T and T, that now stand a better chance of being simple. Note that the factor 2 in the expression of $l(S_i)_{i=1}^t$ ensures it decreases even when we remove an S_i from the list.

Let us first write out the details of the routine that finds an isomorphism as in (4.1). We will call this routine UPDATE and within the main algorithm we will call it whenever we have improved on our sequence of S_i's and we want to update our $\mathfrak{a}$, the sequence of a_i's and the isomorphism $R/\mathfrak{a} \xrightarrow{\sim} \bigoplus_{i=1}^t S_i^{a_i}$.

Let $\mathfrak{S} = \{(S_i)_{i=1}^t \mid t \in \mathbb{Z}_{\geq 0}$, each S_i is a nonzero finite length R-module and each simple R-module occurs as a factor in the composition series of at least one $S_i\}$.

Proposition 4.4.2. *There exists a deterministic polynomial-time algorithm that takes as input a finite ring R and a collection of modules $(S_i)_{i=1}^t \in \mathfrak{S}$ and outputs a sequence of integers $a_1, \ldots, a_t \in \mathbb{Z}_{>0}$, a two-sided nilpotent ideal $\mathfrak{a}'$ of R and an isomorphism $\varphi : R/\mathfrak{a}' \xrightarrow{\sim} \bigoplus_{i=1}^{t'} (S_i')^{a_i}$, where $t' \in \mathbb{Z}_{\geq 0}$, $\mathfrak{a}' = \bigcap_{i=1}^{t'} \mathrm{ann}_R(S_i')$ and $(S_i')_{i=1}^{t'} \in \mathfrak{S}$ is such that $l(S_i')_{i=1}^{t'} \leq l(S_i)_{i=1}^t$.*

Proof. Let $\mathfrak{b}$ be a left ideal of R which we will use as an intermediate variable that at the end of the algorithm will become equal to the desired $\mathfrak{a}'$. We start off by setting $\mathfrak{b} = R$, $t' = t$, $a_i = 0$ and $S_i' = S_i$ for $1 \leq i \leq t$, and $\varphi = 0$. Throughout the algorithm $\varphi : R/\mathfrak{b} \xrightarrow{\sim} \bigoplus_{i=1}^{t'} (S_i')^{a_i}$ and $\bigcap_{i=1}^{t'} \mathrm{ann}_R(S_i') \subseteq \mathfrak{b}$ will be invariant. If for all i we have $\mathfrak{b}S_i' = 0$, then we are done. Otherwise, we choose $1 \leq h \leq t'$ and $s \in S_h'$ such that $\mathfrak{b}s \neq 0$. We define

$$\psi : R \to S_h' \oplus \bigoplus_{i=1}^{t'} (S_i')^{a_i}, \quad \psi(r) = (rs, \varphi(r)).$$

Then $\ker(\psi) = \operatorname{ann}_R(s) \cap \mathfrak{b} \subsetneq \mathfrak{b}$, since $\mathfrak{b}$ did not annihilate s. Let $\overline{\psi}$ be the map induced by ψ on $R/\ker(\psi)$. Then

$$\overline{\psi} \text{ isomorphism} \iff \psi \text{ surjective} \iff S'_h \oplus \{0\} \subseteq \operatorname{im}(\psi) \iff \mathfrak{b}s = S'_h.$$

This comes as a confirmation of our intuition: we are treating the S'_i as if they were simple, so if $0 \neq \mathfrak{b}s \leq S'_h$, then we would want $\mathfrak{b}s$ to be the whole of S'_h.

If ψ is surjective, we make the following replacements: $a_h := a_h + 1$, $\varphi := \psi$ and $\mathfrak{b} := \ker \psi$. Note that now $\mathfrak{b}$ has smaller length than before.

If ψ is not surjective, we replace S'_h by $S'_h/\mathfrak{b}s$, t' by $t'+1$ and put $S'_{t'+1} := \mathfrak{b}s$ and $a_{t'+1} := 0$. In addition, we replace φ by the composition $\pi \circ \varphi$, where π is the canonical map $\pi : \bigoplus_{i=1}^{t'} [\mathrm{old}(S'_i)^{\mathrm{old}\, a_i}] \to \bigoplus_{i=1}^{t'+1} [\mathrm{new}(S'_i)^{\mathrm{new}\, a_i}]$ and we replace $\mathfrak{b}$ by the kernel of our new φ. Note that the new $\mathfrak{b}$ will contain the old $\mathfrak{b}$, but the improvement is now sitting in the new S'_i.

Consider the quantity $l(S'_i)_{i=1}^{t'} = \sum_{i=1}^{t'} (2\,\mathrm{length}(S'_i) - 1)$. Then $0 \leq t' \leq l(S'_i)_{i=1}^{t'}$, since each S'_i has length at least 1. Also, $l(S'_i)_{i=1}^{t'}$ is bounded above by its initial value. At each iteration of the algorithm, we either see a decrease in the value of $l(S'_i)_{i=1}^{t'}$, when we improve our sequence, or we see a decrease in the length of $\mathfrak{b}$, whose length is initially equal to $\mathrm{length}(R)$. Since $\mathrm{length}(R) \leq \log_2(|R|)$ (recall that R was finite), the algorithm runs in polynomial time. □

We now turn to the main algorithm:

Theorem 4.4.3. *There exists a deterministic polynomial-time algorithm that, given a finite ring R and a finite R-module M, computes a set of generators for M of minimum cardinality.*

Proof. We begin by running the UPDATE algorithm presented in Proposition 4.4.2 with input $t = 1$ and $S_1 = {}_R R$. Let $X = \{x_1, \ldots, x_d\}$ be a set of R-generators of M. Then $\overline{X} = \{x_i + \mathfrak{a}M \mid 1 \leq i \leq d\}$ generates $M/\mathfrak{a}M$ over $R/\mathfrak{a}$. This gives a surjective homomorphism $(R/\mathfrak{a})^d \cong \bigoplus_{i=1}^{t} S_i^{a_i d} \twoheadrightarrow M/\mathfrak{a}M$. Relabel the S_i to get a map $\bigoplus_{i \in I} S_i \twoheadrightarrow M/\mathfrak{a}M$. We would like to find a subset $J \subseteq I$ for which this map becomes an isomorphism. In this process, the standard proof of its existence in the case where the S_i are simple (see [60], Chapter XVII, Lemma 2.1) may produce a witness of nonsimplicity of some S_i, in which case we refine the sequence and call the UPDATE subroutine to start over. In the end, we will have produced an isomorphism $\bigoplus_{i=1}^{t} S_i^{c_i} \xrightarrow{\sim} M/\mathfrak{a}M$, for some $c_i \in \mathbb{Z}_{>0}$ (note that the ideal $\mathfrak{a}$ may now be different from the one we started with).

Let $n = \max_i \{\lceil \frac{c_i}{a_i} \rceil\}$. If for all $i \neq h$ we have $\mathrm{Hom}_R(S_i, S_h) = 0$, then there is a surjective map $(R/\mathfrak{a})^n \twoheadrightarrow M/\mathfrak{a}M$ and n is minimal with this property. Since $\mathfrak{a} \subseteq \mathrm{J}(R)$, we can lift this to a map $R^n \twoheadrightarrow M$ and produce n generators of M. If, however, for some $i \neq h$ we have that $\mathrm{Hom}_R(S_i, S_h)$ contains a nonzero element f, then we can once again refine our sequence (either using $\ker(f) \neq 0$ or $\operatorname{im}(f) \neq S_h$ or removing one of S_i or S_h from the list if they are isomorphic) and start over by calling the UPDATE routine on our newly improved sequence.

To establish the running time, consider again the quantity

$$l(S_i)_{i=1}^{t} = \sum_{i=1}^{t} (2\,\mathrm{length}(S_i) - 1).$$

At each iteration of the algorithm, we either produce an output or we improve on our sequence, which results in a decrease in $l(S_i)_{i=1}^{t}$. □

Note 4.4.4. As we have described the algorithm in Theorem 4.4.3, after we have improved our sequence, calling the UPDATE subroutine overwrites our so far acquired knowledge about $\mathfrak{a}$. There is possibly a way of saving some of this information by running a modified version of the UPDATE routine, which would thus make the main algorithm slightly more efficient. However, being more precise about this here would obscure the idea of the algorithm.

4.5 Remark on implementation and performance

Since our goal was to place the module isomorphism problem in the complexity class P, we have not been concerned with calculating running time exponents or performing a detailed complexity analysis. A more careful organisation of the subroutines of the algorithms presented, or indeed considering randomised variations thereof may yield better running times, but such endeavours are beyond our scope at this moment.

Chapter 5

A miscellaneous collection of algorithms

5.1 Testing if a ring is a field

One basic question we may ask ourselves when presented with a finite ring is if it is not in fact a field.

Lemma 5.1.1. *Let p be a prime and R a finite commutative $\mathbb{F}_p$-algebra. Then the following are equivalent:*

(i) *The map $F : R \to R$ given by $x \mapsto x^p$ is injective.*
(ii) *R has no nonzero nilpotent elements.*
(iii) *R is a field.*

Proof. Note that F is an $\mathbb{F}_p$-linear map.

(i)⇒(ii): Suppose there exists $0 \neq x \in R$ and $n \in \mathbb{Z}_{>1}$ such that $x^{n-1} \neq 0$ and $x^n = 0$. Choose $d \in \mathbb{Z}_{\geq 0}$ maximal such that $dp < n$. Then $(d+1)p \geq n$ and $d+1 < n$, so $0 \neq x^{d+1} \in \ker(F)$.

(ii)⇒(iii): If R has no nonzero nilpotent elements, then R is semisimple (see Theorem 1.4.9, part (iii)). Since R is commutative, it must be a field.

(iii)⇒(i) is clear. □

Theorem 5.1.2. *There exists a deterministic polynomial-time algorithm that, given a finite ring R, determines whether or not R is a field.*

Proof. If $\mathrm{char}(R)$ is not a prime or R is not commutative, then R is not a field. This can be tested using Theorem 1.1.3 and Theorem 3.3.1. Otherwise R is a commutative $\mathbb{F}$-algebra and we use Lemma 5.1.1, part (i) to test whether it is a field. □

Note 5.1.3. Another deterministic polynomial-time algorithm for testing if a finite ring is a field is given in [4], Section 5, Theorem 4.

Note 5.1.4. A deterministic polynomial-time algorithm for the case where $R^+ = (\mathbb{Z}/p\mathbb{Z})^n$, for some prime p and some $n \in \mathbb{Z}_{>0}$ is given in [23], Section 4.

5.2 Testing if a ring is simple

We have seen in Theorem 1.1.3 that primality testing is in P. It is therefore natural to ask whether we can construct a deterministic polynomial-time algorithm which decides whether a finite ring is simple.

Lemma 5.2.1. *Let R be a semisimple ring. Then the centre of R is a field if and only if R is simple.*

Proof. Since R is semisimple, it is a finite product of simple rings. Hence the centre of R is the product of the centres of these simple rings, and so, is a finite product of fields. This product is then a field itself if and only if R was simple to begin with. □

Theorem 5.2.2. *There exists a deterministic polynomial-time algorithm that, given a finite ring R, determines whether R is simple, and if it is, outputs a prime p and two positive integers m and n such that $R \cong \mathcal{M}_n(\mathbb{F}_{p^m})$.*

Proof. Compute $k := \mathrm{Z}(R)$ and test if it is a field using Theorem 5.1.2. If k is not a field, then R cannot be simple. If it is a field, then R is a finite-dimensional algebra over k, and we can compute its Jacobson radical using Theorem 3.4.1. If $\mathrm{J}(R) \neq 0$, then R cannot be simple. If $\mathrm{J}(R) = 0$, then R is a semisimple algebra whose centre is a field, which by Lemma 5.2.1 implies R is simple. Proceed by computing $m \in \mathbb{Z}_{>0}$, the dimension of k over $\mathbb{F}_p$, and $n \in \mathbb{Z}_{>0}$, the size of the matrix ring, both of which can be done in polynomial time.

□

Note 5.2.3. Given two finite rings R and R', we can now decide if they are simple and isomorphic: simply compare the size of the matrices which will be produced by Theorem 5.2.2 and test if the centres of R and R' are isomorphic fields, using Theorem 3.5.2.

The algorithm above does not explicitly exhibit an isomorphism between R and $\mathcal{M}_n(\mathrm{Z}(R))$. We can use Theorem 3.5.2 to get an isomorphism of fields $\mathrm{Z}(R) \cong \mathbb{F}_{p^m}$, for some $m \in \mathbb{Z}_{>0}$, but we can say no more than that.

The problem of exhibiting an isomorphism $R \cong \mathcal{M}_n(\mathrm{Z}(R))$ is often referred to as the *explicit isomorphism problem*, and has received recent attention in [45, 46]. In the case that $\mathrm{Z}(R)$ is a finite field and n is a power of 2, this problem has a deterministic polynomial-time solution. In general, the problem of finding an isomorphism between finite algebras over finite fields is not believed to be NP-hard, but is at least as hard as the graph isomorphism problem ([41, 52]).

5.3 Testing if a module is simple

Testing simplicity of a module can also be done in polynomial time.

Theorem 5.3.1. *There exists a deterministic polynomial-time algorithm that, given a finite ring R and a finite R-module M, determines whether M is simple or not.*

Proof. By Schur's Lemma (Theorem 1.6.1), if M is simple, then $\mathrm{End}_R(M)$ is a division ring, so we start by computing $n := \mathrm{char}(\mathrm{End}_R(M)) = \exp(M^+)$ using Theorem 3.3.1. If n is not a prime, then we conclude that M is not simple. Otherwise, if n is a prime, then nR is a two-sided ideal in R, and $R' := R/nR$ is an algebra over a finite field, so we may compute its Jacobson radical using Theorem 3.4.1. Now M is an R'-module, since nR annihilates M.

If $\mathrm{J}(R')$ does not annihilate M, then M is not simple over R', and hence M is not simple over R. Otherwise, M is an $R'/\mathrm{J}(R')$-module and by Theorem 1.6.4, it is now enough to test whether $\mathrm{End}_{R'/\mathrm{J}(R')}(M)$ is a field, which can be done by Theorem 5.1.2. □

Second proof. Alternatively, compute

$$I := \mathrm{ann}(M) = \ker(R \to \mathrm{End}_{\mathbb{Z}}(M^+), r \mapsto r \cdot m),$$

where "$\cdot$" denotes the action of R on M. Then M is a faithful R/I-module and so if M is simple, we claim that R/I is simple as a ring. To see this, suppose M is simple. Then the Jacobson radical of R/I annihilates M, but since M is faithful, $\mathrm{J}(R/I) = 0$, hence R/I is semisimple. Now M is a faithful simple module over a semisimple ring, so R/I must in fact be simple.

We thus begin by testing simplicity of R/I as a ring, using Theorem 5.2.2. If R/I is not simple, then M cannot be simple and we are done. Otherwise, the algorithm in Theorem 5.2.2 will output a prime p, and two integers m, n such that $R/I \cong \mathcal{M}_n(\mathbb{F}_{p^m})$ as rings. Now, by Theorem 1.4.2, the only simple $\mathcal{M}_n(\mathbb{F}_{p^m})$-module, up to isomorphism, is $(\mathbb{F}_{p^m})^n$, and the other modules over $\mathcal{M}_n(\mathbb{F}_{p^m})$ are direct sums of $(\mathbb{F}_{p^m})^n$. Moreover, R/I-modules are exactly the R-modules annihilated by I. Hence M is simple over R if and only if $|M| = p^{nm}$. □

5.4 Testing if a module is projective

For many future algorithms it will be very useful to be able to test if a module is projective.

Theorem 5.4.1. *There exists a deterministic polynomial-time algorithm that, given a finite ring R and a R-module M, together with a generating set of cardinality d, for some $d \in \mathbb{Z}_{\geq 0}$, determines if M is R-projective or not, and if it is, produces a splitting of the natural surjection $R^d \twoheadrightarrow M$.*

Proof. Recall that M is projective if and only if the natural surjection $f : R^d \twoheadrightarrow M$ has a left inverse. The latter can be tested using Proposition 2.5.1, which will also produce a left inverse. □

Second proof. Another way to determine whether M is projective comes as a consequence of Theorem 4.1.1, since M is projective if and only if M is a direct summand of R^d. We compute the largest isomorphic common direct summand of R^d and M, say S. If $M \cong S$, then M is projective and the isomorphism $M \to S$, which is also produced by the algorithm, induces a splitting of $R^d \to M$. Otherwise the algorithms concludes that M is not projective. □

5.5 Constructing projective covers

Recall the definition of a projective cover given in Definition 1.6.24 and the fact that over left-artinian rings, all modules have a projective cover, unique up to isomorphism (Theorem 1.6.25).

Theorem 5.5.1. *There exists a deterministic polynomial-time algorithm that, given a finite ring R and a finite R-module M, outputs a projective cover of M.*

Proof. Use the algorithm in the proof of Theorem 4.1.3 to construct a sequence of R-modules $(S_i)_{i=1}^t$, two sequences of integers $(a_i)_{i=1}^t$ and $(c_i)_{i=1}^t$, and a two-sided nilpotent ideal $\mathfrak{a}$ such that

$$R/\mathfrak{a} \cong \bigoplus_{i=1}^{t} S_i^{a_i} \quad \text{and} \quad M/\mathfrak{a}M \cong \bigoplus_{i=1}^{t} S_i^{c_i},$$

where for all $1 \leq i \leq t$, we have that $a_i > 0$ and $c_i \geq 0$. Relabel, to write $R/\mathfrak{a} \cong \bigoplus_{i\in I} S_i$ and $M/\mathfrak{a}M \cong \bigoplus_{j\in J} S_j$. Then for each $j \in J$ we have a surjective map $g_j : M \twoheadrightarrow S_j$, such that $\mathfrak{a}M = \bigcap_{j\in J} \ker(g_j)$.

For each $j \in J$, pick $j' \in I$ such that $S_j \cong S_{j'}$. Since $S_{j'}$ is a direct summand of $R/\mathfrak{a}$, there exists an idempotent $\overline{e}_j \in R/\mathfrak{a}$ such that $S_{j'} \cong (R/\mathfrak{a})\overline{e}_j$ (see Theorem 1.5.4). To find $\overline{e}_j$, look at the image of 1 under the isomorphism $R/\mathfrak{a} \cong \bigoplus_{i\in I} S_i$ and identify the entry corresponding to j'. Replace all other entries by zeros and let $\overline{e}_j$ be the preimage of this element under the same isomorphism.

Thus, the decomposition $R/\mathfrak{a} \cong \bigoplus_{i\in I} S_i$ gives rise to a sequence of idempotent elements $\overline{e}_1, \ldots, \overline{e}_{|J|}$ in $R/\mathfrak{a}$ such that for all $j \in J$ we have $S_{j'} \cong (R/\mathfrak{a})\overline{e}_j$.

By Proposition 1.5.8, these idempotents can then be lifted deterministically in polynomial time to idempotents $e_1, \ldots, e_{|J|}$ in R. For all $j \in J$, let $P_j = Re_j$. Since e_j is an idempotent in R, we can write $R = Re_j \oplus R(1 - e_j)$, so P_j is projective. Hence we can construct a sequence of maps f_j such that for each $j \in J$ the following diagram commutes:

$$\begin{array}{ccc} & & P_j \\ & {\scriptstyle f_j}\swarrow & \downarrow{\scriptstyle \pi_j} \\ M & \xrightarrow[g_j]{} & S_j, \end{array}$$

where $\pi_j : P_j \to S_j$ is the natural projection map. This can be done by solving a system of linear equations over R. Let $P = \bigoplus_{j\in J} P_j$ and let $f : P \to M$ be the direct sum of the f_j.

We claim that for each $j \in J$, the pair (P_j, π_j) is a projective cover of S_j. Clearly π_j is surjective and P_j is projective. We need to show that $\ker(\pi_j) \subseteq_s P_j$. Let $N \leq P_j$ be a submodule such that $\ker(\pi_j) + N = P_j$. By construction, $\ker(\pi_j) = \mathfrak{a}P_j$. Since $\mathfrak{a} \subseteq \mathrm{J}(R)$, by Nakayama's Lemma we must have $N = P_j$. Since taking projective covers commutes with direct sums, $(P, \bigoplus_{j\in J} \pi_j)$ is a projective cover for $M/\mathfrak{a}M$. Since $\mathfrak{a}$ is nilpotent, f is surjective and (P, f) is a projective cover of M. □

5.6 Constructing injective hulls

Recall the definition of injective hulls given in Definition 1.6.28 and the fact that injective hulls exist for modules over any ring. Moreover, two injective hulls of a module M are isomorphic (Theorems 1.6.29).

To construct injective hulls, we will make use of the *character module* (see Definition 1.9.1). Recall that for a finite ring R, the character functor defines a duality between ${}^{\mathrm{fg}}_{R}\mathfrak{M}$ and $\mathfrak{M}^{\mathrm{fg}}_{R}$. (see Theorem 1.9.2). Moreover, the following holds.

Proposition 5.6.1. *Let R be a left-noetherian ring and M a finitely generated R-module. Then*

(i) *M is projective in ${}_{R}\mathfrak{M}$ if and only if M is projective in ${}^{\mathrm{fg}}_{R}\mathfrak{M}$.*
(ii) *M is injective in ${}_{R}\mathfrak{M}$ if and only if M is injective in ${}^{\mathrm{fg}}_{R}\mathfrak{M}$.*

Proof. The "only if" directions of both (i) and (ii) are clear. We prove the converse statements below.

Suppose M is projective in ${}_{R}\mathfrak{M}$. Then M is a direct summand of a finitely generated free R-module, so it is projective in ${}^{\mathrm{fg}}_{R}\mathfrak{M}$.

Suppose M is injective in ${}_{R}\mathfrak{M}$. Since R is left-noetherian, all its left ideals are finitely generated. Hence by Theorem 1.6.12, part (iii) (Baer's test), M is injective in ${}^{\mathrm{fg}}_{R}\mathfrak{M}$. □

Note 5.6.2. The left-noetherian condition on R is not needed for part (i).

Corollary 5.6.3. *Let R be a finite ring and M an R-module. Then M is injective over R if and only if $\widehat{M}$ is projective over R^{o}, the opposite ring of R.*

Note 5.6.4. If R is a finite ring and M is a finite R-module, then we may take

$$\widehat{M} = \mathrm{Hom}_{\mathbb{Z}}(M, \frac{1}{e}\mathbb{Z}/\mathbb{Z}), \tag{5.1}$$

where e is a multiple of $\exp(M^+)$.

Theorem 5.6.5. *There exists a deterministic polynomial-time algorithm that, given a finite ring R and a finite R-module M, computes an injective hull of M.*

Proof. Compute $\widehat{M} = \mathrm{Hom}_{\mathbb{Z}}(M, \frac{1}{e}\mathbb{Z}/\mathbb{Z})$, where e is a multiple of $\exp(M^+)$. Using Theorem 5.5.1, compute a projective cover (P, f) of $\widehat{M}$ over R^{o}. Now set $I := \widehat{P} = \mathrm{Hom}_{\mathbb{Z}}(P, \frac{1}{e'}\mathbb{Z}/\mathbb{Z})$, where e' is a multiple of $\exp(P^+)$. Then (I, g) is an injective hull of M by Theorem 1.9.2. The algorithm also produces a map $g : \widehat{\widehat{M}} \hookrightarrow I$, given by precomposition with f, such that $\mathrm{im}(g) \supseteq_e I$. □

5.7 Testing if a module is injective

Recall that a module is injective if and only if it is isomorphic to its injective hull (Theorem 1.6.30).

Theorem 5.7.1. *There exists a deterministic polynomial-time algorithm that, given a finite ring R and a finite R-module M, determines if M is injective.*

Proof. Construct an injective hull I of M using Theorem 5.6.5. Now check if the map $g : M \hookrightarrow I$ produced by Theorem 5.6.5 is bijective. □

5.8 Testing if a ring is quasi-Frobenius

Recall that a finite ring R is quasi-Frobenius if R is left self-injective (Theorem 1.7.1, Definition 1.7.2).

Theorem 5.8.1. *There exists a deterministic polynomial-time algorithm that, given a finite ring R, determines whether R is quasi-Frobenius.*

Proof. Use Theorem 5.7.1 to determine if the left-regular R-module ${}_RR$ is injective. □

5.9 Constructive tests for existence of injective and surjective module homomorphisms

In this section we discuss the algorithmic problem of testing for existence and of finding injective and surjective homomorphisms between two finite length modules over a ring R. If R is a finite-dimensional algebra over a field, this problem can be cast in the context of matrix completion, and has been shown to be NP-hard in [42]. In view of the results of [10, 15] and of Theorem 4.1.2, this result is striking. It is not however an isolated type of result: the subgraph isomorphism problem is an NP-hard problem, while the graph isomorphism problem is believed to be NP-intermediate.

While in the general case, testing constructively for existence of injective and surjective module homomorphisms is NP-hard, with certain restrictions on the modules considered, the problem turns out to be tractable. We are interested in the case where R is a finite ring and one of the modules is either projective or injective over R, for which the problem simplifies somewhat.

Theorem 5.9.1. *There exists a deterministic polynomial-time algorithm that, given a finite ring R and two finite R-modules M and N, one of which is projective, determines whether there exists a surjection $M \twoheadrightarrow N$. If one exists, the algorithm exhibits one such.*

Proof. If N is projective, then it suffices to test whether N is a direct summand of M, which can be done by Theorem 4.1.3. This will also produce a surjection $M \twoheadrightarrow N$.

If M is projective, then we proceed by constructing a projective cover (P, f) of N. Note that existence of a surjection $M \twoheadrightarrow N$ is equivalent to existence of a surjection $M \twoheadrightarrow P$. If there exists a surjection $M \twoheadrightarrow P$, simply compose it with f to get a surjection $M \twoheadrightarrow N$. Conversely, if there exists a surjection $M \twoheadrightarrow N$, then, since P is a projective cover of N, there exists a surjective map $g : M \twoheadrightarrow P$ making the following diagram commute

$$\begin{array}{ccc} & & P \\ & \overset{g}{\nearrow} & \downarrow f \\ M & \longrightarrow\!\!\!\!\!\rightarrow & N. \end{array}$$

This reduces the problem to the previous case. □

Second proof for the case where M is projective. If M is projective, we can also decide existence of a surjection $M \to N$ in a more direct manner. Use the algorithm in the proof of Theorem 4.1.3 to construct a two-sided nilpotent ideal $\mathfrak{a} \subset R$ and a sequence of R-modules $(S_i)_{i=1}^t$ that is "compatible" both with M and with N, i.e. such that

$$M/\mathfrak{a}M \cong \bigoplus_{i=1}^{t} S_i^{a_i} \quad \text{and} \quad N/\mathfrak{a}N \cong \bigoplus_{i=1}^{t} S_i^{b_i},$$

for some $a_i, b_i \in \mathbb{Z}_{\geq 0}$. This is done by running the algorithm in the proof of Theorem 4.1.3 on the ring R and the module M, and including N as one of the candidates for the isomorphism classes of simple R-modules in the UPDATE subroutine. Note that, by construction, the algorithm ensures that for all $i \neq j$, we have $\mathrm{Hom}_R(S_i, S_j) = 0$.

We claim that existence of a surjection $M/\mathfrak{a}M \twoheadrightarrow N/\mathfrak{a}N$, which can be easily tested by comparing a_i and b_i for each $1 \leq i \leq t$, is equivalent to existence of a surjection $M \to N$. If there exists a surjection $M \to N$, then clearly it induces a surjection $M/\mathfrak{a}M \twoheadrightarrow N/\mathfrak{a}N$. Conversely, if there exists a surjection $\overline{f} : M/\mathfrak{a}M \twoheadrightarrow N/\mathfrak{a}N$, then, since M is projective, there exists a map $f : M \to N$ making the following diagram commute:

$$\begin{array}{ccc} M & \overset{f}{\dashrightarrow} & N \\ \downarrow & & \downarrow \\ M/\mathfrak{a}M & \underset{\overline{f}}{\longrightarrow\!\!\!\!\!\rightarrow} & N/\mathfrak{a}N, \end{array}$$

and f is surjective since $\mathfrak{a}$ is nilpotent. □

Note 5.9.2. The case where $M \cong R^n$, for some $n \in \mathbb{Z}_{>0}$ can be settled using the algorithm for computing the minimum number of generators of a module, given in Theorem 4.1.3.

Dually to Theorem 5.9.1, we have the following result:

Theorem 5.9.3. *There exists a deterministic polynomial-time algorithm that, given a finite ring R and two finite R-modules M and N, one of which is injective, determines whether or not there exists an injection $M \hookrightarrow N$. If one exists, the algorithm exhibits one such.*

Proof. Let $k := \frac{1}{e}\mathbb{Z}/\mathbb{Z}$, where $e \in \mathbb{Z}_{>0}$ is a multiple of $\exp(M^+)$ and $\exp(N^+)$, and apply Theorem 5.9.1 to modules $\mathrm{Hom}_k(N, k)$ and $\mathrm{Hom}_k(M, k)$.

□

The remaining cases, not treated by Theorems 5.9.1 and 5.9.3, are constructive tests for existence of the following R-module homomorphisms:

- $P \hookrightarrow M$, for P a projective module,
- $M \hookrightarrow P$, for P a projective module,

and their respective duals,

- $N \twoheadrightarrow I$, for I an injective module,
- $I \twoheadrightarrow N$, for I an injective module.

Mimicking the construction given in the proof of Theorem 1.2 of [42], we settle these as being NP-hard, even when R is a finite local commutative ring and $P = R$. This is done by a reduction from an instance of the *nonsingular matrix completion problem*, which is known to be NP-hard.

The nonsingular matrix completion problem is an algorithmic question that can be formulated as follows: given a square matrix A, whose entries are homogeneous linear polynomials in $\mathbb{F}[x_1, \ldots, x_n]$, for some field $\mathbb{F}$, decide if there exist values from $\mathbb{F}$ that can be assigned to the variables $x_1, \ldots, x_n$, so as to make A nonsingular. The constructive version of this problem asks for values of $x_1, \ldots, x_n$ making A nonsingular to be exhibited.

Nonsingular matrix completion problems arise naturally in spaces of linear transformations. Let $n \in \mathbb{Z}_{>0}$ and $\mathbb{F}$ be a field. Let $A \subset \mathcal{M}_n(\mathbb{F})$ be a linear subspace and let $\{A_1, \ldots, A_m\}$ be a basis of A over $\mathbb{F}$. Deciding existence of (resp. finding) a nonsingular matrix in A is equivalent to deciding existence of (resp. finding) a sequence $c_1, \ldots, c_m \in \mathbb{F}$ such that $\sum_{i=1}^{m} c_i A_i$ is nonsingular.

The complexity of the nonsingular matrix completion problem is very much dependent on the size of the field $\mathbb{F}$ (see [40, 42]). If $\mathbb{F}$ is "large enough", then the Schwartz-Zippel lemma (see [81, 89]) provides an efficient randomized solution. However, over finite fields, nonsingular matrix completion is NP-complete ([14, 40]).

Theorem 5.9.4. *There exists a deterministic polynomial-time reduction from the decision (resp. constructive) version of nonsingular matrix completion to the problem of deciding existence of (resp. finding) an injective module homomorphism from a finite commutative local ring R containing a field, to an R-module M.*

Proof. Let $\mathbb{F}$ be a finite field and let U, V be two finite-dimensional $\mathbb{F}$-vector spaces of the same dimension. Let $0 \neq L \leq \mathrm{Hom}_{\mathbb{F}}(U, V)$ be a linear subspace. Consider the ring

$$R = \mathbb{F} \oplus U,$$

with componentwise addition and multiplication given by

$$(a, x)(b, y) = (ab, ay + bx).$$

Then R is a commutative local ring, with maximal ideal U, and $U^2 = 0$.

Put

$$M = L \oplus V.$$

We make M into an R-module by defining an action:

$$(a, u) \cdot (l, v) := (al, av + l(u)),$$

for all $a \in \mathbb{F}$, $u \in U$, $l \in L$ and $v \in V$.

Note that any homomorphism $R \to M$ is determined by the image of 1_R. Let $\psi : R \to M$ be an R-module homomorphism, and suppose $1 \mapsto (l, v)$, for some $(l, v) \in M$. Then

$$\mathrm{im}(\psi) = \mathbb{F}(l, v) + (0, lU)$$

and $\psi \in \mathrm{Hom}_R(R, M)$ is injective if and only if $l \in L$ is an isomorphism.

□

Theorem 5.9.5. *There exists a deterministic polynomial-time reduction from the decision (resp. constructive) version of nonsingular matrix completion to the problem of deciding existence of (resp. finding) an injective module homomorphism from an R-module M to R, where R is a finite commutative local ring containing a field.*

Proof. Let $\mathbb{F}$ be a finite field and let U, V be two finite-dimensional $\mathbb{F}$-vector spaces of the same dimension. Let $0 \neq L \subseteq \mathrm{Hom}_{\mathbb{F}}(U, V)$ be a linear subspace. Consider the ring

$$R = \mathbb{F} \oplus L \oplus U \oplus V,$$

with componentwise addition and multiplication given by

$$(f, l, u, v) \cdot (f', l', u', v') = (ff', fl' + f'l, fu' + f'u, fv' + fv + l(u') + l'(u)).$$

Then R is a commutative ring with unique maximal ideal $L \oplus U \oplus V$. Note that $L \oplus V$ is a two-sided ideal in R. Put

$$M := R/(L \oplus V) \cong \mathbb{F} \oplus U.$$

Note that $U^2 = 0$, so M also has the structure of a local commutative ring, with maximal ideal U.

Note that

$$\mathrm{Hom}_R(M, R) \cong \mathrm{ann}_R(L \oplus V) = L \oplus U_0 \oplus V, \tag{5.2}$$

where $U_0 = \bigcap_{f \in L} \ker(f)$ and the isomorphism in (5.2) is given by mapping $f \mapsto f(\overline{1})$. Let ϕ be an R-module homomorphism. Then ϕ corresponds uniquely to an element $(0, l_0, u_0, v_0) \in L \oplus U_0 \oplus V$ and

$$\mathrm{im}(\phi) \cong \mathbb{F}(0, l_0, u_0, v_0) + (0, 0, 0, l_0 U).$$

Hence $\phi \in \mathrm{Hom}_R(M, R)$ is injective if and only if $l_0 \in L$ is an isomorphism.

□

We consider now another weaker variant of the problem of testing constructively for injective and surjective module homomorphisms. Suppose we are given a finite ring R and two modules M and N. Instead of looking for a surjection $M \twoheadrightarrow N$, we may ask if there is an integer k such that there exists a surjective homomorphism $f : M^k \twoheadrightarrow N$. If such a pair (k, f) exists, we would like to exhibit it. Note that we do not ask for k to be minimal with this property. This problem turns out to have an easy solution.

Theorem 5.9.6. *There exists a deterministic polynomial-time algorithm that, given a finite ring R and two finite R-modules M, N, decides if there exists a pair (k, f), where $k \in \mathbb{Z}_{\geq 0}$ and $f : M^k \to N$ is a surjective R-module homomorphism. If it exists, the algorithm exhibits such a pair.*

Proof. Compute a set S of $\mathbb{Z}$-generators of $\mathrm{Hom}_R(M, N)$. If $N = \sum_{f \in S} f(M)$, then output $(|S|, \sum_{f \in S} f)$. Otherwise conclude that there does not exist a pair as required.

□

Dually, we have the following result:

Theorem 5.9.7. *There exists a deterministic polynomial-time algorithm that, given a finite ring R and two finite R-modules M, N, decides if there exists a pair (k, f), where $k \in \mathbb{Z}_{\geq 0}$ and $f : M \to N^k$ is an injective R-module homomorphism. If it exists, the algorithm exhibits such a pair.*

Proof. Compute a set S of $\mathbb{Z}$-generators of $\mathrm{Hom}_R(M, N)$. If $\bigcap_{f \in S} \ker(f) = \{0\}$, then output $(|S|, \prod_{f \in S} f)$. Otherwise conclude that there does not exist a pair as required.

□

Chapter 6

Approximating the Jacobson radical of a finite ring

When trying to answer questions about left-artinian rings and modules over them, it is often convenient to reduce the problem at hand to the semisimple case, where structures are much more manageable, and then "lift". However, this approach requires the computation of the Jacobson radical of the ring, which we cannot efficiently carry out in general. But how "close" can we get to semisimplicity with a deterministic polynomial-time algorithm in the case of finite rings?

In this chapter, we give two deterministic polynomial-time algorithms that, given a finite ring A, produce two-sided nilpotent ideals $\mathfrak{j}_A$, such that $A/\mathfrak{j}_A$ is "almost semisimple". We think of such ideals $\mathfrak{j}_A$ as approximations to the Jacobson radical.

6.1 Introduction

When considering the module isomorphism problem for finite rings, not being able to compute Jacobson radicals was the main obstacle in the way of generalising methods that had worked in the case of finite-dimensional algebras over finite fields (cf. [15]). The side-exit algorithm of Theorem 4.1.3 was designed to construct an approximation of the Jacobson radical which was good enough for the purpose at hand, namely determining the minimum number of generators of a module. Motivated by this, we design deterministic polynomial-time algorithms that compute good working approximations of the Jacobson radical of a finite ring, that is, two-sided nilpotent ideals such that when we quotient the ring by them, we are left with something that is "almost" semisimple.

The notion we will use to approximate semisimplicity is that of separability. Given a commutative ring R, an R-algebra S is said to be *separable* over R if S is projective as an $S \otimes_R S^{\mathrm{o}}$-module, where S^{o} denotes the opposite ring of S. A ring is said to be separable if it is separable as a $\mathbb{Z}$-algebra. Section 6.2 explores the structure and properties of separable algebras and attempts to make an argument for why they

are a good "approximation" to semisimple algebras. Section 6.2.6 gives a complete classification of finite rings that are separable over $\mathbb{Z}$, as finite products of matrix rings over certain commutative rings.

It turns out that finite separable rings are automatically projective over a certain subring, which we will refer to as the *generalised prime subring* (see Section 6.2.5). With this in mind, we make the following definition.

Definition 6.1.1. *Let A be a finite ring. We say an ideal $\mathfrak{j}_A \subset A$ is an* approximation of the Jacobson radical of A *if it satisfies the following conditions:*

(A1) *$\mathfrak{j}_A$ is a nilpotent two-sided ideal of A,*
(A2) *$A/\mathfrak{j}_A$ is separable,*
(A3) *The prime subring and generalised prime subring of $A/\mathfrak{j}_A$ coincide.*

If $\mathfrak{j}_A$ is an approximation of the Jacobson radical of A, then the ring $A/\mathfrak{j}_A$ has many of the good properties that semisimple rings have: it has "many" projective and injective modules, it is quasi-Frobenius, it is a symmetric algebra over its prime subring and it is isomorphic to a product of matrix rings over commutative local rings. However, as opposed to semisimplicity, which can be neither tested nor enforced (see Note 3.4.2), separability is a much friendlier notion in algorithmic contexts. In Sections 6.3 and 6.4, we prove the following results.

Theorem 6.1.2. *There exists a deterministic polynomial-time algorithm that, given a finite ring A, computes an approximation of the Jacobson radical of A.*

We prove Theorem 6.1.2 by exhibiting two algorithms that produce approximations of Jacobson radicals of finite rings.

Proposition 6.1.3. *Let $\mathcal{F}$ be the class of finite rings. The two families of ideals $(\mathfrak{j}_A)_{A\in\mathcal{F}}$ and $(\mathfrak{j}'_A)_{A\in\mathcal{F}}$, produced by the two algorithms described in the proof of Theorem 6.1.2 are functorial under isomorphisms, i.e. if $\phi : A \to B$ is an isomorphism of finite rings, then $\phi(\mathfrak{j}_A) = \mathfrak{j}_B$ and $\phi(\mathfrak{j}'_A) = \mathfrak{j}'_B$.*

Theorem 6.1.4. *There exists a deterministic polynomial-time algorithm that, given a finite ring A, computes the generalised prime subring of A.*

In Section 6.3.6, we look at some basic examples of the running of these algorithms.

6.2 Separability

6.2.1 Separable algebras

We begin with a study of separable algebras and some of their basic properties. We argue that the notion of separability is a good starting point in our quest to approximate semisimplicity. The main references for this section are [20, 54].

Definition 6.2.1. *Let R be a commutative ring and S an R-algebra. We say S is* separable *over R if S is projective as a module over $S^{\mathrm{e}} := S \otimes_R S^{\mathrm{o}}$, where S^{o} denotes the opposite ring of S and the module structure is given by $(s \otimes s')t = sts'$, for $s, s', t \in S$.*

Note 6.2.2. For any R-algebra S, the ring S^{e} is called the *enveloping algebra* of S, which justifies the choice of notation. A left S^{e}-module is the same as an S-S-bimodule whose induced R-structures coincide.

Let $\phi : S^{\mathrm{e}} \to S$ be the S^{e}-module homomorphism given by $a \otimes a' \mapsto aa'$. Note that $\ker(\phi)$ is then generated as an S^{e}-module by elements of the form $s \otimes 1 - 1 \otimes s$.

Theorem 6.2.3 ([20], Chapter II, Proposition 1.1)**.** *Let R be a commutative ring and S an R-algebra. Then the following are equivalent:*

(i) *S is separable over R.*
(ii) *The exact sequence of S^{e}-modules $0 \to \ker(\phi) \to S^{\mathrm{e}} \xrightarrow{\phi} S \to 0$ splits.*
(iii) *There exists an element $e \in S^{\mathrm{e}}$ such that $\phi(e) = 1$ and $\forall s \in S$, $(s \otimes 1)e = (1 \otimes s)e$.*

Note 6.2.4. The element e in (iii) is necessarily an idempotent, since $e^2 - e = (e - (1 \otimes 1))e \in \ker(\phi)e = 0$, and it is referred to as a *separability idempotent*. It arises as the image of $1 \in S$ under a splitting of ϕ and is in general not unique.

Note 6.2.5.

(i) Let $\tau : S^{\mathrm{e}} \to (S^{\mathrm{o}})^{\mathrm{e}}$ be the map given by $x \otimes y \mapsto y \otimes x$. Then τ is a ring isomorphism. For this, note that for all $x_1, x_1, y_1, y_2 \in S$:

$$\tau((x_1 \otimes y_1)(x_2 \otimes y_2)) = \tau(x_1x_2 \otimes y_2y_1) = y_2y_1 \otimes x_1x_2$$

and

$$\tau(x_1 \otimes y_1)\tau(x_2 \otimes y_2) = (y_1 \otimes x_1)(y_2 \otimes x_2) = y_2y_1 \otimes x_1x_2.$$

(ii) The map $\tau' : S^{\mathrm{e}} \to S^{\mathrm{e}}$ given by $x \otimes y \mapsto y \otimes x$ is an involutive ring anti-automorphism, since $(\tau')^2 = \mathrm{id}$ and $\tau'((x_1 \otimes y_1)(x_2 \otimes y_2)) = y_2y_1 \otimes x_1x_2 = (y_2 \otimes x_2)(y_1 \otimes x_1) = \tau'(x_2 \otimes y_2)\tau'(x_1 \otimes y_1)$. A separability idempotent that satisfies $\tau'(e) = e$ is called a *symmetric separability idempotent*.

Proposition 6.2.6. *Let R be a commutative ring and S be a separable R-algebra. Then S^{o} is a separable R-algebra.*

Proof. Let $S^{\mathrm{e}} = S \otimes_R S^{\mathrm{o}}$ and $(S^{\mathrm{o}})^{\mathrm{e}} = S^{\mathrm{o}} \otimes_R S$. Consider the two exact sequences

$$\begin{array}{ccccccccc} 0 & \longrightarrow & S^{\mathrm{e}} & \underset{\phi}{\overset{\sigma}{\leftrightarrows}} & S & \longrightarrow & 0 \\ & & \tau\downarrow & & \downarrow\mathrm{id} & & \\ 0 & \longrightarrow & (S^{\mathrm{o}})^{\mathrm{e}} & \xrightarrow{\psi} & S^{\mathrm{o}} & \longrightarrow & 0, \end{array}$$

where $\phi : S^{\mathrm{e}} \to S$ is given by $\phi(s \otimes s') = ss'$, $\psi : (S^{\mathrm{o}})^{\mathrm{e}} \to S^{\mathrm{o}}$ is given by $\psi(s' \otimes s) = ss'$ and τ is the ring isomorphism described in Note 6.2.5. Let σ be a splitting of ϕ, which exists since S is a separable algebra over R. It is easy to see that the above diagram commutes and since σ is S^{e}-linear, $\tau \circ \sigma$ is $(S^{\mathrm{o}})^{\mathrm{e}}$-linear. Hence $\tau \circ \sigma$ gives a splitting of ψ, showing that S^{o} is separable over R. □

Note 6.2.7. If e is a separability idempotent of S over R, then $\tau(e)$ is a separability idempotent of S^{o} over R.

The following proposition describes separability over fields.

Theorem 6.2.8.

(i) ([54], Chapter III, Proposition 3.4). *Let $R \subset S$ be a finite field extension. Then S is separable over R if and only if for all $s \in S$, the minimal polynomial of s over R has distinct roots in a splitting field over R.*

(ii) ([20], Chapter II, Theorem 2.5). *Let R be a field and S an R-algebra. Then S is separable over R if and only if S is finite-dimensional over R and for every field extension $R \subset K$, we have that $S \otimes_R K$ is semisimple.*

(iii) ([54], Chapter III, Theorem 3.1). *Let R be a field and S an R-algebra. Then S is separable over R if and only if S is isomorphic to an algebra of the form $\prod_{i=1}^{n} \mathcal{M}_{n_i}(D_i)$, for some $n \in \mathbb{Z}_{\geq 0}$ and $n_i \in \mathbb{Z}_{>0}$, with D_i finite-dimensional division algebras over R and $\mathrm{Z}(D_i) \supset R$ finite-dimensional separable field extensions.*

Corollary 6.2.9. *A separable algebra over a field is semisimple.*

Proof. Put $K = R$ in Proposition 6.2.8, part (ii). □

Since finite extensions of perfect fields are separable, semisimplicity and separability are equivalent notions for finite-dimensional algebras over perfect fields.

Corollary 6.2.10. *Let R be a perfect field and S a finite-dimensional R-algebra. Then S is a separable R-algebra if and only if S is semisimple.*

Example 6.2.11. Let R be a commutative ring and let $n \in \mathbb{Z}_{\geq 0}$. Then the matrix ring $\mathcal{M}_n(R)$ is a separable R-algebra. To see this, let E_{ij} denote the matrix with $(i,j)^{\text{th}}$ entry equal to 1 and all other entries equal to 0. Then for any fixed i, the element $e_i = \sum_{j=1}^{n} E_{ij} \otimes E_{ji}$ is a separability idempotent.

Example 6.2.12. Let R be a commutative ring and let G be a finite group such that $|G|$ is a unit in R. Then the group algebra $R[G]$ is a separable R-algebra. To see that, note that $e := |G|^{-1} \sum_{g \in G} g \otimes g^{-1}$ is a separability idempotent.

Parts (i) and (ii) of Proposition 6.2.8 establish that the definition of separability is compatible with the classical definitions of separability in the cases of finite field extensions and finite-dimensional algebras over a field. The reason the extra finiteness condition is required is that separable algebras, as we have defined them, turn out to have a lot of structure, as shown in the following series of results.

Proposition 6.2.13 ([20], Chapter II, Proposition 2.1)**.** *Let R be a commutative ring and S a separable R-algebra. Suppose that S is projective as an R-module. Then S is finitely generated as an R-module.*

Proposition 6.2.14 ([20], Chapter II, Proposition 1.12, Theorem 3.8)**.** *Let R be a commutative ring.*

(i) *Let A be a separable commutative algebra over R. Suppose S is a separable A-algebra. Then S is an R-algebra and it is separable over R.*
(ii) *Let S be a separable R-algebra and A be any R-subalgebra of the centre of S. Then S is separable over A.*
(iii) *Suppose S is an R-algebra. Then S is separable over R if and only if S is separable over its centre, and its centre is separable over R.*

Proposition 6.2.15 ([54], Chapter III, Theorem 5.1)**.** *Let S be a ring. If S is separable over its centre, then S is projective as a module over its centre.*

Proposition 6.2.16 ([54], Chapter III, Proposition 1.7)**.**

(i) *Let R be a commutative ring and R_1, R_2 be two commutative R-algebras. Let S_1 be a separable R_1-algebra and S_2 a separable R_2-algebra. Then $S_1 \otimes_R S_2$ is a separable $R_1 \otimes_R R_2$-algebra, with $(r_1 \otimes r_2)(s_1 \otimes s_2) = r_1 s_1 \otimes r_2 s_2$. Moreover, $\mathrm{Z}(S_1 \otimes S_2) = \mathrm{Z}(S_1) \otimes \mathrm{Z}(S_2)$.*
(ii) *Let R_1, R_2 be two commutative rings. Let S_1 be an R_1-algebra and S_2 an R_2-algebra. Then $S_1 \times S_2$ is separable over $R_1 \times R_2$ if and only if S_1 is separable over R_1 and S_2 is separable over R_2.*
(iii) *Let R be a commutative ring and S_1, S_2 two R-algebras. Then $S_1 \times S_2$ is separable over R if and only if both S_1 and S_2 are separable over R.*

Corollary 6.2.17. *Let R be a commutative ring, R' a commutative R-algebra and S a separable R-algebra. Then $S \otimes_R R'$ is separable over over R'.*

Proof. In Proposition 6.2.16, part (i), take $R_1 := R$, $R_2 := R'$, $S_1 := S$ and $S_2 := R'$. □

Theorem 6.2.18 ([20], Chapter II, Theorem 7.1)**.** *Let R be a commutative ring and S an R-algebra that is finitely generated as an R-module. Then the following are equivalent:*

(i) *S is separable over R.*
(ii) *For every maximal ideal $\mathfrak{m}$ of R, we have that $S \otimes_R R_{\mathfrak{m}}$ is separable over $R_{\mathfrak{m}}$.*
(iii) *For every maximal ideal $\mathfrak{m}$ of R, the quotient $S/\mathfrak{m}S$ is separable over $R/\mathfrak{m}$.*

Separability is testable deterministically in polynomial time (cf. Note 3.4.2).

Theorem 6.2.19. *There exists a deterministic polynomial-time algorithm that, given a finite commutative ring R and a finite R-algebra S, decides whether or not S is separable over R.*

Proof. Using Proposition 2.4.1, we compute the enveloping algebra $S^{\mathrm{e}} = S \otimes_R S^{\mathrm{o}}$, after which we test projectivity of S over S^{e} using Theorem 5.4.1. □

6.2.2 Azumaya and finite-étale algebras

There are two distinguished classes of separable algebras that deserve special attention: Azumaya and finite-étale algebras.

Recall the definition of a progenerator (Definition 1.8.6).

Theorem 6.2.20 ([54], Chapter III, Theorem 6.1, [20], Chapter II, Theorem 3.4)**.** *Let R be a commutative ring and S an R-algebra. Then the following are equivalent.*

(i) *S is separable over R and $\mathrm{Z}(S) = R$.*
(ii) *S is an R-progenerator and the map $\alpha : S^{\mathrm{e}} \to \mathrm{End}_R(S)$, given by $s \otimes s' \mapsto (f : t \mapsto sts')$, is an isomorphism of R-algebras.*
(iii) *S is an S^{e}-progenerator and $\mathrm{Z}(S) = R$.*
(iv) *There exist an R-algebra T and an R-progenerator P such that $S \otimes_R T \cong \mathrm{End}_R(P)$ as R-algebras.*

Definition 6.2.21. *An R-algebra S satisfying the conditions of Theorem 6.2.20 is called an* Azumaya algebra *over R.*

Note 6.2.22. From Theorem 6.2.20, part (ii), it is easy to see that if S is Azumaya over R, then S^{o} is also Azumaya over R. This gives another, more conceptual way of showing that separability is stable under taking opposites. Suppose S is separable as an algebra over a commutative ring R. Then S is Azumaya over $\mathrm{Z}(S)$ and $\mathrm{Z}(S)$ is separable over R. Now $R^{\mathrm{o}} = R$ and $\mathrm{Z}(S) = \mathrm{Z}(S^{\mathrm{o}})$. So $\mathrm{Z}(S^{\mathrm{o}})$ is separable over R^{o} and S^{o} is Azumaya over $\mathrm{Z}(S^{\mathrm{o}})$. Hence S^{o} is separable over R, by Theorem 6.2.14, part (iii).

Example 6.2.23 ([26], §8)**.** Over a field, an algebra is Azumaya if and only if it is central simple.

Example 6.2.24 ([20], Chapter II, Proposition 4.1)**.** Let R be a commutative ring. Then the endomorphism ring of any R-progenerator is Azumaya over R.

Proposition 6.2.25 ([80], Proposition 3.9)**.** *Let R be a commutative ring and A an R-algebra that is Azumaya of constant rank over R. Then there exists a faithfully flat ring extension S of R, and $n \in \mathbb{Z}_{>0}$ such that $A \otimes_R S \cong \mathcal{M}_n(S)$.*

Corollary 6.2.26. *Let R be a commutative ring and A an Azumaya R-algebra. Then the rank of A over R, as a function on $\mathrm{Spec}(R)$, is a square.*

Proof. This follows from Proposition 6.2.25 and the fact that extension of scalars does not change the rank.

□

In the commutative setting, the notion we are interested in is that of a *finite-étale* algebra.

Definition 6.2.27. *Let R be a commutative ring. An R-algebra S is* finite-étale *over R if S is commutative, separable as an R-algebra and projective as an R-module.*

We state a couple of results describing the behaviour of finite-étale algebras and Azumaya algebras with respect to tensor products and direct products. These are consequences of Proposition 6.2.16.

Proposition 6.2.28. *Let R be a commutative ring and let R_1, R_2 be two commutative R-algebras.*

(i) *Let S_1 be an Azumaya R_1-algebra and S_2 an Azumaya R_2-algebra. Then $S_1 \otimes_R S_2$ is Azumaya over $R_1 \otimes_R R_2$.*
(ii) *Let S_1 be a finite-étale R_1-algebra and S_2 a finite-étale R_2-algebra. Then $S_1 \otimes_R S_2$ is finite-étale over $R_1 \otimes_R R_2$.*

Corollary 6.2.29. *Let R be a commutative ring and R' a commutative R-algebra.*

(i) *Let S be an Azumaya R-algebra. Then $S \otimes_R R'$ is Azumaya over R'.*
(ii) *Let S be a finite-étale R-algebra. Then $S \otimes_R R'$ is finite-étale over R'.*

Proposition 6.2.30. *Let R_1, R_2 be two commutative rings. Let S_1 be an R_1-algebra and S_2 an R_2-algebra. Then*

(i) *$S_1 \times S_2$ is Azumaya over $R_1 \times R_2$ if and only if S_1 is Azumaya over R_1 and S_2 is Azumaya over R_2.*
(ii) *$S_1 \times S_2$ is finite étale over $R_1 \times R_2$ if and only if S_1 is finite-étale over R_1 and S_2 is finite-étale over R_2.*

Moreover, if $R := R_1 = R_2$, then $S_1 \times S_2$ is finite étale over R if and only if S_1 and S_2 are both finite-étale over R.

Note 6.2.31. If $R := R_1 = R_2 \neq 0$ and S_1, S_2 are Azumaya R-algebras, then $S_1 \times S_2$ is not Azumaya over R.

6.2.3 The Brauer group

For a commutative ring R, we can define an equivalence relation on the collection of Azumaya R-algebras such that the equivalence classes form an abelian group with binary operation given by taking tensor products over R.

Definition 6.2.32. *Let R be a commutative ring. Let $\mathcal{B}(R)$ be a collection of Azumaya R-algebras such that every Azumaya R-algebra is isomorphic to exactly one element of $\mathcal{B}(R)$. Let*

$$\mathcal{B}^{\mathrm{o}}(R) = \{A \in \mathcal{B} \mid A \cong \mathrm{End}_R(P) \text{ as } R\text{-algebras, for some } R\text{-progenerator } P\}.$$

Define an equivalence relation $\sim$ on $\mathcal{B}(R)$ by:

$$A \sim B \iff \text{there exist } Y, Z \in \mathcal{B}^{\mathrm{o}}(R) \text{ such that } A \otimes_R Y \cong B \otimes_R Z \text{ as } R\text{-algebras.}$$

Denote by $[A]$ the equivalence class of $A \in \mathcal{B}(R)$ under $\sim$. The set of all such equivalence classes, denoted by $\mathrm{Br}(R)$, together with binary operation given by $[A] \cdot [B] = [A \otimes_R B]$ is an abelian group called the Brauer group *of R. The identity is given by $[R]$ and inverses are given by $[A]^{-1} = [A^{\mathrm{o}}]$.*

Note 6.2.33. Since all Azumaya R-algebras are finitely generated projective as R-modules, $\mathcal{B}(R)$ is indeed a set. It is also easy to check that $\sim$ is an equivalence relation and that $\mathrm{Br}(R)$ is an abelian group.

Example 6.2.34. (Brauer groups)

1. If k is a finite field, $\mathrm{Br}(k)$ is trivial (see [82], Chapter X, §7).
2. $\mathrm{Br}(\mathbb{Z})$ is trivial (see [26], page 196).
3. If k is a finite commutative ring, then $\mathrm{Br}(k)$ is trivial (see [87], Proposition 4.1).

For more on Brauer groups, see [20], Chapter III, Section 5.

6.2.4 Separable projective algebras

It is often convenient to look at algebras that are both finitely generated projective as modules and separable as algebras over the underlying commutative ring.

Definition 6.2.35. *Let R be a commutative ring. An R-algebra S that is separable as an R-algebra and projective as an R-module is said to be* separable projective *over R.*

This notion can be linked to the notions of Azumaya and finite-étale.

Theorem 6.2.36. *Let k be a commutative ring and S a k-algebra. Let $R := \mathrm{Z}(S)$. Then the following are equivalent:*

(i) *S is Azumaya over R and R is finite-étale over k.*
(ii) *S is separable projective over k.*

Proof. (i)⇒(ii) By Theorem 6.2.20, part (i), Definition 6.2.27 and Proposition 6.2.14, part (iii), we have that S is separable as a k-algebra. From Corollary 6.2.15 we know that S is projective as an R-module and so by transitivity of projectivity, S is projective as a k-module.

(ii)⇒(i) By Proposition 6.2.14, part (ii), we have that S is separable over R so it is Azumaya over R, and R is separable over k. All that remains to be established is that R is projective as a k-module.

From Theorem 6.2.20 we know that S is an R-progenerator, so

$$R = \sum_{f \in \mathrm{Hom}_R(S,R)} f(S).$$

Since S is finitely generated and projective as an R-module, $\mathrm{Hom}_R(S, R)$ is finitely generated as an R-module, so we can restrict the sum to a finite set of generators. In particular, there exists a surjective R-homomorphism $S^n \twoheadrightarrow R$ for some $n \in \mathbb{Z}_{>0}$. But R is R-projective, so this map splits, giving $S^n \cong R \oplus Q$ for some R-module Q. Now S is k-projective by hypothesis, so S^n is also k-projective and since R is a direct summand of S^n, we have that R is also k-projective.

□

Over a semisimple ring, every module is projective. The following proposition says that a separable projective algebra over a commutative ring has "many" projective modules.

Theorem 6.2.37 ([31], Proposition 2.3)**.** *Let R be a commutative ring, S a separable R-algebra and M a finitely generated S-module. Then any exact sequence of S-modules $0 \to M_1 \to M_2 \to M \to 0$ that splits over R, splits over S.*

Proposition 6.2.38. *Let R be a commutative ring and S a separable R-algebra that is projective as an R-module. Let M be an S-module. Then*

$$M \text{ is projective as an } S\text{-module} \iff M \text{ is projective as an } R\text{-module.} \tag{6.1}$$

Proof. ($\Rightarrow$) This direction is easy and follows by transitivity of projectivity: M is S-projective and S is R-projective, so M is R-projective.

($\Leftarrow$) This direction follows from Theorem 6.2.37, Proposition 5.6.1 and Note 5.6.2. □

Note 6.2.39. The "if" direction of (6.1) is a strong statement and only requires S to be separable over R. We will refer to this property as "projectivity lift". Another proof of this fact is also given in [20], Chapter II, Proposition 2.3, which uses the existence and properties of the separability idempotent. We sketch it here. Suppose M is an R-projective S-module. Let $f : N \to M$ be an S-module epimorphism. Since M is R-projective, there exists R-homomorphism $g : M \to N$ such that $fg = \mathrm{id}_M$. Note that $\mathrm{Hom}_R(M, N)$ is an S^{e}-module via $(a \otimes b) \cdot \phi(m) = a\phi(bm)$. Suppose e is a separability idempotent for S. Then $e \cdot g$ is an S-module homomorphism and $f(e \cdot g) = \mathrm{id}_M$. Hence M is S-projective.

Proposition 6.2.40. *Let R be a commutative ring and S a separable R-algebra that is projective as an R-module. Let M be an S-module. Then*

$$M \text{ is injective as an } S\text{-module} \iff M \text{ is injective as an } R\text{-module.} \tag{6.2}$$

Proof. ($\Rightarrow$) Recall from Proposition 1.6.20 that $\mathrm{Hom}_S({}_SS_R, {}_SM_R) \cong {}_RM_R$, so it is enough to show that $\mathrm{Hom}_S(S, M)$ is injective as an R-module. By tensor-hom adjunction (see Section 1.6.6), we have

$$\mathrm{Hom}_R(-, \mathrm{Hom}_S(S, M)) \cong \mathrm{Hom}_S(S \otimes_R -, M).$$

Since S is projective over R, the functor $S \otimes_R -$ is exact, and since M is injective over S, the functor $\mathrm{Hom}_S(-, M)$ is exact. Hence $\mathrm{Hom}_R(-, \mathrm{Hom}_S(S, M))$ is exact, i.e. $\mathrm{Hom}_S(S, M)$ is an injective R-module.

($\Leftarrow$) Consider an exact sequence of S-modules $0 \to I \to M \to C \to 0$, where M is an S-module and C is a cyclic S-module. Since I is R-injective the sequence is R-split, so because C is cyclic over S, the sequence is S-split by Theorem 6.2.37. The result now follows from Theorem 1.6.12, part (iv), which states that I is an injective S-module if and only if every short exact sequence $0 \to I \to M \to C \to 0$, where M is an S-module and C is a cyclic S-module, is S-split. □

Note 6.2.41. We will refer to the property induced by the "if" direction as "injectivity lift".

Corollary 6.2.42. *Let S be a finite ring that is separable projective over its prime subring. Then S is a quasi-Frobenius ring.*

Proof. Since S is finite, its prime subring is isomorphic to $R = \mathbb{Z}/n\mathbb{Z}$, where $n = \mathrm{char}(S) \in \mathbb{Z}_{>0}$. Since R is quasi-Frobenius (see Example 1.7.3), an R-module is injective if and only if it is projective. Since S admits both projectivity and injectivity lift from R by Propositions 6.2.38 and 6.2.40, it follows that S itself is quasi-Frobenius. □

We record some other properties of separable projective algebras.

Proposition 6.2.43. *Let A be a finite semisimple ring. Then A is separable projective over its prime subring.*

Proof. Since A is semisimple, the characteristic of A is squarefree. By Proposition 6.2.16, part (ii), Proposition 1.6.9 and the fact that A is semisimple, we may assume that A has prime subring $\mathbb{F}_p$, for some prime p, and that $A \cong \mathcal{M}_n(D)$, for some $n \in \mathbb{Z}_{>0}$, where D is a finite field extension of $\mathbb{F}_p$. Now A is separable projective over D by Example 6.2.23 and D is separable projective over $\mathbb{F}_p$, since finite extensions of perfect fields are separable. Hence A is separable projective over $\mathbb{F}_p$. □

Theorem 6.2.44. *Let A be a nonzero finite ring. Then A is semisimple if and only if A is separable projective over its prime subring and* $\mathrm{char}(A)$ *is squarefree.*

Proof. Let $n := \mathrm{char}(A)$. The "if" direction follows from Proposition 6.2.43 and the fact that for any $d \in \mathbb{Z}_{>0}$ such that $d^2 \mid n$, we have $0 \neq \frac{n}{d}A \subseteq \mathrm{J}(A)$. The other direction follows from Theorem 6.2.18, part (iii) and Proposition 6.2.16, part (ii). □

Moreover, separable projective algebras that are faithful as modules over the base ring, are symmetric (see Definition 1.8.1 and Theorem 1.8.2).

Theorem 6.2.45 ([24], Theorem 4.2)**.** *Let k be a commutative ring and A a separable projective k-algebra that is faithful as a module over k. Then A is a symmetric k-algebra.*

6.2.5 Separable rings

Let A be a ring. Then A is a $\mathbb{Z}$-algebra, as well as an algebra over its prime subring. By Proposition 6.2.14, parts (i) and (ii), we have that A is separable over $\mathbb{Z}$ if and only if A is separable over its prime subring.

Definition 6.2.46. *We say a ring is* separable *if it is separable as a $\mathbb{Z}$-algebra.*

Theorem 6.2.47. *There exists a deterministic polynomial-time algorithm that, given a finite ring A, decides whether or not it is separable.*

Proof. We begin by computing the prime subring of A using Theorem 3.3.1, and then use Theorem 5.4.1 to test separability over the prime subring. □

Suppose A_1, A_2 are two finite rings that are separable projective over their prime subrings. Then it is not necessarily true that $A_1 \times A_2$ will also be separable projective over its prime subring. However, the class of separable rings is closed under taking products by Proposition 6.2.16, part (iii). Being closed under taking products is also an important property of the class of semisimple rings. For a finite ring A, we would like to identify a subring $S \subset A$ such that A is separable over $\mathbb{Z}$ if and only if A is separable projective over S. That is the aim of this section.

Let A be a finite ring. Then A has a unique block decomposition, $A = \prod_{i \in I} A_i$, where I is the set of centrally primitive idempotents of A and each A_i is connected (see Definition 1.5.3, Theorem 1.5.5), and hence has prime power characteristic (see Theorem 1.5.5). We group together the A_i according to their characteristics to get

$$A = \prod_{i \in I} A_i = \prod_{\substack{p \text{ prime} \\ e \in \mathbb{Z}_{>0}}} \left(\prod_{\substack{i \\ \mathrm{char}(A_i) = p^e}} A_i \right). \tag{6.3}$$

Let

$$B_{p,e} := \prod_{\substack{i \\ \mathrm{char}(A_i) = p^e}} A_i.$$

Definition 6.2.48. *Let A be a finite ring. We define the* generalised prime subring *of A, denoted by $\mathcal{P}_A$, to be the product of the prime subrings of $B_{p,e}$.*

Proposition 6.2.49. *Let A be a finite ring and let k be its prime subring. Then $\mathcal{P}_A$ is separable as a k-algebra.*

Proof. The ring $\mathcal{P}_A$ is a product of rings, each of which is separable over k by Proposition 6.2.16, part (ii). The result now follows from Proposition 6.2.16, part (iii). □

Lemma 6.2.50. *Let A be a finite ring. Then $\mathcal{P}_A = \mathcal{P}_{\mathrm{Z}(A)}$.*

Proof. This follows since a block decomposition of A induces a block decomposition of $\mathrm{Z}(A)$ (see Theorem 1.5.6), together with the fact that any ring has the same prime subring as its centre. □

Proposition 6.2.51. *Let A be a finite separable ring. Then A is projective as a module over $\mathcal{P}_A$.*

Proof. By Proposition 6.2.16 we may assume that A is connected. Further, by Proposition 6.2.13 and Proposition 6.2.14, part (iii), we may assume that A is commutative. Hence A is local. Suppose A has prime subring $\mathbb{Z}/p^n\mathbb{Z}$, for some prime p and some $n \in \mathbb{Z}_{>0}$. Then A/pA is a separable $\mathbb{F}_p$-algebra. By Corollary 6.2.9, we have that A/pA is semisimple, so it must be a finite field extension of $\mathbb{F}_p$. Suppose the degree of this extension is d. We are left with showing that $A^+ \cong (\mathbb{Z}/p^n\mathbb{Z})^d$.

Consider the map $A \to p^{n-1}A$, given by $f : a \mapsto p^{n-1}a$. This map is surjective. Moreover, $pA \subseteq \ker(f)$. Since A/pA is a field, pA is a maximal ideal of A, so it must be the case that $pA = \ker(f)$. Hence $|A/pA| = |p^{n-1}A| = p^d$.

Since $A/pA \cong \mathbb{F}_p^d$ as $\mathbb{F}_p$-vector spaces, we may choose a basis $\{a_1, \ldots, a_d\} \subset A$ of A/pA over $\mathbb{F}_p$. Since $p^n A = 0$, it follows that $\{a_1, \ldots, a_d\}$ generate A over $\mathbb{Z}/p^n\mathbb{Z}$. Then the map $(\mathbb{Z}/p^n\mathbb{Z})^d \twoheadrightarrow A$, given by sending the generator of the i^{th} copy of $\mathbb{Z}/p^n\mathbb{Z}$ to a_i, is surjective. To see that it is also injective it is enough to show that the cardinalities agree. Consider the chain of ideals $A \supset pA \supset \ldots \supset p^{n-1}A \supset \{0\}$. Since $|p^i A/p^{i+1}A| = p^d$, for all $0 \leq i \leq n-1$, we have that $|A| = p^{nd}$. Hence $(\mathbb{Z}/p^n\mathbb{Z})^d \cong A$. □

Theorem 6.2.52. *Let A be a finite ring. Then A is separable if and only if A is separable projective over $\mathcal{P}_A$.*

Proof. Let k be the prime subring of A. The "only if" direction follows from Proposition 6.2.49 and the fact that A is separable over $\mathbb{Z}$ if and only if it is separable over k. The "if" direction follows from Proposition 6.2.14, parts (i), (ii) and Proposition 6.2.51. □

Note 6.2.53. If A is separable, but not finite, it is not necessarily projective over any proper subring. To see this, consider the $\mathbb{Z}$-algebra $\mathbb{Q}$. Then $\mathbb{Q}$ is certainly separable over $\mathbb{Z}$, since $\mathbb{Q} \otimes_{\mathbb{Z}} \mathbb{Q} \cong \mathbb{Q}$, but it cannot be projective over any proper subring, since then it would have to be finitely generated over it (see Proposition 6.2.13).

Suppose $R \subseteq \mathbb{Q}$ is a subring such that $\mathbb{Q}$ is finitely generated as a module over R. Then $\mathbb{Q}$ is integral over R, i.e. for every $q \in \mathbb{Q}$, there exists a monic polynomial $f \in R[X]$ such that $f(q) = 0$. But then it is easy to see that $\mathbb{Q}^* \cap R = R^*$, where R^* denotes the unit group of R. Hence $R^* = R \backslash \{0\}$, so R is a field and then it must be the case that $R = \mathbb{Q}$.

6.2.6 Classification of finite separable rings

For finite-dimensional semisimple algebras over a field and separable algebras over a field, we have explicit descriptions in terms of matrix rings over certain division rings (see Theorem 1.4.5 and Theorem 6.2.8, part (iii), respectively). The aim of this section is to provide a similar classification result for finite separable rings. It turns out that finite separable rings are isomorphic to products of finitely many matrix algebras over certain special commutative rings, called *Witt rings*.

Recall that a ring C is said to be *connected* (or *indecomposable*) if C has exactly two central idempotents, namely 0 and 1. Note that a connected ring is nonzero.

Here, we develop the theory concerning Witt rings that we require. We restrict our attention to *truncated Witt rings over finite fields*, as this is a sufficient level of generality for our purposes. For more on Witt rings, see Chapter II: §6 of [82] or Chapter VI: Exercises 46-51 of [60].

Let p be a prime and $e \in \mathbb{Z}_{>0}$. Let

$$\underline{\mathbf{C}} = \text{category of finite local commutative } \mathbb{Z}/p^e\mathbb{Z}\text{-algebras},$$
$$\underline{\mathbf{D}} = \text{category of finite fields of characteristic } p,$$

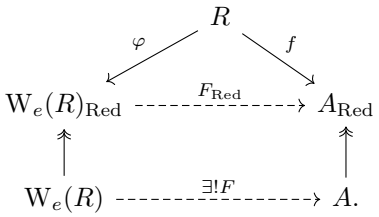
and consider the covariant functor

$$\begin{aligned} \mathrm{Red} : \underline{\mathbf{C}} &\longrightarrow \underline{\mathbf{D}} \\ A &\longmapsto A_{\mathrm{Red}} := A/\sqrt{0_A}, \end{aligned} \tag{6.4}$$

where $\sqrt{0_A}$ denotes the nilradical of A (which equals the maximal ideal of A).

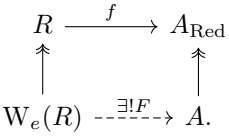
Theorem 6.2.54. *The functor* Red *has a left adjoint.*

Proof. We will show that if $R \in \underline{\mathbf{D}}$, then there exists a pair $(\mathrm{W}_e(R), \varphi)$, with $\mathrm{W}_e(R) \in \underline{\mathbf{C}}$ and $\varphi : R \to \mathrm{W}_e(R)_{\mathrm{Red}}$ a ring homomorphism, such that for every $A \in \underline{\mathbf{C}}$ and every ring homomorphism $R \xrightarrow{f} A_{\mathrm{Red}}$, there exists a unique ring homomorphism $F : \mathrm{W}_e(R) \xrightarrow{F} A$ such that $f = F_{\mathrm{Red}} \circ \varphi$, i.e. such that the following diagram commutes:



Moreover, φ is an isomorphism, and the pair $(\mathrm{W}_e(R), \varphi)$ is unique up to unique isomorphism.

We may assume that $R = \mathbb{F}_p[X]/(\overline{g}(X))$, where $g(X) \in \mathbb{Z}/p^e\mathbb{Z}$ is a monic polynomial of degree d, for some $d \in \mathbb{Z}_{>0}$, and $\overline{g}(X) := (g(X) \mod p) \in \mathbb{F}_p[X]$ is irreducible. Let $\mathrm{W}_e(R) = (\mathbb{Z}/p^e\mathbb{Z})[X]/(g(X))$. Then $\mathrm{W}_e(R)_{\mathrm{Red}} \cong R$. Let $\varphi : R \xrightarrow{\sim} \mathrm{W}_e(R)_{\mathrm{Red}}$ be the natural isomorphism. We are left with showing that for all $A \in \underline{\mathbf{D}}$ and $f : R \to A_{\mathrm{Red}}$, there is a unique $F : \mathrm{W}_e(R) \to A$, making the following diagram commute:



This follows from Hensel's lemma. The last part follows by properties of universal objects.

This shows that there is a functor

$$\begin{aligned} \mathrm{W}_e : \underline{\mathbf{D}} &\longrightarrow \underline{\mathbf{C}} \\ R &\longmapsto \mathrm{W}_e(R) \end{aligned}$$

that is left adjoint to Red. □

Definition 6.2.55. *Let p be a prime and $d, e \in \mathbb{Z}_{>0}$. The* e-truncated Witt ring *of $\mathbb{F}_{p^d}$ is defined to be $\mathrm{W}_e(\mathbb{F}_{p^d})$.*

Note 6.2.56. The ring $\mathrm{W}_e(\mathbb{F}_{p^d})$ has cardinality p^{de}.

Example 6.2.57. We have $\mathrm{W}_1(\mathbb{F}_{p^d}) = \mathbb{F}_{p^d}$ and $\mathrm{W}_e(\mathbb{F}_p) = \mathbb{Z}/p^e\mathbb{Z}$.

From the above construction, we have the following result.

Proposition 6.2.58. *Let p be a prime and $e, d \in \mathbb{Z}_{>0}$. Then the ring $\mathrm{W}_e(\mathbb{F}_{p^d})$ is local, with maximal ideal $p\,\mathrm{W}_e(\mathbb{F}_{p^d})$, which has cardinality $p^{d(e-1)}$. Moreover, the set of ideals of $\mathrm{W}_e(\mathbb{F}_{p^d})$ is $\{p^i\,\mathrm{W}_e(\mathbb{F}_{p^d}) \mid 0 \leq i \leq e\}$.*

Note 6.2.59. Another way of constructing Witt rings is via Galois theory for commutative rings. In the case of finite rings, the notion of a Witt ring is then replaced by that of a *Galois ring*. This is the more common terminology in literature on separability (see [9, 20]).

Proposition 6.2.60. *Let R be a finite commutative ring. Then R is local separable if and only if $R \cong \mathrm{W}_e(\mathbb{F}_{p^d})$, for some prime p and some $e, d \in \mathbb{Z}_{>0}$.*

Proof. For the "only if" direction, suppose that R is local separable. Then it has prime power characteristic, p^e, for some prime p and some $e \in \mathbb{Z}_{>0}$. By Theorem 6.2.18, part (iii), we have that R is separable over $\mathbb{Z}/p^e\mathbb{Z}$ if and only if R/pR is separable over $\mathbb{F}_p$. By Corollary 6.2.9, we have that R/pR is semisimple. Hence $R/pR \cong \mathbb{F}_{p^d}$. We have the following commutative diagram:

$$\begin{array}{ccc} R & \twoheadrightarrow & \mathbb{F}_{p^d} \\ \uparrow & & \uparrow \\ \mathbb{Z}/p^e\mathbb{Z} & \twoheadrightarrow & \mathbb{F}_p. \end{array}$$

By the proof of Theorem 6.2.54, we have that $R \cong \mathrm{W}_e(\mathbb{F}_{p^d})$.

The "if" direction follows by construction of truncated Witt rings. □

We can now turn to the classification of finite separable rings. By Proposition 6.2.16, part (ii) it suffices to classify finite connected separable rings.

Proposition 6.2.61. *Let $\mathcal{P}$ be the set of primes. There is a bijection between the sets*

$$\{\text{finite commutative local separable rings}\}/\cong \quad \longleftrightarrow \quad \mathcal{P} \times \mathbb{Z}_{>0} \times \mathbb{Z}_{>0},$$

where

$$[A] \longmapsto (p, d, e), \tag{6.5}$$

if for $\mathfrak{m}$, the maximal ideal of A, we have

$$\begin{aligned} \mathrm{char}(A/\mathfrak{m}) &= p, \\ [A/\mathfrak{m} : \mathbb{F}_p] &= d, \\ \mathrm{char}(A) &= p^e. \end{aligned}$$

The inverse of the map in (6.5) is given by

$$\mathrm{W}_e(\mathbb{F}_{p^d}) \longleftarrow\!\shortmid (p, d, e). \tag{6.6}$$

Proof. This is a consequence of Theorem 6.2.54, Definition 6.2.55 and Proposition 6.2.60. □

Theorem 6.2.62. *Let $\mathcal{P}$ denote the set of primes. Then there is a bijection between the sets*

$$\{\text{finite connected separable rings}\}/\cong \quad \longleftrightarrow \quad \mathcal{P} \times \mathbb{Z}_{>0} \times \mathbb{Z}_{>0} \times \mathbb{Z}_{>0},$$

where

$$[A] \longmapsto (p, d, e, n), \tag{6.7}$$

if

$$\begin{aligned} |A| &= p^{den^2}, \\ |\operatorname{Z}(A)| &= p^{de}, \\ |\operatorname{Z}(A)/p\operatorname{Z}(A)| &= p^d. \end{aligned}$$

The inverse of the map in (6.7) is given by

$$\mathcal{M}_n(\mathrm{W}_e(\mathbb{F}_{p^d})) \longleftarrow\!\shortmid (p, d, e, n). \tag{6.8}$$

Proof. By Theorem 6.2.52, a finite connected ring is separable if and only if it is separable projective over its prime subring, which by Theorem 6.2.36 is equivalent to being Azumaya over its centre and its centre being separable over its prime subring. Given Proposition 6.2.61, we are thus left with classifying Azumaya algebras over truncated Witt rings.

The map given in (6.7) is well-defined, since $\mathrm{W}_e(\mathbb{F}_{p^d})$ is a local ring, and hence the degree of any Azumaya $\mathrm{W}_e(\mathbb{F}_{p^d})$-algebra is well-defined and is a square (Corollary 6.2.26). Injectivity follows from Example 6.2.34, part 3. Surjectivity follows from the fact that matrix rings over commutative rings are separable (Example 6.2.11).

Note that an Azumaya algebra over a commutative local ring is free as a module over that ring (since projective modules over local rings are free). The fact that the maps (6.7) and (6.8) are mutual inverses now follows from Proposition 6.2.61, Theorem 6.2.20, part (ii), and Example 6.2.34, part 3. □

Corollary 6.2.63. *Let A be a finite separable ring. Then $A \cong A^{\mathrm{o}}$ as rings.*

We know from Theorem 1.2.11 that for R a commutative ring, $n \in \mathbb{Z}_{>0}$ and $S = \mathcal{M}_n(R)$, every two-sided ideal of S is of the form $\mathcal{M}_n(I)$, for some two-sided ideal I of R. Conversely, if I is a two-sided ideal of R, then $\mathcal{M}_n(I)$ is a two-sided ideal of S.

Corollary 6.2.64. *The set of two-sided ideals of a finite connected separable ring is in bijection with the set of ideals of its prime subring.*

Proof. Recall from Proposition 6.2.58 that all ideals of $\mathrm{W}_e(\mathbb{F}_{p^d})$, where p is a prime, are generated by powers of p. In particular, there are exactly $e+1$ ideals, with maximal ideal generated by p. The result now follows from Theorem 6.2.62. □

6.2.7 The trace map and the trace radical

In this section we introduce the notions of trace map and trace ideal, which are closely related to separability.

Let k be a commutative ring and P a finitely generated k-module. Consider the map

$$\mathrm{Hom}_k(P,k) \otimes_k P \xrightarrow{\varphi} \mathrm{End}_k(P), \quad f \otimes x \mapsto (y \mapsto f(y)x). \tag{6.9}$$

Then φ induces a map

$$\begin{aligned} \square : (\mathrm{Hom}_k(P,k) \otimes_k P) \times (\mathrm{Hom}_k(P,k) \otimes_k P) &\longrightarrow \mathrm{Hom}_k(P,k) \otimes_k P, \\ (f \otimes x, g \otimes y) &\longmapsto f(y)g \otimes x, \end{aligned} \tag{6.10}$$

which makes the following diagram commute:

$$\begin{array}{ccc} \mathrm{End}_k(P) \times \mathrm{End}_k(P) & \xrightarrow{\text{composition}} & \mathrm{End}_k(P) \\ {\scriptstyle \varphi\times\varphi}\uparrow & & \uparrow{\scriptstyle \varphi} \\ (\mathrm{Hom}_k(P,k) \otimes_k P) \times (\mathrm{Hom}_k(P,k) \otimes_k P) & \xrightarrow{\square} & \mathrm{Hom}_k(P,k) \otimes_k P. \end{array} \tag{6.11}$$

Recall the definition of a dual basis.

Lemma 6.2.65 ([57], Lemma 2.9, Remark 2.11)**.** (Dual Basis Lemma) *Let k be a ring and let P be a k-module. Then*

(i) *P is projective if and only if there exists a collection $\{x_i, f_i\}_{i\in I}$ for some index set I, with $x_i \in P$ and $f_i \in \mathrm{Hom}_k(P,k)$ such that*

$$\forall x \in P: \ f_i(x) = 0 \text{ for almost all } i, \text{ and } x = \sum_{i=1}^{n} f_i(x)x_i.$$

(ii) *P is finitely generated projective if and only if there exist $n \in \mathbb{Z}_{>0}$ and a collection $\{x_i, f_i\}_{i=1}^n$ with $x_i \in P$ and $f_i \in \mathrm{Hom}_k(P,k)$ such that*

$$\forall x \in P: \quad x = \sum_{i=1}^{n} f_i(x)x_i,$$

i.e.

$$\varphi\left(\sum_{i=1}^{n} f_i \otimes x_i\right) = \mathrm{id}_P,$$

where φ is the map defined in (6.9).

Note 6.2.66. Part (ii) above is equivalent to requiring $1 \in \mathrm{im}(\varphi)$, where φ is the map defined in (6.9).

Definition 6.2.67. *Let k be a ring and P a projective k-module. A collection $\{x_i, f_i\}_i$ as described in the previous lemma is called a* dual basis *of P.*

Lemma 6.2.68. *Let k be a commutative ring, M a k-module and N a finitely generated projective k-module. Then*

$$\begin{aligned}\psi : \mathrm{Hom}_k(M,k) \otimes_k N &\to \mathrm{Hom}_k(M,N),\\ f \otimes x &\mapsto (y \mapsto f(y)x)\end{aligned}$$

is a k-module isomorphism.

Proof. If $N = k$, then ψ is clearly an isomorphism. Then, by properties of tensor products and Hom's, for any $n \in \mathbb{Z}_{>0}$, we have that $N = k^n$ also gives an isomorphism. But then ψ remains an isomorphism for N a finitely generated projective k-module, since then N is a direct summand of some multiple of k. □

Note 6.2.69. Lemma 6.2.68 remains true if it is M that is finitely generated projective as a k-module (the proof of this follows the same lines).

Now let P be a finitely generated projective k-module. Then the map φ as defined in (6.9) is an isomorphism. Further, if we consider the map

$$\mathrm{Hom}_k(P,k) \otimes_k P \xrightarrow{\psi} k, \quad f \otimes x \mapsto f(x), \tag{6.12}$$

then we get an induced map $\mathrm{tr}_{P/k} := \psi\varphi^{-1}$:

$$\begin{array}{ccc} \mathrm{Hom}_k(P,k) \otimes_k P & \xrightarrow{\sim} & \mathrm{End}_k(P). \\ \downarrow & \swarrow_{\mathrm{tr}_{P/k}} & \\ k & & \end{array} \tag{6.13}$$

Definition 6.2.70. *Let k be a commutative ring and P a finitely generated projective k-module. Then the map $\mathrm{tr}_{P/k}$ of diagram (6.13) is called the* trace *of P over k. The quantity $\mathrm{rk}_{P/k} := \mathrm{tr}_{P/k}(\mathrm{id})$ is called the* Hattori-Stallings rank *of P over k.*

Proposition 6.2.71. *Let k be a commutative ring and P a finitely generated k-module. Then P is projective if and only if $\mathrm{Hom}_k(P,k) \otimes_k P$ is a ring with multiplication given by $\square$, as defined in (6.10).*

Proof. If P is finitely generated projective, then $\mathrm{Hom}_k(P,k) \otimes_k P \cong \mathrm{End}_k(P)$, so it is a ring and the map $\square$ is simply composition of maps by diagram (6.11).

For the other direction, note that $\mathrm{Hom}_k(P,k) \otimes_k P$ being a ring implies the existence of an element $\alpha = \sum_{i=1}^n f_i \otimes x_i \in \mathrm{Hom}_k(P,k) \otimes_k P$ such that for all $\beta \in \mathrm{Hom}_k(P,k) \otimes_k P$, we have $\square(\alpha, \beta) = \beta$. But then for any $x \in P$,

$$\mathrm{id} \otimes x = \square(\sum_{i=1}^n f_i \otimes x_i, \mathrm{id} \otimes x) = \sum_{i=1}^n f_i(x)\, \mathrm{id} \otimes x_i,$$

Applying ψ (as defined in (6.12))to both sides, we get

$$x = \sum_{i=1}^{n} f_i(x)x_i.$$

Hence $\{x_i, f_i\}_{i=1}^n$ is a dual basis of P and so P is finitely generated projective by Lemma 6.2.65. □

The following result says that the trace map $\mathrm{tr}_{P/k}$ behaves "as expected".

Proposition 6.2.72 ([7], Section 1)**.** *Let k be a commutative ring and P a finitely generated projective k-module.*

(i) *Let $e_1, e_2 \in \mathrm{End}_k(P)$. Then $\mathrm{tr}_{P/k}(e_1 + e_2) = \mathrm{tr}_{P/k}(e_1) + \mathrm{tr}_{P/k}(e_2)$.*
(ii) *Let $e \in \mathrm{End}_k(P)$ and $c \in k$. Then $\mathrm{tr}_{P/k}(ce) = c\,\mathrm{tr}_{P/k}(e)$.*
(iii) *Let $e_1, e_2 \in \mathrm{End}_k(P)$. Then $\mathrm{tr}_{P/k}(e_1 \circ e_2) = \mathrm{tr}_{P/k}(e_2 \circ e_1)$.*
(iv) (The trace is compatible with base change) *Let k' be a commutative ring and $\alpha : k \to k'$ a ring homomorphism. Let $P' = P \otimes_k k'$ and $e \in \mathrm{End}_k(P)$. Then*

$$\mathrm{tr}_{P'/k'}(e \otimes_k 1_{k'}) = \alpha(\mathrm{tr}_{P/k}(e)).$$

Let us now consider the case when P is in fact a k-algebra.

Definition 6.2.73. *Let k be a commutative ring and let A be a k-algebra that is finitely generated and projective as a k-module. Then the map $A \to \mathrm{End}_k(A)$ given by $a \mapsto (x \mapsto ax)$ induces a map $\mathrm{Tr}_{A/k} : A \to k$ in the following diagram:*

$$\begin{array}{ccccc} \mathrm{Hom}_k(A,k) \otimes_k A & \xrightarrow{\sim} & \mathrm{End}_k(A) & \longleftarrow & A. \\ \downarrow & & \nearrow_{\mathrm{tr}_{A/k}} & \swarrow_{\mathrm{Tr}_{A/k}} & \\ k & & & & \end{array}$$

We call $\mathrm{Tr}_{A/k}$ the trace map *of A over k.*

Note 6.2.74. It is easy to see that $\mathrm{Tr}_{A/k}$ is a k-module homomorphism. Moreover, by Proposition 6.2.72, part (iii), for all $a, b \in A$, we have $\mathrm{Tr}_{A/k}(ab) = \mathrm{Tr}_{A/k}(ba)$.

We give a second definition of the trace map, using a more element-oriented approach.

Definition 6.2.75. *Let k be a commutative ring and A a k-algebra that is finitely generated projective as a k-module. Let $\{x_i, f_i\}$ be a dual basis of A over k. We define*

$$\mathrm{tr}_{A/k} : \mathrm{End}_k(A) \to k, \quad g \mapsto \sum_{i=1}^{n} f_i(g(x_i))$$

and

$$\mathrm{Tr}_{A/k} : A \to k, \quad r \mapsto \sum_{i=1}^{n} f_i(rx_i).$$

Note 6.2.76. The above definition is independent of the choice of dual basis. One way to see this is the following proposition.

Proposition 6.2.77. *The two definitions of* $\mathrm{Tr}_{A/k}$ *agree.*

Proof. Let $\{a_i, f_i\}_{i=1}^n$ be a dual basis of A as a finitely generated projective module over k. Let $\mathrm{Tr}^{(1)}_{A/k}$ be the trace map as in Definition 6.2.73 and $\mathrm{Tr}^{(2)}_{A/k}$, the trace map as in Definition 6.2.75. For any $a \in A$, consider the two trace maps:

$$A \longrightarrow \mathrm{End}_k(A) \longrightarrow k$$

$$\mathrm{Tr}^{(1)}_{A/k}: \qquad a \longmapsto e_a := (x \mapsto ax) \longmapsto \mathrm{tr}_{A/k}(e_a)$$

$$\mathrm{Tr}^{(2)}_{A_k}: \qquad a \longmapsto \sum_{i=1}^n f_i(aa_i).$$

Consider the element $\sum_{i=1}^n f_i \otimes aa_i \in \mathrm{Hom}_k(A,k) \otimes_k A$. This maps to

$$(y \mapsto \sum_{i=1}^n f_i(y)aa_i = \sum_{i=1}^n af_i(y)a_i = ay) = e_a$$

under the isomorphism $\varphi : \mathrm{Hom}_k(A,k) \otimes_k A \xrightarrow{\sim} \mathrm{End}_k(A)$, which in turn maps to $\mathrm{tr}_{A/k}(a) = \mathrm{Tr}^{(1)}_{A/k}(a) \in k$. Also, it maps to $\sum_{i=1}^n f_i(aa_i) \in k$ under $\psi : \mathrm{Hom}_k(A,k) \otimes_k A \to k$. Hence $\mathrm{Tr}^{(1)}_{A/k}(a) = \mathrm{tr}_{A/k}(e_a) = \sum_{i=1}^n f_i(aa_i) = \mathrm{Tr}^{(2)}_{A/k}(a)$. □

Example 6.2.78. Let k be a commutative ring and let A be a k-algebra such that $A \cong k^n$ as k-modules, for some $n \in \mathbb{Z}_{>0}$. Let $\mathcal{B} = \{b_i \mid 1 \le i \le n\}$ be a basis of A over k. Then a dual basis of A over k is given by $\{b_i, f_i\}_{i=1}^n$, where $f_j : \sum_{i=1}^n a_i b_i \mapsto a_j$. It is easy to see that $\mathrm{Tr}_{A/k}(1) = n \cdot 1$.

Example 6.2.79. Let k be a commutative ring and let $A = \mathcal{M}_n(k)$. Then

$$\mathrm{Tr}_{A/k} = n \cdot (\text{usual trace}).$$

Example 6.2.80. Let A be a finite-dimensional algebra over a field k. Then nilpotent elements of A have trace zero. This is because nilpotent matrices over a field have trace zero.

Example 6.2.81. Let A be a finite-dimensional algebra over a field $\mathbb{F}_p$, where p is a prime and let S be a finite A-module. Since S is a vector space over $\mathbb{F}_p$, we have a ring homomorphism $\rho : A \to \mathrm{End}_{\mathbb{F}_p}(S)$ given by sending an element $a \in A$ to the endomorphism of S corresponding to the action of a on S. Define $\mathrm{Tr}^{(S)} : A \to \mathbb{F}_p$ to be the map $\mathrm{tr}_{S/\mathbb{F}_p} \circ \rho$, where $\mathrm{tr}_{S/\mathbb{F}_p}$ is the trace of S over $\mathbb{F}_p$. Note that $\mathrm{Tr}^{(A)} = \mathrm{Tr}_{A/\mathbb{F}_p}$

is the usual trace map, as defined in 6.2.73. If $0 \to S \to T \to U \to 0$ is an exact sequence of A-modules, then

$$\mathrm{Tr}^{(T)} = \mathrm{Tr}^{(S)} + \mathrm{Tr}^{(U)}.$$

To see this, suppose C_1, C_2 are bases for S and U respectively. Then the matrix of $\mathrm{Tr}^{(T)}$ can be represented as an upper triangular block matrix, where the two diagonal blocks are the matrices of $\mathrm{Tr}^{(S)}$ and $\mathrm{Tr}^{(U)}$ with respect to C_1 and C_2 respectively. Moreover, if a basis of T contains a basis of S, then the rest is a basis of U.

Definition 6.2.82. *Let k be a commutative ring and A a k-algebra that is finitely generated projective as a k-module. The* trace radical *of A over k is the kernel of the right A-module homomorphism:*

$$\psi : A \to \mathrm{Hom}_k(A, k), \quad a \mapsto \mathrm{Tr}_{A/k} \cdot a := (x \mapsto \mathrm{Tr}_{A/k}(ax)). \tag{6.14}$$

In other words,

$$I_{A/k} := \{a \in A \mid \mathrm{Tr}_{A/k}(aA) = 0\}.$$

Note 6.2.83.

(i) By Proposition 6.2.72, part (iii), the trace radical is a two-sided ideal.
(ii) By Proposition 6.2.72, part (iii), if $\mathrm{Tr}_{A/k}$ generates $\mathrm{Hom}_k(A, k)$ as a right A-module, then it generates it as a left A-module.
(iii) Suppose k is a field. Since $\dim_k(A) = \dim_k(\mathrm{Hom}_k(A, k))$, we have that $A \cdot \mathrm{Tr}_{A/k} = \mathrm{Hom}_k(A, k)$ if and only if $I_{A/k} = 0$.

Lemma 6.2.84. *Let A be a finite ring. Suppose $A = \prod_{i=1}^{l} A_i$, for some $l \in \mathbb{Z}_{>0}$ and A_i finite rings. Let $n_i := \mathrm{char}(A_i)$ and suppose that for all $i \neq j$ we have* $\gcd(n_i, n_j) = 1$ *and that each A_i is free as a module over $\mathbb{Z}/n_i\mathbb{Z}$. Then*

$$I_{A/(\mathbb{Z}/\mathrm{char}(A)\mathbb{Z})} = \prod_i I_{A/(\mathbb{Z}/n_i\mathbb{Z})}.$$

Proof. Write $k := \mathbb{Z}/\mathrm{char}(A)\mathbb{Z}$. First note that $\mathrm{char}(A) = \prod_i n_i$ and that A is indeed projective over k, so that $I_{A/k}$ is well-defined. The result now follows from Proposition 6.2.72, part (iv). □

Theorem 6.2.85. *There exists a deterministic polynomial-time algorithm that, given a finite commutative ring k and a finite k-algebra A that is projective as a k-module, computes the trace radical $I_{A/k}$.*

Proof. We begin by computing $\mathrm{Hom}_k(A, k)$, using Proposition 2.4.1. Then $I_{A/k}$ is computed as the kernel of the map ψ from Definition 6.2.82. □

6.2.8 Strongly separable algebras

We now study the connections between the trace radical and separability.

Proposition 6.2.86 ([64], Proposition 6.11)**.** *Let k be a commutative ring and R a commutative k-algebra. Then R is finite-étale over k if and only if R is finitely generated projective as a module over k and $\mathrm{Tr}_{R/k}$ generates $\mathrm{Hom}_k(R,k)$.*

Note 6.2.87. This is usually taken to be the definition of finite-étale.

Let k be a commutative ring. We would like to characterise k-algebras A that are finitely generated and projective as k-modules and have the property that $\mathrm{Tr}_{A/k}$ generates $\mathrm{Hom}_k(A,k)$ as a right A-module, but are not necessarily commutative.

Theorem 6.2.88 ([21], Theorem 1, [49], Theorem 3.4)**.** *Let k be a commutative ring and let A be a k-algebra with centre $R := \mathrm{Z}(A)$ such that A is a finitely generated projective k-module. Then the following are equivalent:*

(i) *The trace map $\mathrm{Tr}_{A/k}$ generates $\mathrm{Hom}_k(A,k)$ as a right A-module.*
(ii) *A is k-separable and $A = R \oplus [A,A]$ as R-modules, where $[A,A]$ is the R-submodule of A generated by elements of the form $ab - ba$, with $a, b \in A$.*
(iii) *A is k-separable and $\mathrm{Tr}_{A/R}(1)$ is a unit in $k \cdot 1_A$, the image of k in R.*
(iv) *A has a symmetric separability idempotent over k.*

Definition 6.2.89. *An algebra satisfying any of the conditions of Theorem 6.2.88 is called a* strongly separable *algebra.*

Note 6.2.90. We see that strongly separable algebras are a special kind of symmetric algebras, namely ones for which a nonsingular, symmetric, associative bilinear map $B : A \times A \to k$ is given by $B(a,b) = \mathrm{Tr}_{A/k}(ba)$. We have seen in Theorem 6.2.45 that any separable algebra that is finitely generated, projective and faithful as a module over its base ring is symmetric, but the trace map need not be nonsingular, and thus may not give rise to such a map B.

Example 6.2.91. Let k be a finite field and let $A = \mathcal{M}_n(k)$. Suppose $\mathrm{char}\, k$ divides n. Then A is separable (and symmetric) over k, but $\mathrm{Tr}_{A/k} = n \cdot (\text{the usual trace}) = 0$, so A is not strongly separable over k. To see that A is a symmetric k-algebra, we must look at the usual trace map, which we denote by tr_0. Consider the map $B : A \times A \to k$ given by $B(a,b) = \mathrm{tr}_0(ab)$. This is now bilinear, symmetric, associative and nonsingular, as required for it to witness the fact that A is symmetric as a k-algebra.

Example 6.2.92. (Strongly separable algebras)

1. Let $n \in \mathbb{Z}_{>0}$ and k be a commutative ring. If $n \cdot 1$ is a unit in k, then $\mathcal{M}_n(k)$ is strongly separable over k with symmetric separability idempotent $n^{-1} \sum_{i,j=1}^{n} E_{ij} \otimes E_{ji}$, where E_{ij} denotes the $n \times n$ matrix whose $(i,j)^{\text{th}}$ entry is equal to 1 and all other entries are equal to 0.

2. Let $n \in \mathbb{Z}_{>0}$ and k be a finite commutative ring. Put $A = \mathcal{M}_n(k)$. If $n \cdot 1$ is not a unit in k, then $n \cdot 1$ is a zero-divisor. Since $\mathrm{Tr}_{A/k} = n \cdot$ (usual trace), we have that $I_{A/k} \neq 0$ and hence A is not strongly separable.
3. Let G be a finite group, k a commutative ring, and put $A := kG$. If $|G|$ is a unit in k, then A is strongly separable over k, with symmetric separability idempotent $|G|^{-1} \sum_{g \in G} g \otimes g^{-1}$.
4. ([3], Corollary 3.1) Let k be a field with $\mathrm{char}(k) = 0$ and A a k-algebra. Then A is strongly separable if and only if it is finite-dimensional and semisimple.

6.3 An approximation of the Jacobson radical

We have seen in Sections 6.2.4 and 6.2.5 that separable projective algebras and separable rings have many nice properties. Until now, however, we have mainly stayed on theoretical ground. We would now like to be able to algorithmically reduce any finite ring to this "state". In other words, given a finite ring, we would like to quotient out by some two-sided ideal and obtain a ring that is separable. Our goal in a perfect world would have been to quotient out by the Jacobson radical and obtain a semisimple ring. Since computing the Jacobson radical is in general out of our reach (see Note 3.4.2), we will have to content ourselves with quotienting out by something that is *almost* the Jacobson radical and obtaining something that is *almost* semisimple, more precisely, something that is separable.

6.3.1 Defining an approximation

Definition 6.3.1. *Let A be a finite ring and $\mathfrak{j}_A \subset A$ an ideal. We say $\mathfrak{j}_A$ is an* approximation of the Jacobson radical *of A if*

(A1) *$\mathfrak{j}_A$ is a two-sided nilpotent ideal of A.*
(A2) *$A/\mathfrak{j}_A$ is finite separable.*
(A3) *The prime subring and generalised prime subring of $A/\mathfrak{j}_A$ coincide.*

Note 6.3.2. Let A be a finite ring. Then by Theorem 1.4.9, Proposition 6.2.43 and Theorem 6.2.52, the Jacobson radical is an approximation of itself.

Note 6.3.3. Approximations of Jacobson radicals are not unique. Let p be a prime and let $A = \mathbb{Z}/p^2\mathbb{Z}$. Then A is finite separable with prime subring and generalised prime subring equal to $\mathbb{Z}/p^2\mathbb{Z}$. Hence 0 and $\mathrm{J}(A) = p\mathbb{Z}/p^2\mathbb{Z}$ are both approximations of the Jacobson radical of A.

Theorem 6.3.4. *Let A be a finite ring and $\mathfrak{j}_A$ a two-sided ideal of A such that $\mathfrak{j}_A$ is nilpotent and $A/\mathfrak{j}_A$ is separable projective over its prime subring. Suppose, moreover, that the characteristic of A is a power of some prime p. Then*

$$(A/\mathfrak{j}_A)/(p(A/\mathfrak{j}_A)) = A/\mathrm{J}(A),$$

and

$$(A/\mathfrak{j}_A)^+ \cong (\mathbb{Z}/p^e\mathbb{Z})^r,$$

where $r = \dim_{\mathbb{F}_p}(A/\operatorname{J}(A))$ *and* $e \in \mathbb{Z}_{>0}$ *is such that* $p^e = \operatorname{char}(A/\mathfrak{j}_A)$.

Proof. Let $e \in \mathbb{Z}_{>0}$ and p be a prime such that $\operatorname{char}(A/\mathfrak{j}_A) = p^e$. First note that, since $A/\mathfrak{j}_A$ is separable over $\mathbb{Z}/p^e\mathbb{Z}$, we have that $(A/pA)/((\mathfrak{j}_A + pA)/pA)$ is semisimple by Corollary 6.2.10 and Theorem 6.2.18, part (iii). This implies that $(\mathfrak{j}_A + pA)/pA \supseteq \operatorname{J}(A/pA)$. Moreover, $\mathfrak{j}_A$ is nilpotent. We thus have that $(\mathfrak{j}_A + pA)/pA = \operatorname{J}(A/pA)$. Hence

$$\begin{aligned}(A/\mathfrak{j}_A)/(p(A/\mathfrak{j}_A)) &= (A/pA)/((\mathfrak{j}_A + pA)/pA)\\ &= (A/pA)/\operatorname{J}(A/pA)\\ &= A/\operatorname{J}(A).\end{aligned}$$

By Nakayama's Lemma, the minimum number of generators of $A/\mathfrak{j}_A$ as a $\mathbb{Z}/p^e\mathbb{Z}$-module is equal to the dimension of $A/\operatorname{J}(A)$ over $\mathbb{F}_p$. But $A/\mathfrak{j}_A$ is projective over $\mathbb{Z}/p^e\mathbb{Z}$, so it is free of finite rank. Thus the rank of $A/\mathfrak{j}_A$ as a $\mathbb{Z}/p^e\mathbb{Z}$-module is equal to $\dim_{\mathbb{F}_p}(A/\operatorname{J}(A))$.

□

Example 6.3.5. Let p be a prime and M an $\mathbb{F}_p$-vector space of dimension 1. Let

$$A = \mathbb{Z}/p^2\mathbb{Z} \oplus M$$

be the ring with componentwise addition and multiplication given by

$$(a, x) \cdot (b, y) = (ab, ay + bx).$$

In particular, A is a commutative ring with $M^2 = 0$. Moreover, $\operatorname{J}(A) = p\mathbb{Z}/p^2\mathbb{Z} \oplus M$.

For any approximation $\mathfrak{j}$ of the Jacobson radical of A we must have $A/\mathfrak{j} \cong \mathbb{Z}/p^2\mathbb{Z}$ or $A/\mathfrak{j} \cong \mathbb{F}_p$. If $A/\mathfrak{j} \cong \mathbb{F}_p$, then it must be the case that $\mathfrak{j} = \operatorname{J}(A)$.

Let S be the set of all approximations of the Jacobson radical of A. We have bijections between the following sets

$$\begin{aligned}S\backslash\{\operatorname{J}(A)\} &\longleftrightarrow \{\text{ring homomorphisms } A \to \mathbb{Z}/p^2\mathbb{Z}\}\\ &\longleftrightarrow \{\text{group homomorphisms } M \to p\mathbb{Z}/p^2\mathbb{Z}\}.\end{aligned}$$

The latter set has p elements and each of these gives rise to an approximation of the Jacobson radical of A of the same size.

Note 6.3.6. Example 6.3.5 shows that, even though the set of all approximations of the Jacobson radical of a finite ring always has a maximal element with respect to inclusion, given by the Jacobson radical, it does not necessarily have a minimal element.

The aim of this section is to describe deterministic polynomial-time algorithms that produce approximations of the Jacobson radical of a finite ring. We are interested in algorithms that have the additional property that, when run on two isomorphic rings, they output isomorphic approximations of their Jacobson radicals (induced by the same isomorphism), even when the ring isomorphism is unknown (cf. Section 3.5).

We will treat rings in a differentiated manner, depending on the size of the primes dividing their characteristic. We define convenient notions of "small" and "large" primes that allow us to split the ring into two parts and deal with them separately. The case of small primes is easy to deal with, since we can actually compute the Jacobson radical and thus arrange for genuine semisimplicity. The case of large primes requires more work and what allows us to deal with them is Theorem 6.2.88, part (iii).

6.3.2 Trace radical vs. Jacobson radical

We start by proving a series of results about the trace ideal (see Definition 6.2.73) that we will make use of within our algorithms.

Proposition 6.3.7. *Let A be a finite-dimensional algebra over a field k. Then the trace ideal $I_{A/k}$ contains the Jacobson radical $\mathrm{J}(A)$.*

Proof. Since A is left-artinian, its Jacobson radical is nilpotent and so by Example 6.2.80, all its elements lie in $I_{A/k}$. □

Theorem 6.3.8. *Let A be a finite-dimensional algebra over the finite field $\mathbb{F}_p$, where p is a prime and $p > \dim_{\mathbb{F}_p}(A)$. Then $I_{A/\mathbb{F}_p} = \mathrm{J}(A)$.*

Proof. The inclusion $\mathrm{J}(A) \subseteq I_{A/\mathbb{F}_p}$ is given by Proposition 6.3.7. For the other inclusion, use Example 6.2.81 and induction to write

$$\mathrm{Tr}_{A/\mathbb{F}_p} = \sum_{\substack{S \text{ simple} \\ \text{up to } \cong}} \mathrm{length}_S(A) \cdot \mathrm{Tr}^{(S)}, \tag{6.15}$$

where the sum is taken over isomorphism classes of simple A-modules and $\mathrm{length}_S(A)$ is the number of times that S occurs in a composition series of A.

Recall from Note 1.4.6 that

$$A/\mathrm{J}(A) \cong \prod_{\substack{S \text{ simple} \\ \text{up to } \cong}} \mathrm{End}_{\mathrm{End}_A(S)}(S) \tag{6.16}$$

as $\mathbb{F}_p$-algebras. Since $I_{A/\mathbb{F}_p}$ is a two-sided ideal of A (by Note 6.2.83) and $I_{A/\mathbb{F}_p} \supseteq \mathrm{J}(A)$, we have that $I_{A/\mathbb{F}_p}/\mathrm{J}(A)$ is a two-sided ideal of $A/\mathrm{J}(A)$, and so it is a subproduct of (6.16).

Suppose $I_{A/\mathbb{F}_p} \neq \mathrm{J}(A)$. Then there is at least one simple S_0 occurring in this subproduct. Then $I_{A/\mathbb{F}_p}$ contains all elements of A that act as 0 on all simple A-modules not isomorphic to S_0. Consider such an element that acts as 1 on S_0 and as 0 on all other simples. Since it lies in the product of (6.16), it is represented by an element $r \in A$. Then by (6.15), the trace of r is $\mathrm{length}_{S_0}(A) \cdot \dim_{\mathbb{F}_p}(S_0)$. This quantity is strictly positive and less than or equal to $\dim_{\mathbb{F}_p}(A)$, so, for $p > \dim_{\mathbb{F}_p}(A)$, it is nonzero in $\mathbb{F}_p$. Hence $r \notin I_{A/\mathbb{F}_p}$, giving a contradiction.

□

Corollary 6.3.9. *There exists a deterministic polynomial-time algorithm that, given a finite algebra over a finite field $\mathbb{F}_p$, where p is a prime satisfying $p > \dim_{\mathbb{F}_p}(A)$, computes the Jacobson radical* $\mathrm{J}(A)$.

Proof. From Theorem 6.3.8 we have $I_{A/\mathbb{F}_p} = \mathrm{J}(A)$, and the trace ideal can be computed deterministically in polynomial time by Theorem 6.2.85. □

Note 6.3.10. We already knew this ([18, 27]).

Definition 6.3.11. *Let A be a finite ring and let $n := \mathrm{char}(A)$. We say a prime $p \mid n$ is* a small prime for A *if $p \leq \dim_{\mathbb{F}_p}(A/pA)$. We say a prime $p \mid n$ is* a large prime for A *if $p > \dim_{\mathbb{F}_p}(A/pA)$.*

Note 6.3.12. When it is clear what ring we are referring to, we will simply refer to a prime as being large or small.

Note 6.3.13. Let A be a finite ring. Then a prime $p \mid \mathrm{char}(A)$ is large if the number of cyclic direct summands of A^+ of size divisible by p (a quantity which is independent of the decomposition), is larger than p.

Proposition 6.3.14. *Let A be a finite ring and let m be its characteristic. Suppose m is divisible only by large primes and that A is projective as a $\mathbb{Z}/m\mathbb{Z}$-module. Let*

$$n' = \mathrm{rad}(m) := \prod_{\substack{p \mid m \\ p \text{ prime}}} p.$$

Then $n'A \subseteq \mathrm{J}(A)$ and

$$\begin{aligned} \mathrm{J}(A)/n'A &= \mathrm{J}(A/n'A) \\ &= I_{(A/n'A)/(\mathbb{Z}/n'\mathbb{Z})} \\ &= \left(I_{A/(\mathbb{Z}/m\mathbb{Z})} : \frac{m}{n'}A\right)/n'A, \end{aligned} \tag{6.17}$$

where $I_{A/(\mathbb{Z}/m\mathbb{Z})}$ is the trace radical of A over $\mathbb{Z}/m\mathbb{Z}$, as before, and

$$\left(I_{A/(\mathbb{Z}/m\mathbb{Z})} : \frac{m}{n'}A\right) := \{x \in A \mid \frac{m}{n'}x \in I_{A/(\mathbb{Z}/m\mathbb{Z})}\}.$$

Proof. To simplify notation, we write $I_A := I_{A/(\mathbb{Z}/m\mathbb{Z})}$ and $I_{A/n'A} := I_{(A/n'A)/(\mathbb{Z}/n'\mathbb{Z})}$. We will use similar abbreviations for the trace maps. Note that $A/n'A$ is projective as a module over $\mathbb{Z}/n'\mathbb{Z}$, so the trace map and the trace radical are well-defined.

It is easy to see that $n'A$ is nilpotent in A, since some power of n' is divisible by m. The first equality of (6.17) follows since $n'A$ is nilpotent and the second equality follows by Proposition 1.4.11, Theorem 6.3.8 and Lemma 6.2.84. For the last part,

note that by Proposition 6.2.72, part (iv), we have $\mathrm{Tr}_{A/n'A} \equiv \mathrm{Tr}_A \mod n'$. Hence

$$\begin{aligned} x + n'A \in I_{A/n'A} &\iff \mathrm{Tr}_{A/n'A}((A/n'A)(x + n'A)) = 0 \\ &\iff \mathrm{Tr}_{A/n'A}(Ax + n'A) = 0 \\ &\iff \mathrm{Tr}_A(Ax) \subset n'\mathbb{Z}/m\mathbb{Z} \\ &\iff \frac{m}{n'}\,\mathrm{Tr}_A(Ax) = 0 \\ &\iff \frac{m}{n'}x \in I_A \\ &\iff x + n'A \in \left(I_A : \frac{m}{n'}A\right)/n'A. \end{aligned}$$

□

Proposition 6.3.15. *Let A be a finite ring and let $\mathrm{char}(A) := p^e$, for some prime p and some $e \in \mathbb{Z}_{>0}$. Suppose that A is free as a $\mathbb{Z}/p^e\mathbb{Z}$-module and p is a large prime. Let $B := A/I_{A/(\mathbb{Z}/p^e\mathbb{Z})}$ and let $0 < e' \leq e$. Then*

(i) *if B is free as a $\mathbb{Z}/p^e\mathbb{Z}$-module, then B is separable over $\mathbb{Z}/p^e\mathbb{Z}$.*
(ii) *if $B/B[p^{e-e'}]$ is free as a $\mathbb{Z}/p^{e'}\mathbb{Z}$-module, then $B/B[p^{e-e'}]$ is separable over $\mathbb{Z}/p^{e'}\mathbb{Z}$, where $B[p^{e-e'}] = \ker(B \to B, b \mapsto p^{e-e'}b)$.*

Proof. (i) Since B is free over $\mathbb{Z}/p^e\mathbb{Z}$, by Theorem 6.2.18, part (iii), we have that B is separable over $\mathbb{Z}/p^e\mathbb{Z}$ if and only if B/pB is separable over $\mathbb{F}_p$. By Theorem 6.3.8, since $p > \dim_{\mathbb{F}_p}(A/pA)$, we have

$$B/pB = (A/pA)/I_{(A/pA)/\mathbb{F}_p} = (A/pA)/\,\mathrm{J}(A/pA),$$

so B/pB is semisimple, and thus B is separable over $\mathbb{Z}/p^e\mathbb{Z}$ (see Corollary 6.2.9).

(ii) Let $C := B/B[p^{e-e'}]$ and consider the canonical map $\pi : A \twoheadrightarrow C$. First note that

$$\begin{aligned} \ker(\pi) &:= \ker(A \twoheadrightarrow B := A/I_{A/(\mathbb{Z}/p^e\mathbb{Z})} \twoheadrightarrow C := B/B[p^{e-e'}]) \\ &= (I_{A/(\mathbb{Z}/p^e\mathbb{Z})} : p^{e-e'}A), \end{aligned}$$

where

$$(I_{A/(\mathbb{Z}/p^e\mathbb{Z})} : p^{e-e'}A) := \{x \in A \mid p^{e-e'}x \in I_{A/(\mathbb{Z}/p^e\mathbb{Z})}\}.$$

Further,

$$\begin{aligned} \pi^{-1}C[p^{e'-1}] &= \{x \in A \mid p^{e'-1}x \in (I_{A/(\mathbb{Z}/p^e\mathbb{Z})} : p^{e-e'}A)\} \\ &= (I_{A/(\mathbb{Z}/p^e\mathbb{Z})} : p^{e-1}A). \end{aligned}$$

Since C is free as a $\mathbb{Z}/p^{e'}\mathbb{Z}$-module, we have that $C[p^{e'-1}] = pC$, so

$$\pi^{-1}C[p^{e'-1}] = pA + \ker(\pi).$$

Hence

$$(I_{A/(\mathbb{Z}/p^e\mathbb{Z})} : p^{e-e'}A) + pA = (I_{A/(\mathbb{Z}/p^e\mathbb{Z})} : p^{e-1}A). \tag{6.18}$$

Since C is free as a $\mathbb{Z}/p^{e'}\mathbb{Z}$-module, by Theorem 6.2.18, part (iii), we have that C is separable over $\mathbb{Z}/p^{e'}\mathbb{Z}$ if and only if C/pC is semisimple. But, since p is large,

$$\begin{aligned} C/pC &= (A/pA)/((I_{A/(\mathbb{Z}/p^e\mathbb{Z})} : p^{e-e'}A) + pA)/pA) \\ &= (A/pA)/(I_{A/(\mathbb{Z}/p^e\mathbb{Z})} : p^{e-1}A)/pA), \quad \text{by (6.18)} \\ &= (A/pA)/I_{(A/pA)/\mathbb{F}_p}, \quad \text{by (6.17)} \\ &= (A/pA)/\operatorname{J}(A/pA), \quad \text{by Theorem 6.3.8.} \end{aligned}$$

□

Proposition 6.3.16. *Let A be a finite ring and let $n := \operatorname{char}(A)$. Suppose that A is projective as a $\mathbb{Z}/n\mathbb{Z}$-module. If all primes p dividing n are large, then $I_{A/(\mathbb{Z}/n\mathbb{Z})} \subseteq J(A)$.*

Proof. Let S be a simple A-module. Then S has exponent p for some prime $p \mid n$. So S is a simple module over A/pA. By Proposition 6.2.72, part (iv), the following diagram commutes:

$$\begin{array}{ccc} A & \longrightarrow & A/pA \\ \Big\downarrow{\scriptstyle \operatorname{Tr}_{A/(\mathbb{Z}/n\mathbb{Z})}} & & \Big\downarrow{\scriptstyle \operatorname{Tr}_{(A/pA)/\mathbb{F}_p}} \\ \mathbb{Z}/n\mathbb{Z} & \longrightarrow & \mathbb{F}_p. \end{array}$$

Since p is large, by Theorem 6.3.8 we have

$$(I_{A/(\mathbb{Z}/n\mathbb{Z})} + pA)/pA \subseteq \operatorname{J}(A/pA). \tag{6.19}$$

Let $i \in I_{A/(\mathbb{Z}/n\mathbb{Z})}$. By (6.19), the image of i in A/pA lies in $\operatorname{J}(A/pA)$, so it annihilates all simple A/pA-modules. In particular, it annihilates S. Hence i annihilates all simple A-modules. So $i \in \operatorname{J}(A)$. □

Proposition 6.3.17. *Let A be a finite ring and let $m := \operatorname{char}(A)$. Suppose A is projective over $\mathbb{Z}/m\mathbb{Z}$ and that all primes dividing m are large. Then $\operatorname{char}(A) = \operatorname{char}(A/I_{A/(\mathbb{Z}/m\mathbb{Z})})$.*

Proof. Write $I := I_{A/(\mathbb{Z}/m\mathbb{Z})}$. By Proposition 6.2.72, for all primes $p \mid m$, we have that $\operatorname{Tr}_{A/(\mathbb{Z}/m\mathbb{Z})}(1) \equiv \dim_{\mathbb{F}_p}(A/pA) \mod p$. Since all primes dividing m are large, we have $p > \dim_{\mathbb{F}_p}(A/pA)$, and so $p \nmid \operatorname{Tr}_{A/(\mathbb{Z}/m\mathbb{Z})}(1)$ for all $p \mid m$. Hence $\operatorname{Tr}_{A/(\mathbb{Z}/m\mathbb{Z})}(1)$ is a unit in A. Now by Proposition 6.3.16 we have $I \subseteq \operatorname{J}(A)$, so $\operatorname{Tr}_{A/(\mathbb{Z}/m\mathbb{Z})}(1) \notin I$, otherwise it would be a nilpotent element of A. Hence $\operatorname{Tr}_{A/(\mathbb{Z}/m\mathbb{Z})}(1)$ is an element of additive order m in A/I. □

6.3.3 Separating small primes from large primes

To develop algorithms from the theory in the previous sections, we first need to address the problem of separating small primes from large primes.

Proposition 6.3.18. *There exists a deterministic polynomial-time algorithm that, given a finite ring A of characteristic n, outputs two positive integers n_1 and n_2, such that:*

(i) $n_1 \cdot n_2 = n$,
(ii) $\gcd(n_1, n_2) = 1$,
(iii) *all primes dividing n_1 are small and all primes dividing n_2 are large.*

Moreover, the algorithm produces a prime factorisation of n_1.

Proof. Suppose the finite abelian group A^+ is given to the algorithm as the direct sum of cyclic groups $A^+ \cong \bigoplus_{i=1}^{t} \mathbb{Z}/d_i\mathbb{Z}$, where for all $1 \leq i \leq t$, we have $d_i \in \mathbb{Z}_{>1}$. Let $S := \{p \mid p \text{ prime}, p \leq t\}$. Note that if a prime p is small, then certainly $p \leq t$, since $\dim_{\mathbb{F}_p}(A/pA) = \#\{i \mid p \mid d_i\} \leq t$. So the set S will certainly contain all small primes. Now we decide which of the primes which occur in the factorisation of the d_i are actually small, by checking the condition $p \leq \dim_{\mathbb{F}_p}(A/pA)$. Finally, we gather all small primes, with their exponents, into n_1.

To see that the algorithm is polynomial-time, note that $t \leq \log_2(|A|)$. □

Note 6.3.19. In the process of running the algorithm of Proposition 6.3.18, we have obtained a complete factorisation of n_1, i.e. we know what all the small primes dividing n are, and what their multiplicities are. Note that any two isomorphic rings have the same small primes, with the same multiplicities.

Lemma 6.3.20. *There exists a deterministic polynomial-time algorithm that, given a finite ring A, computes a two-sided nilpotent ideal I of A and two positive integers n_1' and n_2' such that*

(i) $n_1' \cdot n_2' = \operatorname{char}(A/I)$,
(ii) $\gcd(n_1', n_2') = 1$,
(iii) *all primes dividing n_1' are small for A/I and all primes dividing n_2' are large for A/I,*
(iv) *n_1' is squarefree.*

Proof. We start by computing the characteristic of A using Theorem 3.3.1 and then apply Proposition 6.3.18 to compute n_1 and n_2, i.e. to separate small primes for A from large primes for A. This also gives a complete factorisation of n_1. Let $m \in \mathbb{Z}_{>1}$ be the largest integer such that $m^2 \mid n_1$ and write $n = m \cdot l$. Then $lA \neq 0$, but $(lA)^2 = 0$, so we have found a nilpotent two-sided ideal, $I = lA \subset A$. Put $n_1' = n_1/m$ and $n_2' = n_2$. Note also that the small (resp. large) primes for A are the same as the small (resp. large) primes for A/I. □

Theorem 6.3.21. *There exists a deterministic polynomial-time algorithm that, given a finite ring A, computes a two-sided nilpotent ideal $\mathfrak{i}$ such that*

$$A/\mathfrak{i} \cong A_1 \times A/n_2A,$$

where A_1 is semisimple and all primes dividing n_2 are large for $A/\mathfrak{i}$.

Proof. By Lemma 6.3.20, we may write

$$A/I \cong A/n_1A \times A/n_2A,$$

for some two-sided nilpotent ideal $I \subseteq A$, where n_2 is divisible only by large primes for A/I and n_1 is squarefree. For each prime $p \mid n_1$, the ring A/pA is an algebra over $\mathbb{F}_p$, so by Theorem 3.4.1 we may compute its Jacobson radical and factor it out of A, making the first component genuinely semisimple.

□

6.3.4 Algorithms

Theorem 6.3.22. *There exists a deterministic polynomial-time algorithm that, given a finite ring A, computes an approximation of the Jacobson radical of A.*

Proof. By Theorem 6.3.21, we may assume that $\mathrm{char}(A)$ is only divisible by large primes, i.e. for all primes $p \mid n$, we have $p > \dim_{\mathbb{F}_p}(A/pA)$. We proceed by building a sequence of rings that will terminate with a strongly separable ring (see Theorem 6.2.88 and Definition 6.2.89). We start by putting

$$A_0 := A.$$

To continue, we would like to make use of the trace and the trace radical and for this, we must ensure that we are working with a projective module. We thus quotient out by a multiple of A_0 to get

$$A_1 := A_0/mA_0,$$

where $m \mid \mathrm{char}(A)$ is such that mA_0 is a nilpotent two-sided ideal of A_0, and A_1 is projective as a module over $\mathbb{Z}/m\mathbb{Z}$. The way to do this deterministically in polynomial time has been described in Proposition 2.8.1.

We proceed by computing the trace radical $I := I_{A_1/(\mathbb{Z}/m\mathbb{Z})}$ using Theorem 6.2.85. By Proposition 6.3.8, we have $I \subseteq \mathrm{J}(A_1)$. The next term in our sequence is

$$A_2 := A_1/I.$$

If $I = 0$, then the map

$$\begin{aligned} \phi : A_1 &\to \mathrm{Hom}_{\mathbb{Z}/m\mathbb{Z}}(A_1, \mathbb{Z}/m\mathbb{Z}) \\ a &\mapsto (a \cdot \mathrm{Tr}_{A_1/(\mathbb{Z}/m\mathbb{Z})} : b \mapsto \mathrm{Tr}_{A_1/(\mathbb{Z}/m\mathbb{Z})}(ba)) \end{aligned}$$

is injective, and since A_1 is a finite projective $\mathbb{Z}/m\mathbb{Z}$-module, ϕ is an isomorphism of $\mathbb{Z}/m\mathbb{Z}$-modules, i.e. the trace map generates $\mathrm{Hom}_{\mathbb{Z}/m\mathbb{Z}}(A_1, \mathbb{Z}/m\mathbb{Z})$ as a left module, which is equivalent to A being strongly separable over $\mathbb{Z}/m\mathbb{Z}$ by Theorem 6.2.88.

If $I \neq 0$, then we treat A_2 as we did A_0, i.e. we make it projective as a module and calculate the trace radical of the resulting ring. We continue in this manner until the trace radical becomes equal to 0. In the end we will have produced a strongly separable algebra over its prime subring.

By Proposition 2.8.1, for every prime p dividing the characteristic of the final ring, a unique power of p occurs as an invariant of the underlying abelian group of the ring. Hence the generalised prime subring of the final ring is equal to its prime subring.

Set $\mathtt{j}_A$ to be the kernel of the map induced by the successive quotienting. Since $\mathrm{length}(A_i) < \mathrm{length}(A_{i-1})$, for all i, the algorithm terminates in polynomial time. □

Note 6.3.23. Suppose $\mathrm{char}(A) = n = \prod_{i=1}^{t} p_i^{e_i}$, for some $n \in \mathbb{Z}_{>0}$, p_i distinct primes and $e_i \in \mathbb{Z}_{>0}$. Then the number of iterations performed by the algorithm is bounded above by $\sum_{i=1}^{t} e_i - t$. This is because at each step we have to make the ring at hand projective over its prime subring in such a way that the radical of the characteristic is not changed.

Calculating the trace radical over and over again is a costly operation. There is a more economic way of proceeding, which we describe in the remaining part of this section. We begin with an auxiliary result.

Proposition 6.3.24. *Let A be a finite ring and let $\mathrm{char}(A) := m$, for some $m \in \mathbb{Z}_{>0}$. Suppose that A is projective as a $\mathbb{Z}/m\mathbb{Z}$-module and that all primes dividing m are large. Let $B := A/I_{A/(\mathbb{Z}/m\mathbb{Z})}$. Let $m' \mid m$ be such that $\mathrm{rad}(m) = \mathrm{rad}(m/m')$ and $B/B[m']$ is projective as a $\mathbb{Z}/\frac{m}{m'}\mathbb{Z}$-module. Then $B[m']$ is nilpotent and $B/B[m']$ is separable over $\mathbb{Z}/\frac{m}{m'}\mathbb{Z}$.*

Proof. Let p be a prime dividing m. Suppose its exponent in the prime factorisation of m is e, and its exponent in the prime factorisation of m' is $e - e' \geq 0$. Let $C := B/B[m']$ and $\varphi : A \to B := A/I_{A/(\mathbb{Z}/m\mathbb{Z})}$ be the canonical map. Then

$$\begin{aligned}\varphi^{-1}(B[p^{e-e'}]) &= \{x \in A \mid p^{e-e'}x \in I_{A/(\mathbb{Z}/p^e\mathbb{Z})}\} \\ &= \left(I_{A/(\mathbb{Z}/p^e\mathbb{Z})} : p^{e-e'}A\right).\end{aligned}$$

We have already seen in Propositions 6.3.15 and 6.3.14 that

$$(I_{A/(\mathbb{Z}/p^e\mathbb{Z})} : p^{e-e'}A) + pA = (I_{A/(\mathbb{Z}/p^e\mathbb{Z})} : p^{e-1}A). \tag{6.20}$$

and

$$(\varphi^{-1}(B[p^{e-e'}]) + pA)/pA = \mathrm{J}(A/pA),$$

so that $\varphi^{-1}(B[p^{e-e'}])$ is nilpotent, and hence $B[p^{e-e'}]$ is nilpotent.

Glueing along all primes (using Lemma 6.2.84), we get that

$$\left(\varphi^{-1}(B[m']) + n'A\right)/n'A = \mathrm{J}(A/n'A), \tag{6.21}$$

and $B[m']$ is nilpotent. Hence $C/n'C = (A/n'A)/\operatorname{J}(A/n'A)$ is semisimple. In particular, since n' is squarefree, $C/n'C$ is separable over $\mathbb{Z}/n'\mathbb{Z}$, by Theorem 6.2.44. Since $n' = \operatorname{rad}(m) = \operatorname{rad}(m/m')$, we have by Theorem 6.2.18, part (iii) and Proposition 6.2.16, part (ii), that C is separable over $\mathbb{Z}/\frac{m}{m'}\mathbb{Z}$.

□

Second proof of Theorem 6.3.22. By Theorem 6.3.21, we may assume that the characteristic of A is only divisible by large primes. We will construct a sequence of rings that terminates at a separable state. We begin as in the proof of Theorem 6.3.22:

$$A_0 := A, \quad A_1 := A_0/mA_0, \quad A_2 := A_1/I_{A_1/(\mathbb{Z}/m\mathbb{Z})},$$

where $m \mid \operatorname{char}(A_0)$ is such that mA_0 is a two-sided nilpotent ideal of A_0, and A_1 is projective over $\mathbb{Z}/m\mathbb{Z}$. We then proceed by putting

$$A_3 := A_2/A_2[m'],$$

where m' is computed using the deterministic polynomial-time algorithm described in Proposition 2.8.2. Then $m' \mid m$ is the least integer such that A_3 is projective over $\mathbb{Z}/\frac{m}{m'}\mathbb{Z}$. By Proposition 6.3.24 (taking $B := A_2$ and $C := A_3$), the ideal $A_2[m']$ is nilpotent and A_3 is separable over $\mathbb{Z}/\frac{m}{m'}\mathbb{Z}$.

By Propositions 2.8.1 and 2.8.2, for every prime p dividing the characteristic of the final ring, a unique power of p occurs as an invariant of the underlying abelian group of the ring. Hence the generalised prime subring of the final ring is equal to its prime subring.

Set $\mathfrak{j}_A = \ker(A \to A_3)$ under the map induced by the successive quotienting. □

Note 6.3.25. A natural question that arises is whether we have to compute any trace radical at all, i.e. whether given a finite ring A, the ideal given by $A[m']$ is not already an approximation of the Jacobson radical. The answer is no. To see this, consider the ring $A = \mathbb{Z}/p\mathbb{Z} \times \mathbb{Z}/p^2\mathbb{Z}$, where p is a prime. Then $\operatorname{char} A = p^2$ and $m' = p$. Now $\dim_{\mathbb{F}_p}(A/\operatorname{J}(A)) = 2$, but $A/A[m'] \cong \mathbb{F}_p$ has rank $1 < 2$. This contradicts Theorem 6.3.4.

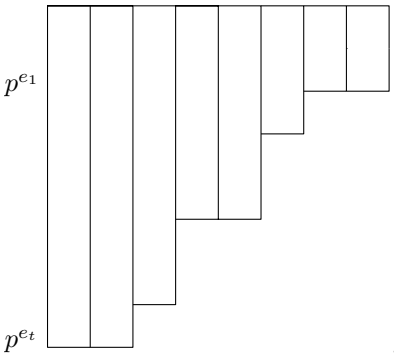
We can also produce a connected example of this type of failure. Let p be a prime and consider the ring $A = (\mathbb{Z}/p^2\mathbb{Z})[X]/(pX^2, X^3)$. Then $\operatorname{char}(A) = p^2$ and $m' = p$. Again, $\dim_{\mathbb{F}_p}(A/\operatorname{J}(A)) = 3$, but $\dim_{\mathbb{F}_p}(A/A[m']) = 2$.

6.3.5 An illustration

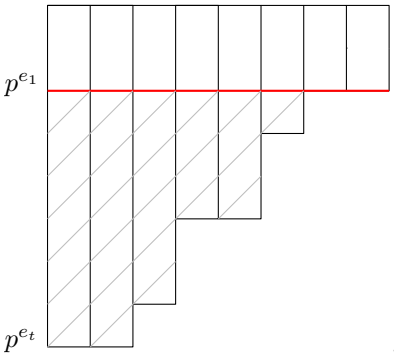
For more clarity, let us explore a graphical illustration of the above results. For simplicity, we restrict to the case that $m = p^e$ for some large prime p and some $e \in \mathbb{Z}_{>0}$. Suppose

$$A^+ \cong \bigoplus_{i=1}^{t} \mathbb{Z}/p^{e_i}\mathbb{Z},$$

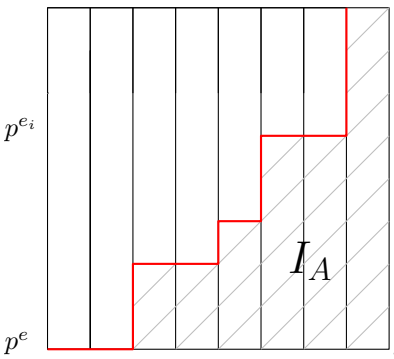
where $e_{i+1} \geq e_i$, $e_1 > 0$ and $e = e_t$. We will represent finite abelian p-groups by



where the number of vertical boxes is equal to the number of cyclic direct summands of A and the height of each such box is equal to the corresponding invariant. Now any finitely generated projective $\mathbb{Z}/p^e\mathbb{Z}$-module is free of finite rank, so is represented by a rectangle. To make A projective over its prime subring, we quotient out by $p^{e_1}A$, where p^{e_1} is the smallest nonzero invariant appearing in the decomposition of A^+ (or by $p^{e'}A$, for any $e' \leq e_1$).

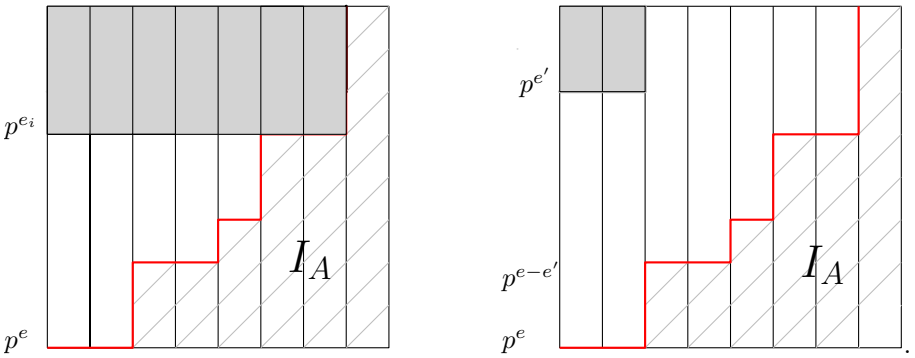


Suppose we have already made A projective as a module over $\mathbb{Z}/m\mathbb{Z}$, so that now $A^+ \cong (\mathbb{Z}/p^e\mathbb{Z})^t$. Then the next step in the algorithm is to quotient out by the trace radical to obtain $A_2 := A/I_A$ with quotient map $\varphi : A \to A_2$. We represent this graphically as:

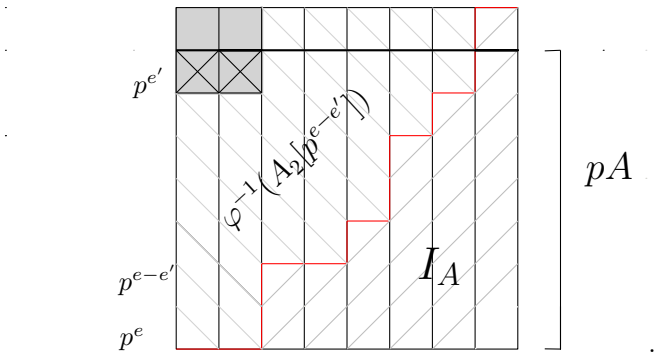


where the shaded lower-right part represents the trace radical. Note that the trace radical must touch the bottom line since when quotienting out by the trace radical, the characteristic remains unchanged (see Proposition 6.3.17). However, it may not touch the top line.

In the first algorithm, we proceed by making the remaining $A_2 := A/I_A$ projective over its prime subring, which we achieve by quotienting out by $p^{e_i}A_2$. This corresponds to looking at the upper-left grey rectangle in the first picture below, which we now treat as our initial box. In the second algorithm we proceed directly by quotienting out by $A_2[p^{e-e'}]$, where $p^{e-e'}$ is the second largest invariant in A_2 (or by quotienting out by $A_2[p^{e-f}]$, for any $0 < f \leq e'$). This leaves us looking at the upper-left grey rectangle in the second picture below, which is now separable projective over its prime subring by the second proof of Theorem 6.3.22.



Note 6.3.26. The equality in (6.20) can be easily seen from the following diagram, since both sides of the equality are represented by the upper-left grey area left unhatched:



6.3.6 Examples

In this section we look at some specific instances of trace computation and running of the algorithms given as proofs of Theorem 6.3.22. We will only consider examples where the characteristic of the given ring is divisible exclusively by large primes.

Note 6.3.27. If A is a finite ring of prime characteristic, then by Theorem 6.2.44, both our algorithms will output the Jacobson radical of A.

Note that rings that are finite products rings of the form $\mathbb{Z}/n\mathbb{Z}$, with $n \in \mathbb{Z}_{>0}$ and componentwise addition and multiplication, are separable over their prime subrings (see Proposition 6.2.16, part (iii)). Thus, if they are projective over their prime subring, they are strongly separable over their prime subring (see Proposition 6.2.86), so their trace ideal is trivial.

Example 6.3.28 (Integers modulo n)**.** Let $A = (\mathbb{Z}/5\mathbb{Z})^2 \times (\mathbb{Z}/3^2\mathbb{Z})$. Note that $5 > 2$ and $3 > 1$, so the primes occurring are large. The prime subring of A is $k = \mathbb{Z}/45\mathbb{Z}$ and A is projective as a k-module and hence A is strongly separable over k. Thus, if A is given to the algorithms proving Theorem 6.3.22, they will find that $I_{A/(\mathbb{Z}/45\mathbb{Z})} = 0$, and will therefore output $\mathfrak{j}_A = 0$.

It is easy to check that a projective basis of A over k is given by $\{F_i, x_i\}_{i=1}^3$, where

$$F_1 : (a, b, c) \mapsto 9a, \quad F_2 : (a, b, c) \mapsto 9b, \quad F_3 : (a, b, c) \mapsto 5c,$$

and

$$x_1 = (4, 0, 0), \quad x_2 = (0, 4, 0), \quad x_3 = (0, 0, 2).$$

Hence by Definition 6.2.75,

$$\mathrm{Tr}_{A/(\mathbb{Z}/45\mathbb{Z})}(a, b, c) = 9 \cdot 4a + 9 \cdot 4b + 5 \cdot 2c,$$

and so $\mathrm{Tr}_{A/(\mathbb{Z}/45\mathbb{Z})}(1) = 82$, which is indeed congruent to $(2 \mod 5)$ and to $(1 \mod 9)$, as predicted by Proposition 6.2.72. Moreover, $82 \equiv 37 \mod 45$, so that $\mathrm{Tr}_{A/(\mathbb{Z}/45\mathbb{Z})}(1)$ is a unit in $\mathbb{Z}/45\mathbb{Z}$ (cf. Theorem 6.2.88, part (iii)).

Example 6.3.29. Consider again the ring of Example 6.3.5,

$$A = \mathbb{Z}/p^2\mathbb{Z} \oplus M,$$

where $p > 2$ is a prime, M is a 1-dimensional $\mathbb{F}_p$-vector space, addition is componentwise and multiplication is given by

$$(a, x) \cdot (b, y) = (ab, ay + bx).$$

Then A is a commutative local ring with maximal ideal $p\mathbb{Z}/p^2\mathbb{Z} \oplus M$. The first step of our algorithms gives $A_1 := \mathbb{Z}/p\mathbb{Z} \oplus M$. Then $I_{A_1/\mathbb{F}_p} = M$ and $A_2 := \mathbb{Z}/p\mathbb{Z}$. Hence both algorithms output $\mathfrak{j}_A = pA + M$.

Example 6.3.30 (Group rings)**.** Let $k = \mathbb{Z}/n\mathbb{Z}$, for some $n \in \mathbb{Z}_{>0}$, and G a finite group. Let $A = k[G]$ be the group ring of G over k. We know from Example 6.2.92, part (3), that if $|G| \cdot 1$ is a unit in k, then A is strongly separable over k. Suppose all primes dividing n are large. Then $p > \dim_{\mathbb{F}_p}(A/pA) = \dim_{\mathbb{F}_p}(\mathbb{F}_p[G]) = |G|$, so $|G|$ is a unit in k and our algorithms output $\mathfrak{j}_A = 0$.

Let $k = \mathbb{Z}/n\mathbb{Z}$ and $A = \mathcal{M}_m(k)$, for some $m \in \mathbb{Z}_{>0}$. We know from Example 6.2.92, parts 1 and 2, that $\mathcal{M}_m(k)$ is strongly separable if and only if m is a unit in k. Moreover, $\mathrm{Tr}_{A/k} = m \cdot (\text{usual trace})$.

Example 6.3.31 (Matrix rings)**.** The smallest example of a matrix ring over a commutative ring whose characteristic is divisible only by large primes is $A = \mathcal{M}_2(\mathbb{F}_5)$. In this case, A is simple, but since the primes occurring are large, the algorithms will not be able to detect this. Since 2 is a unit in $\mathbb{F}_5$, the ring A is strongly separable, so the algorithms will output $\mathfrak{j}_A = 0$.

6.3.7 Remarks

Functoriality

Proposition 6.3.32. *Let $\mathcal{F}$ be the class of finite rings. The two families of ideals $(\mathfrak{j}_A)_{A\in\mathcal{F}}$ and $(\mathfrak{j}'_A)_{A\in\mathcal{F}}$, produced by the two algorithms described in the two proofs of Theorem 6.1.2 are functorial under isomorphisms, i.e. if $\phi : A \to B$ is an isomorphism of finite rings, then $\phi(\mathfrak{j}_A) = \mathfrak{j}_B$ and $\phi(\mathfrak{j}'_A) = \mathfrak{j}'_B$.*

Proof. It is clear by construction (Propositions 2.8.1 and 2.8.2) that two isomorphic rings will yield the same m and m'. Trace ideals are compatible with ring isomorphisms by Proposition 6.2.72, part (iv). □

Comparison between proofs of Theorem 6.3.22

We have already noted that the algorithm given in the second proof of Theorem 6.3.22 performs only 3 steps. But what can we say about the number of iterations (trace radical computations) needed in the first algorithm?

If A is a finite ring, let us write $\mathfrak{j}_1^A$ for the approximation of the Jacobson radical of A produced by the first proof of Theorem 6.3.22, and $\mathfrak{j}_2^A$ for the approximation of the Jacobson radical of A produced by the second proof. So far we have only seen examples where $\mathfrak{j}_1^A = \mathfrak{j}_2^A$. A natural question to ask is how $\mathfrak{j}_1$ and $\mathfrak{j}_2$ compare (with respect to inclusion), or indeed whether they are comparable at all.

We give partial answers to these questions in this section.

Let $e \in \mathbb{Z}_{>0}$ and let $p > 2$ be a prime. Let $e' \in \mathbb{Z}_{>0}$ be such that $2e' < e$. Set

$$A = (\mathbb{Z}/p^e\mathbb{Z})[X]/(X^2 - p^{e'}X).$$

Since A is already projective over its prime subring, in the notation of our algorithms, we have $A_0 := A = A_1$. Write $k := \mathbb{Z}/p^e\mathbb{Z}$, $\mathrm{Tr}_A := \mathrm{Tr}_{A/k}$ and $I_A := I_{A/k}$. Then

$$\mathrm{Tr}_A(1) = 2,$$
$$\mathrm{Tr}_A(X) = p^{e'}.$$

The matrix representing the map

$$A \to \mathrm{Hom}(A, \mathbb{Z}/p^e\mathbb{Z})$$

is then given by

$$F = \begin{pmatrix} 2 & p^{e'} \\ p^{e'} & p^{2e'} \end{pmatrix}.$$

Note that by Theorem 6.3.4, the rank of $A/\mathfrak{j}_1^A$ and that of $A/\mathfrak{j}_2^A$ must be equal to 1.

Suppose that $e = 2e'N + r$, for some $N \in \mathbb{Z}_{>0}$ and some $1 \leq r \leq 2e'$. Then

$$\mathfrak{j}_1^A = \left(p^r, X - \frac{p^{e'}}{2}\right) \tag{6.22}$$

and

$$\mathfrak{j}_2^A = \left(p^{2e'}, X - \frac{p^{e'}}{2}\right). \tag{6.23}$$

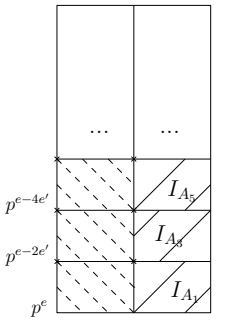
This example illustrates three remarks:

Remark I. The number of iterations performed by the first algorithm is unbounded.

To see this, note that

$$I_A = k \cdot p^{e-2e'}\left(X - \frac{p^{e'}}{2}\right)$$

and the rank of A/I_A is still 2. For the next trace radical computation, we simply replace e by $e-2e'$. Hence if $e > 2e'N$, then the first algorithm performs at least N trace radical computations. Graphically, the algorithm computing $\mathfrak{j}_1^A$ can be represented by the following diagram:



Algorithm 1

Remark II. The cardinality of $A/\mathfrak{j}_2^A$ may be larger than the cardinality of $A/\mathfrak{j}_1^A$.
Compare (6.22) with (6.23).
Remark III. It is possible that $\mathfrak{j}_2^A \subseteq \mathfrak{j}_1^A$.
Compare (6.22) with (6.23).

Good properties

We summarize the good properties of the ring $A/\mathfrak{j}_A$, where $\mathfrak{j}_A$ is an approximation of the Jacobson radical of A.

Theorem 6.3.33. *Let A be a finite ring and $\mathfrak{j}_A$ an approximation of the Jacobson radical of A. Then*

(i) *$A/\mathfrak{j}_A$ is separable,*
(ii) *the prime subring and the generalised prime subring of $A/\mathfrak{j}_A$ coincide,*
(iii) *$A/\mathfrak{j}_A$ admits projectivity and injectivity lift from its prime subring,*
(iv) *$A/\mathfrak{j}_A$ is a quasi-Frobenius ring,*
(v) *$A/\mathfrak{j}_A$ is a symmetric algebra over its prime subring.*

Applications and further questions

The exploration of possible applications of the computation of an approximation of the Jacobson radical of a finite ring is a problem for future research. We record here an application to testing simplicity of a finite module M over a finite ring R. More basic algorithms for this were given in Chapter 5, Section 5.3.

Third proof of Theorem 5.3.1. We begin by computing an approximation $\mathfrak{j}_R$ of the Jacobson radical of R using either of the proofs of Theorem 6.3.22. By Proposition 1.3.14, it is enough to test whether M is simple as an $R/\mathfrak{j}_R$-module, which by Theorem 1.6.4 reduces to testing whether $\mathrm{End}_{R/\mathfrak{j}_R}(M)$ is a field. This can be done by Theorem 5.1.2. □

It is also an interesting question to decide if for a finite ring A, we have $\mathfrak{j}_2^A \subseteq \mathfrak{j}_1^A$ in general. This has not been contradicted by the examples we have considered.

6.4 Computing the generalised prime subring

Let A be a finite ring and denote its generalised prime subring by $\mathcal{P}_A$ (see Definition 6.2.48). The two algorithms proving Theorem 6.3.22 each produce approximations $\mathfrak{j}_A$ of the Jacobson radical of A. In particular, the prime subring of the ring $A/\mathfrak{j}_A$ is equal to its generalised prime subring (see Definition 6.3.1). In what follows, we give a deterministic polynomial-time algorithm that, given a finite ring A, computes $\mathcal{P}_A$.

By Lemma 6.2.50, the generalised prime subring of a finite ring is equal to the generalised prime subring of its centre. Hence we may restrict to the case that A is a

finite commutative ring. Then

$$A = \prod_{\mathfrak{m} \text{ maximal}} A_{\mathfrak{m}},$$

where the product is taken over maximal ideals of A and $A_{\mathfrak{m}}$ denotes the localisation of A at $\mathfrak{m}$. Let $e_{\mathfrak{m}}$ be the primitive idempotent corresponding to $A_{\mathfrak{m}}$. Then $e_{\mathfrak{m}}$ has order a prime power, equal to $\exp(A_{\mathfrak{m}}^+)$.

Let $Q = \{\exp(A_{\mathfrak{m}}^+) \mid \mathfrak{m} \subseteq A \text{ maximal ideal}\}$. Define a map $\mathbb{Z}_{>0} \to A$, $q \mapsto e_q$, where

$$e_q = \sum_{\exp(A_{\mathfrak{m}}^+)=q} e_{\mathfrak{m}}. \tag{6.24}$$

Note that if $q \in \mathbb{Z}_{>0} \backslash Q$, then $e_q = 0$. Moreover,

$$\sum_{q \in Q} e_q = 1.$$

Let

$$B := \sum_{q \in Q} \mathbb{Z} e_q \cong \prod_{q \in Q} \mathbb{Z}/q\mathbb{Z}. \tag{6.25}$$

Then B is a subring of A, since the e_q are orthogonal idempotents of sum 1 (see Theorem 1.5.4). It is easy to see from Definition 6.2.48 that $B = \mathcal{P}_A$.

Proposition 6.4.1. *Let A be a finite commutative ring such that*

$$A^+ \cong \bigoplus_{d \in D} (\mathbb{Z}/d\mathbb{Z})^{n_d}.$$

where $n_d \in \mathbb{Z}_{>0}$ for all $d \in D$. For $d \in D$, let $A[d] = \ker(A \to A, a \mapsto da)$ and

$$I_d := \bigcap_n A[d]^n.$$

Then for all $d \in D$, there exists a unique element $f_d \in I_d$ such that for all $x \in I_d$, we have $f_d x = x$. Moreover,

$$\mathcal{P}_A = \sum_{d \in D} \mathbb{Z} f_d.$$

Proof. Let $d \in D$. First note that $A[d]$ is an ideal of A, and hence so is I_d. Note that

$$A[d] = \prod_{\mathfrak{m}} A_{\mathfrak{m}}[d] = \left(\prod_{\substack{\mathfrak{m} \\ \exp(A_{\mathfrak{m}}^+) \mid d}} A_{\mathfrak{m}} \right) \times \left(\prod_{\substack{\mathfrak{m} \\ \exp(A_{\mathfrak{m}}^+) \nmid d}} \mathfrak{a}_{d,\mathfrak{m}} \right),$$

for some ideal $\mathfrak{a}_{d,\mathfrak{m}} \subsetneq A_{\mathfrak{m}}$ contained in the maximal ideal of $A_{\mathfrak{m}}$. Then

$$I_d = \left(\prod_{\substack{\mathfrak{m} \\ \exp(A_{\mathfrak{m}}^+) | d}} A_{\mathfrak{m}} \right) \times 0,$$

and the identity in the first factor is $\sum_{\substack{q \in Q \\ q|d}} e_q$. Let

$$f_d := \sum_{\substack{q \in Q \\ q|d}} e_q. \tag{6.26}$$

For uniqueness, suppose $f'_d \in I_d$ is another element such that for all $x \in I_d$, we have $f'_d x = x$. Then $f_d = f'_d f_d = f_d$.

We claim that if $e_q \neq 0$, then there exists $d \in D$ such that d is exactly divisible by q, i.e. $q \mid d$ and $\gcd(q, d/q) = 1$. This is because if $e_q \neq 0$, then there exists $\mathfrak{m}$ such that $\exp(A_{\mathfrak{m}}^+) = q$ and $\mathbb{Z}/q\mathbb{Z}$ is a direct summand of $A_{\mathfrak{m}}^+$, and hence of A^+.

Further, we claim that for every $d \in D$ and every prime $p \mid d$, we have that

$$\sum_{\substack{i \\ p^i | d}} e_{p^i} \in \sum_{d \in D} \mathbb{Z} f_d.$$

To see this, let $d \in D$ and suppose $d = d_1 d_2$, where $d_1, d_2 \in \mathbb{Z}_{>0}$ and $(d_1, d_2) = 1$. Then $f_d = f_{d_1} + f_{d_2}$, so $\mathbb{Z} f_d \subseteq \mathbb{Z} f_{d_1} + \mathbb{Z} f_{d_2}$. Moreover, the additive orders of f_{d_1} and f_{d_2} are coprime, and so the additive order of $f_{d_1} + f_{d_2}$ is equal to the additive order of f_d. Hence $\mathbb{Z} f_d \subseteq \mathbb{Z} f_{d_1} + \mathbb{Z} f_{d_2}$. Suppose that for some prime p and some $r \in \mathbb{Z}_{>0}$, we have that p^r divides d exactly. Take $d_1 = p^r$. Then $f_{d_1} = \sum_{0 \leq i \leq r} e_{p^i} \in \mathbb{Z} f_d$, as required.

We now show that $\sum_{d \in D} \mathbb{Z} f_d = B$, where B is as in (6.25). That $\sum_{d \in D} \mathbb{Z} f_d \subseteq B$ follows from (6.26). For the other inclusion, we will show that for all $q = p^r$, with p a prime and $r \in \mathbb{Z}_{>0}$, we have $e_q \in \sum_{d \in D} \mathbb{Z} f_d$. Fix a prime p. If $e_q = 0$, then we are done. Otherwise, pick $d \in D$ that is exactly divisible by q. Then $e_{p^r} + \sum_{1 \leq i < r} e_{p^i} \in \sum_{d \in D} \mathbb{Z} f_d$. By induction on r, we have $e_q \in \sum_{d \in D} \mathbb{Z} f_d$.

□

Note 6.4.2. Proposition 6.4.1 is true for any choice of D. However, if we choose D such that $d_1 \mid d_2 \mid \ldots \mid d_t$, where $t = |D|$, then the relations between the f_{d_i} become simpler. First note that f_{d_i} is an idempotent for every $1 \leq i \leq t$. To make them orthogonal, put $f'_{d_i} := f_{d_i} - f_{d_{i-1}}$ for every $1 < i \leq t$. Let s_i be the order of f'_{d_i}. Then

$$\mathcal{P}_A = \sum_{i=1}^{t} \mathbb{Z} f'_{d_i} \cong \prod_{i=1}^{t} \mathbb{Z}/s_i\mathbb{Z}$$

as rings.

Theorem 6.4.3. *There exists a deterministic polynomial-time algorithm that, given a finite ring A, computes the generalised prime subring $\mathcal{P}_A$.*

Proof. Since $\mathcal{P}_{\mathrm{Z}(A)} = \mathcal{P}_A$, we may assume $A = \mathrm{Z}(A)$. Suppose A^+ is given to the algorithm as

$$A^+ \cong \bigoplus_{d \in D} (\mathbb{Z}/d\mathbb{Z})^{n_d},$$

where $n_d \in \mathbb{Z}_{>0}$. By computing the Smith normal form of the corresponding group presentation matrix, we can ensure that for $D = \{d_1, \ldots, d_t\}$, we have $d_1 \mid d_2 \mid \ldots \mid d_t$ and $d_1 \neq \pm 1$. For each $1 \leq i \leq t$, compute f_{d_i}, as in Proposition 6.4.1. Turn the set $\{f_{d_i}\}_{1 \leq i \leq t}$ into a set of orthogonal idempotents as in Note 6.4.2. Let s_i be the order of f'_{d_i}. The output of the algorithm consists of the set $\{s_i\}_{1 \leq i \leq t}$, together with a map $\prod_{1 \leq i \leq t} \mathbb{Z}/s_i\mathbb{Z} \to A$, given by $1_{(\mathbb{Z}/s_i\mathbb{Z})} \mapsto f_{d_i}$.

□

Note 6.4.4. If the elements of D are not assumed to divide each other, then the additive relations between the f_d can be computed by solving systems of linear equations over $\mathbb{Z}$.

The computation of the generalised prime subring gives another way of testing whether a finite ring is separable.

Second proof of Theorem 6.2.19. By Theorem 6.2.52, a finite ring is separable if and only if it is separable projective over its generalised prime subring. So compute $\mathcal{P}_A$ using Theorem 6.4.3 and then test projectivity of A over $\mathcal{P}_A$ using Theorem 5.4.1. □

References

[1] M. Agrawal, N. Kayal, and N. Saxena. Primes is in P. *Ann. of Math*, 2:781–793, 2002.

[2] M. Agrawal and N. Saxena. Automorphisms of finite rings and applications to complexity of problems. In Volker Diekert and Bruno Durand, editors, *STACS 2005*, volume 3404 of *Lecture Notes in Computer Science*, pages 1–17. Springer Berlin Heidelberg, 2005.

[3] M. Aguiar. A note on strongly separable algebras. *Boletín de la Academia Nacional de Ciencias (Córdoba, Argentina)*, special issue in honor of Orlando Villamayor(65):51–60, 2000.

[4] V. Arvind, B. Das, and P. Mukhopadhyay. The complexity of black-box ring problems. In *Computing and Combinatorics*, volume 4112 of *Lecture Notes in Computer Science*, pages 126–135. Springer Berlin Heidelberg, 2006.

[5] M. Auslander and O. Goldman. The Brauer group of a commutative ring. *Trans. Am. Math. Soc.*, 97:367–409, 1961.

[6] L. Babai. Graph isomorphism in quasipolynomial time. *preprint*, arXiv:1512.03547, 2015.

[7] H. Bass. *Traces and Euler characteristics*. Lecture Note Series. Cambridge University Press, 1979.

[8] D.J. Bernstein. Factoring into coprimes in essentially linear time. *J. Algorithms*, 54(1):1–30, 2005.

[9] G. Bini and F. Flamini. *Finite Commutative Rings and Their Applications*. The Springer International Series in Engineering and Computer Science. Springer US, 2012.

[10] P.A. Brooksbank and E.M. Luks. Testing isomorphism of modules. *Journal of Algebra*, 320(11):4020–4029, 2008.

[11] P.A. Brooksbank and J.B. Wilson. The module isomorphism problem reconsidered. *Journal of Algebra*, 421:541–559, 2015. Special issue in memory of Ákos Seress.

[12] J.A. Buchmann and H.W. Lenstra. Approximating rings of integers in number fields. *Journal de théorie des nombres de Bordeaux*, 6(2):221–260, 1994.

[13] J.P. Buhler and P. Stevenhagen. *Algorithmic Number Theory: Lattices, Number Fields, Curves and Cryptography.* Mathematical Sciences Research Institute Publications. Cambridge University Press, 2008.

[14] J.F. Buss, G.S. Frandsen, and J.O. Shallit. The computational complexity of some problems of linear algebra. *Journal of Computer and System Sciences*, 58(3):572–596, 1999.

[15] A. Chistov, G. Ivanyos, and M. Karpinski. Polynomial time algorithms for modules over finite dimensional algebras. In *Proceedings of the 1997 International Symposium on Symbolic and Algebraic Computation*, ISSAC '97, pages 68–74, New York, USA, 1997. ACM.

[16] T.-W.J. Chou and G.E. Collins. Algorithms for the solution of systems of linear diophantine equations. *SIAM Journal on Computing*, 11(4):687–708, 1982.

[17] I. Ciocănea-Teodorescu. The module isomorphism for finite rings and related results. *preprint*, arXiv:1512.08365v1, 2015.

[18] A.M. Cohen, G. Ivanyos, and D.B. Wales. Finding the radical of an algebra of linear transformations. *Journal of Pure and Applied Algebra*, 117-118:177–193, 1997.

[19] H. Cohen. *A Course in Computational Algebraic Number Theory.* Graduate Texts in Mathematics. Springer, 1993.

[20] F. DeMeyer and E. Ingraham. *Separable algebras over commutative rings.* Lecture Notes in Mathematics. Springer, Berlin, 1971.

[21] F.R. DeMeyer. The trace map and separable algebras. *Osaka Journal of Mathematics*, 3(1):7–11, 1966.

[22] L.E. Dickson. Algebras and their arithmetics. *Bulletin of the American Mathematical Society*, 30(5-6):247–257, 1924.

[23] R. Eggermont. *Modellen voor eindige lichamen.* Bachelor thesis. Mathematical Institute, Leiden University, 2009.

[24] S. Endo and Y. Watanabe. On separable algebras over a commutative ring. *Osaka Journal of Mathematics*, 4(2):233–242, 1967.

[25] S. Endo and Y. Watanabe. The centers of semi-simple algebras over a commutative ring. ii. *Nagoya Mathematical Journal*, 39:1–6, 1970.

[26] B. Farb and R.K. Dennis. *Noncommutative Algebra.* Graduate Texts in Mathematics. Springer New York, 2012.

[27] K. Friedl and L. Rónyai. Polynomial time solutions of some problems of computational algebra. In *Proceedings of the Seventeenth Annual ACM Symposium on Theory of Computing*, STOC '85, pages 153–162, New York, NY, USA, 1985. ACM.

[28] T. Fritzsche. The Brauer group of character rings. *Journal of Algebra*, 361(0):37–40, 2012.

[29] G. Ganske and B.R. McDonald. Finite local rings. *Rocky Mountain J. Math.*, 3(4):521–540, 1973.

[30] J.L. Hafner and K.S. McCurley. Asymptotically fast triangulation of matrices over rings. In *Proceedings of the First Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '90, pages 194–200, Philadelphia, PA, USA, 1990. Society for Industrial and Applied Mathematics.

[31] A. Hattori. Semisimple algebras over a commutative ring. *Journal of the Mathematical Society of Japan*, 15(4):404–419, 1963.

[32] A. Hattori. On strongly separable algebras. *Osaka Journal of Mathematics*, 2(2):369–372, 1965.

[33] G. Havas, B.S. Majewski, and K.R. Matthews. Extended GCD and Hermite normal form algorithms via lattice basis reduction. *Experiment. Math.*, 7(2):125–136, 1998.

[34] G. Higman. On a conjecture of Nagata. *Proceedings of the Cambridge Philosophical Society*, 52(Part I), January 1956.

[35] M. Hitz, J. Grabmeier, E. Kaltofen, and V. Weispfenning. *Computer Algebra Handbook: Foundations · Applications · Systems*. SpringerLink : Bücher. Springer Berlin Heidelberg, 2012.

[36] D. F. Holt and S. Rees. Testing modules for irreducibility. *Journal of the Australian Mathematical Society (Series A)*, 57:1–16, 8 1994.

[37] D.F. Holt. The meataxe as a tool in computational group theory. In R.T. Curtis and R.A. Wilson, editors, *The Atlas of Finite Groups - Ten Years on*, pages 74–81. Cambridge University Press, 1998. Cambridge Books Online.

[38] D.F. Holt, B. Eick, and E.A. O'Brien. *Handbook of Computational Group Theory*. Discrete Mathematics and Its Applications. CRC Press, 2005.

[39] G. Ivanyos. Modules and maximum rank matrix completion. Presented at the Combinatorics, Groups, Algorithms, and Complexity Conference, Columbus, Ohio, March 21-25, 2010.

[40] G. Ivanyos, M. Karpinski, Y. Qiao, and M. Santha. Generalized Wong sequences and their applications to Edmonds' problems. *Journal of Computer and System Sciences*, 81(7):1373–1386, 2015.

[41] G. Ivanyos, M. Karpinski, L. Rónyai, and N. Saxena. Trading GRH for algebra: Algorithms for factoring polynomials and related structures. *Math. Comput.*, 81(277):493–531, 2012.

[42] G. Ivanyos, M. Karpinski, and N. Saxena. Deterministic polynomial time algorithms for matrix completion problems. *SIAM J. Comput.*, 39(8):3736–3751, 2010.

[43] G. Ivanyos and K.M. Lux. Treating the exceptional cases of the meataxe. *Experimental Mathematics*, 9(3):373–381, 2000.

[44] G. Ivanyos and L. Rónyai. *Computations in Associative and Lie Algebras*, volume 4 of *Algorithms and Computation in Mathematics*. Springer Berlin Heidelberg, 1999.

[45] G. Ivanyos, L. Rónyai, and J. Schicho. Splitting full matrix algebras over algebraic number fields. *Journal of Algebra*, 354(1):211–223, 2012.

[46] G. Ivanyos, L. Rónyai, and J. Schicho. Improved algorithms for splitting full matrix algebras. *JP Journal of Algebra, Number Theory and Applications*, 28(2):141–156, 2013.

[47] N. Jacobson. *Lie Algebras*. Dover Books on Mathematics Series. Dover, 1979.

[48] N. Jacobson. *Basic algebra II*. Basic Algebra. Dover Publications, Incorporated, 2009.

[49] L. Kadison and A.A. Stolin. *Separability and Hopf algebras*. Algebra and Its Applications: International Conference [on] Algebra and Its Applications, March 25-28, 1999, Ohio University, Athens. American Mathematical Society, 2000.

[50] T. Kanzaki. Special type of separable algebra over a commutative ring. *Proc. Japan Acad.*, 40(10):781–786, 1964.

[51] I. Kaplansky. Rings with a polynomial identity. *Bulletin of the American Mathematical Society*, 54(6):575–580, 1948.

[52] N. Kayal and N. Saxena. On the ring isomorphism and automorphism problems. *IEEE Conference on Computational Complexity*, pages 2–12.

[53] N. Kayal and N. Saxena. Complexity of ring morphism problems. *Computational Complexity*, 15(4):342–390, June 2006.

[54] M.-A. Knus and M. Ojanguren. *Théorie de la descente et algèbres d'Azumaya*. Lecture Notes in Mathematics, Vol. 389. Springer-Verlag, Berlin, 1974.

[55] N. Koblitz. *A Course in Number Theory and Cryptography*. Springer-Verlag New York, Inc., New York, NY, USA, 1987.

[56] T.Y. Lam. *Lectures on Modules and Rings*. Graduate Texts in Mathematics. Springer New York, 1999.

[57] T.Y. Lam. *A First Course in Noncommutative Rings*. Graduate Texts in Mathematics. Springer, 2001.

[58] T.Y. Lam. *Serre's Problem on Projective Modules*. Springer Monographs in Mathematics. Springer Berlin Heidelberg, 2010.

[59] S. Landau. Some remarks on computing the square parts of integers. *Information and Computation*, 78(3):246 – 253, 1988.

[60] S. Lang. *Algebra*. Graduate Texts in Mathematics. Springer New York, 2002.

[61] A.K. Lenstra. Factorization of polynomials. In *Computational methods in number theory*, Mathematical Centre Tracts 154-155, pages 169–198, Amsterdam, 1984. Mathematisch Centrum.

[62] A.K. Lenstra. Integer factoring. *Designs, Codes and Cryptography*, 19(2):101–128, 2000.

[63] A.K. Lenstra, H.W. Lenstra, and L. Lovász. Factoring polynomials with rational coefficients. *Math. Ann.*, 261(4):515–534, 1982.

[64] H.W. Lenstra. *Galois Theory for Schemes*. Course notes available from the server of the Universiteit Leiden Mathematics Department, http://websites.math.leidenuniv.nl/algebra/GSchemes.pdf. Electronic third edition: 2008.

[65] H.W. Lenstra. Finding isomorphisms between finite fields. *Mathematics of Computation*, 56(193):329–347, 1991.

[66] H.W. Lenstra. Algorithms in algebraic number theory. *Bull. Amer. Math. Soc.*, 26:211–244, 1992.

[67] H.W. Lenstra. Flags and lattice basis reduction. In *In Proceedings of the third European congress of mathematics*. Birkhuser, 2001.

[68] H.W. Lenstra. Lattices. In *Algorithmic number theory: lattices, number fields, curves and cryptography*, volume 44 of *Math. Sci. Res. Inst. Publ.*, pages 127–181. Cambridge Univ. Press, Cambridge, 2008.

[69] K.M. Lux and M. Szöke. Computing homomorphism spaces between modules over finite dimensional algebras. *Experiment. Math.*, 12(1):91–98, 2003.

[70] K.M. Lux and M. Szöke. Computing decompositions of modules over finite-dimensional algebras. *Experiment. Math.*, 16(1):1–6, 2007.

[71] G. Marks and M. Schmidmeier. Extensions of simple modules and the converse of Schur's lemma. In *Advances in Ring Theory*, Trends in Mathematics, pages 229–237. Birkhäuser Basel, 2010.

[72] M. Orzech and C. Small. *The Brauer group of commutative rings*. Number v. 11 in Lecture notes in pure and applied mathematics. M. Dekker, 1975.

[73] C.H. Papadimitriou. *Computational Complexity*. Theoretical computer science. Addison-Wesley, 1994.

[74] R. Parker. The computer calculation of modular characters (the meat-axe). *Computational Group Theory*, pages 267–274, 1984.

[75] L. Rónyai. Computing the structure of finite algebras. *J. Symb. Comput.*, 9(3):355–373, March 1990.

[76] J. Rotman. *An Introduction to Homological Algebra*. Universitext. Springer New York, 2008.

[77] L.H. Rowen. *Ring Theory*. Number v. 1 in Pure and Applied Mathematics. Academic Press, 1988.

[78] L.H. Rowen. *Ring Theory*. Number v. 2 in Pure and Applied Mathematics. Academic Press, 1988.

[79] L.H. Rowen. *Graduate Algebra: Noncommutative View*. Graduate Algebra. American Mathematical Society, 2008.

[80] D.J. Saltman. *Lectures on Division Algebras*. Number 94 in CBMS Regional Conference Series. American Mathematical Soc., 1999.

[81] J.T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *J. ACM*, 27(4):701–717, October 1980.

[82] J.P. Serre. *Local Fields*. Graduate Texts in Mathematics. Springer New York, 1995.

[83] M. Staromiejski. Polynomial-time locality tests for finite rings. *Journal of Algebra*, 379(0):441–452, 2013.

[84] A. Storjohann. Computation of Hermite and Smith Normal Forms of Matrices. Master's thesis, Department of Computer Science, University of Waterloo, 1994.

[85] A. Storjohann. Near optimal algorithms for computing smith normal forms of integer matrices. In *Proceedings of the 1996 International Symposium on Symbolic and Algebraic Computation*, ISSAC '96, pages 267–274, New York, NY, USA, 1996. ACM.

[86] A. Storjohann. *Algorithms for Matrix Canonical Forms*. PhD thesis, Department of Computer Science, Swiss Federal Institute of Technology, 2000.

[87] M. Szymik. The Brauer group of Burnside rings. *Journal of Algebra*, 324(9):2589–2593, 2010.

[88] J.A. Wood. Duality for modules over finite rings and applications to coding theory. *American Journal of Mathematics*, 121:555–575, 1999.

[89] R. Zippel. Probabilistic algorithms for sparse polynomials. In *Proceedings of the International Symposium on on Symbolic and Algebraic Computation*, EUROSAM '79, pages 216–226, London, UK, 1979. Springer-Verlag.

Index

NP, 3
NP-complete, 3, 51, 72
NP-hard, 3, 70
NP-intermediate, 3, 53, 70
P, 3

algebra, 5
  Azumaya, 80
  finite-étale, 80
  Frobenius, 16
  separable, 76
  separable projective, 82, 86
  strongly separable, 95
  symmetric, 16, 95, 111

basis representation, 46
bimodule, 6, 13, 77
block decomposition, 10
Brauer group, 81

character functor, 17, 69
complexity class, 2
coprime base algorithm, 4, 41

dual basis, 91
duality, 69

essential extension, 14

Fitting's lemma, 7, 58
Fundamental theorem of finite abelian groups, 26

generator, 16

Hermite normal form, 23

idempotents, 9
injective hull, 15, 69
injectivity lift, 84, 111

Jacobson radical, 8, 50, 59, 98
  approximation of, 96, 103, 111

Krull-Remak-Schmidt theorem, 7

lattice, 20
  LLL, 21

MeatAxe, 55
module, 5
  artinian, 6
  finite length, 6
  finitely presented, 13
  flat, 13
  injective, 12, 69, 70, 72
  left/right-regular, 6
  noetherian, 6
  projective, 11, 40, 42, 67, 71, 72, 82, 85
  semisimple, 6
  simple, 6, 67

Nakayama's lemma, 11
nil
  ideal, 5, 56
nilpotent
  ideal, 5, 8, 56
nonsingular matrix completion, 72

progenerator, 17, 80
projective cover, 14, 68
projectivity lift, 83, 111

rank
  Hattori-Stallings, 91
  of a projective module, 13
ring, 4
  centre of, 4, 49
  characteristic of, 4, 49
  connected, 9, 86
  division ring, 7
  Galois, 88
  generalised prime subring, 85, 111
  left-artinian, 56
  left/right artinian, 7
  left/right noetherian, 7
  local, 9
  prime subring, 5, 49
  quasi-Frobenius, 15, 70, 84, 111
  semilocal, 9
  semiprimary, 9
  semisimple, 8, 55, 78, 84
  separable, 84, 111
  simple, 5, 66
  Witt, 86
running time, 1

Schur's lemma, 10
  converse Schur, 11
Smith normal form, 25
splitter, 57
superfluous submodule, 14

trace, 92, 95
  of a projective module, 91
  trace radical, 94, 98
Turing machine, 1

Wedderburn theorem, 8

Algorithms for finite rings – Abstract

In this thesis we are interested in describing algorithms that answer questions arising in ring and module theory. Our focus is on deterministic polynomial-time algorithms and rings and modules that are finite.

The first three chapters prepare the ground for the rest of the text by introducing the underlying ring and module theory, and providing a compendium of algorithms that allow us to perform basic computations with finite abelian groups and finite rings.

The first main result of this thesis concerns the module isomorphism problem: we describe two distinct algorithms that, given a finite ring R and two finite R-modules M and N, determine whether M and N are isomorphic. If they are, the algorithms exhibit such an isomorphism.

In addition, we show how to compute a set of generators of minimal cardinality for a given module, and how to construct projective covers and injective hulls. We also describe tests for module simplicity, projectivity, and injectivity, and constructive tests for existence of surjective module homomorphisms between two finite modules, one of which is projective. As a negative result, we show that the problem of testing for existence of injective module homomorphisms between two finite modules, one of which is projective, is NP-complete.

The last part of the thesis is concerned with finding a good working approximation of the Jacobson radical of a finite ring, that is, a two-sided nilpotent ideal such that the corresponding quotient ring is "almost" semisimple. The notion we use to approximate semisimplicity is that of *separability*.

Algorithmes pour les anneaux finis – Résumé

Cette thèse s'attache à décrire des algorithmes qui répondent à des questions provenant de la théorie des anneaux et des modules. Nous restreindrons essentiellement notre étude à des algorithmes déterministes, en temps polynomial, ainsi qu'aux anneaux et modules finis.

Les trois premiers chapitres préparent le terrain pour le reste du texte. Nous y rappelons les notions nécessaires de la théorie des modules et des anneaux, et nous présentons une collection d'algorithmes qui permettent de réaliser des calculs dans les groupes abéliens et les anneaux finis.

Le premier des principaux résultats de cette thèse concerne le problème de l'isomorphisme entre modules : nous décrivons deux algorithmes distincts qui, étant donné un anneau fini R et deux R-modules M et N finis, déterminent si M et N sont isomorphes. S'ils le sont, les deux algorithmes exhibent un tel isomorphisme.

De plus, nous montrons comment calculer un ensemble de générateurs de taille minimale pour un module donné, et comment construire des couvertures projectives et des enveloppes injectives. Nous décrivons ensuite des tests mettant en évidence le caractère simple, projectif ou injectif d'un module, ainsi qu'un test constructif de l'existence d'un homomorphisme de modules surjectif entre deux modules finis, l'un d'entre eux étant projectif. Par contraste, nous montrons le résultat négatif suivant : le problème consistant à tester l'existence d'un homomorphisme de modules injectif entre deux modules, l'un des deux étant projectif, est NP-complet.

La dernière partie de cette thèse concerne le problème de l'approximation du radical de Jacobson d'un anneau fini. Il s'agit de déterminer un idéal bilatère nilpotent tel que l'anneau quotient correspondant soit "presque" semi-simple. La notion de "semi-simplicité approchée" que nous utilisons est la *séparabilité*.

Algoritmen voor eindige ringen – Samenvatting

We zijn in dit proefschrift geïnteresseerd in het beschrijven van algoritmen voor problemen over ringen en modulen. De nadruk ligt op deterministische algoritmen die in polynomiale tijd werken, en op ringen en modulen die eindig zijn.

De eerste drie hoofdstukken leggen de basis voor de rest: we behandelen voorkennis over ringen en modulen, en stellen een compendium van algoritmen samen die ons basisbewerkingen op eindige abelse groepen en eindige ringen laten uitvoeren.

Het eerste hoofdresultaat van dit proefschrift betreft het isomorfieprobleem voor modulen: we beschrijven twee verschillende algoritmen die voor een eindige ring R en twee eindige R-modulen M en N, in polynomiale tijd beslissen of er een R-moduulisomorfisme tussen M en N bestaat. Als er zo'n isomorfisme bestaat, wordt het door de algoritmen berekend.

Verder laten we zien hoe men een verzameling voortbrengers van minimale grootte voor een gegeven moduul kan berekenen, en hoe men projectieve overdekkingen en injectieve omhulsels kan berekenen. We beschrijven ook methoden om te toetsen of een moduul simpel, projectief of injectief is, en constructieve tests voor het bestaan van een surjectief moduulhomomorfisme tussen twee modulen, waarvan er een projectief is. Als negatief resultaat, laten we zien dat testen of er een injectief moduulhomomorfisme bestaat tussen twee modulen, waarvan er een projectief is, NP-volledig is.

Het laatste deel van dit proefschrift is erop gericht, een goede benadering voor het Jacobson-radicaal te vinden, dat wil zeggen, een tweezijdig nilpotent ideaal waarvan de resulterende quotiëntring "bijna" semisimpel is. Het begrip dat we gebruiken om semisimpliciteit te benaderen is *separabiliteit*.

Acknowledgements

I would like to thank...

Professor Hendrik Lenstra, for his inspiring guidance.
Professor Karim Belabas, for being so generous with his time.
Dr. Gábor Ivanyos, for many useful comments regarding my work.
Professor Peter Stevenhagen.
Professor Bart de Smit.
Raphael Hochard, for the French translation of the abstract of this thesis (page 126).

my parents.

Dr. Radu Gaba, Mihai Bălună and Professor Oliver Riordan, my mathematical parents.

Erik Thörnblad.
Pınar Kılıçer, Dino Festi, Martin Djukanovic, Abtien Javanpeykar and Valerio Dose.
Djordjo Milovic.
Maarten Kampert.

Lucian Corchiş.
Dr. Alvaro Guevara.
Romina Lupşeneanu.

CV

Iuliana Ciocănea-Teodorescu was born June 2nd 1990 in Bucharest, Romania. From 1997 to 2004 she was simultaneously a student of the Goethe German College and the George Enescu Music High School (piano studies), in Bucharest (Romania). In 2009 she graduated from the Mihai Viteazul National College in Bucharest.

Between 2009 and 2013 she was a student of the University of Oxford (UK), where she completed her bachelor's and master's studies in mathematics. Her master's thesis, entitled "Of Galois groups in polynomial time" was supervised by Professor Jonathan Pila.

In 2012 she obtained a DAAD RISE scholarship to undertake a research internship of three months within the Section of Systems Neuroscience, Faculty of Medicine, Dresden University of Technology, Germany.

In 2013 she was awarded an ALGANT-DOC scholarship and started her PhD under the joint supervision of Hendrik W. Lenstra (Leiden University, Netherlands) and Karim Belabas (University of Bordeaux, France).