BACHELOR'S THESIS COMPUTING SCIENCE

RADBOUD UNIVERSITY NIJMEGEN

An Algebraic Cryptanalysis of the ATRAPOS Permutations

Author:
Matthijs Poppe
s1060890

First supervisor/assessor:
Prof. dr. Joan Daemen

Second assessor:
Dr. Simona Samardjiska

August 21, 2025 (last revised October 6, 2025)

Abstract

ATRAPOS is a novel (still in development) family of cryptographic round-based permutations for use in the sponge construction. The ATRAPOS permutations operate natively on elements of $\mathbb{F}_p$ (where $p > 2$ is a prime number). ATRAPOS is designed to provide an efficient alternative for SHA3 as used in the post-quantum asymmetric cryptographic algorithms KYBER and DILITHIUM on platforms where hardware acceleration for multiplication in $\mathbb{F}_p$ is available, where either $p = 3329$ (for KYBER) or $p = 8380417$ (for DILITHIUM).

This thesis investigates the security of ATRAPOS against algebraic attacks using Gröbner basis techniques. To this end, we model the ATRAPOS permutations using sparse systems of polynomials. The complexity of algebraic attacks is determined by a quantity called the "ideal degree" of the ideal generated by these polynomials. We find that the top homogeneous parts of the polynomials corresponding to a single round of ATRAPOS form a so-called "regular sequence" of quadratic polynomials. This property allows us to compute the ideal degree for any number of rounds. Ultimately, we estimate that algebraic attacks against ATRAPOS require $2^{\omega \ell R}$ field operations in $\mathbb{F}_p$ (additions and multiplications), where $2 \leq \omega \leq 3$ is the matrix multiplication exponent, $\ell \geq 3$ is a quantity that depends on p, and R is the number of rounds. Based on this estimate, we determine that at least $R = 4$ (for KYBER) or $R = 10$ (for DILITHIUM) rounds are needed to obtain 128 bits of security against algebraic attacks. Findings from small-scale experiments are consistent with this theorized complexity.

Acknowledgements

First of all, I would like to express my gratitude to my supervisor, Joan Daemen, for introducing me to the ATRAPOS team, for his valuable feedback on my drafts, and for his guidance in writing this thesis. I would also like to thank the ATRAPOS team (Joan Daemen, Silvia Mella, Konstantina Miteloudi, and Morten Øygarden) for allowing me to work on the interesting and challenging problem of analyzing the security of ATRAPOS against algebraic attacks. It has been a great experience working with you, and I look forward to our continued collaboration. In particular, I would like to thank Morten Øygarden for answering my countless questions and for the insightful discussions we had. Lastly, I would like to thank Simona Samardjiska for being my second assessor.

Notation

K	Arbitrary field
$\mathbb{F}_p$	Finite field of size p
$V \oplus W$, $\bigoplus_i V_i$	Direct sum (Definition 1)
$\mathcal{Q} = V/W$	Quotient space of V modulo W (Definition 2)
$[v]_{\mathcal{Q}}$, $[v]$	Equivalence class of v in the quotient space $\mathcal{Q}$ (Definition 2)
$\mathcal{R} = K[x_1, \ldots, x_n]$	Polynomial ring (Definition 4)
f_{top}	Top homogeneous part of f (Definition 6)
$\text{sqfree}(\mathbf{x}^\alpha)$	Square-free part of $\mathbf{x}^\alpha$ (Definition 8)
$\mathcal{SF}$	K-vector space generated by square-free monomials (Definition 8)
$\mathbf{f} = (f_1, \ldots, f_s)$	Polynomial sequence (Definition 10)
$\mathbf{g} \circ \mathbf{f}$	Composition of polynomial sequences (Definition 11)
$\mathbf{f}^{(i)}$	i-th iteration of the polynomial sequence $\mathbf{f}$ (Definition 11)
$V(f_1, \ldots, f_s)$	Affine variety of $f_1, \ldots, f_s$ (Definition 13)
$\mathcal{I}, \mathcal{J}$	Ideal (Definition 15)
$\langle f_1, \ldots, f_s \rangle$	Ideal generated by $f_1, \ldots, f_s$ (Lemma 17)
$d_{\mathcal{I}}$	Ideal degree of $\mathcal{I}$ (Definition 23)
$\cong$	K-vector space or K-algebra isomorphism (Definition 26)
$\text{LM}(f)$	Leading monomial of f (Definition 35)
$\text{LT}(f)$	Leading term of f (Definition 35)
$\text{LT}(\mathcal{I})$	Leading term ideal of $\mathcal{I}$ (Definition 37)
ω	Matrix multiplication exponent (see discussion below Proposition 46)
$\text{HS}_{\mathcal{R}/\mathcal{I}}(t)$	Hilbert series of $\mathcal{R}/\mathcal{I}$ (Definition 52)
ℓ	Width of the 2D ATRAPOS state representation (Section 3.1)
$\overline{n}$	Unique $m \in \{1, \ldots, \ell\}$ such that $n - m \in \ell\mathbb{Z}$ (Section 4.1)
$\mathcal{ASF}$	K-vector space generated by almost square-free monomials (Definition 80)

Contents

1 Introduction **5**
- 1.1 Background . . . 5
- 1.2 Problem Statement . . . 6
- 1.3 Outline . . . 7

2 Preliminaries **8**
- 2.1 Linear Algebra . . . 8
- 2.2 Polynomials . . . 9
- 2.3 Varieties and Ideals . . . 11
- 2.4 Algebras over Fields . . . 15
- 2.5 Monomial Orderings . . . 16
- 2.6 Gröbner Bases . . . 19
- 2.7 Elimination Theory . . . 20
- 2.8 Solving Polynomial Systems . . . 21
- 2.9 Homogeneous Ideals . . . 23
- 2.10 Regular Sequences . . . 25
- 2.11 Security Level . . . 28
- 2.12 Hash Functions . . . 28
- 2.13 The Sponge Construction . . . 29

3 ATRAPOS-SPONGE **32**
- 3.1 Specification . . . 32
  - 3.1.1 The ATRAPOS Permutations . . . 33
- 3.2 CICO problem . . . 33
  - 3.2.1 Polynomial Modeling . . . 35

4 Single-Round Analysis **37**
- 4.1 Polynomial System . . . 37
- 4.2 Direct Sum Claim . . . 40
- 4.3 Proof Sketch for the Direct Sum Claim . . . 42
- 4.4 Proving the Direct Sum Claim . . . 46

5 Multi-Round Analysis **53**
5.1 Introduction 53
5.2 Compositions of Homogeneous Regular Systems 55
5.3 Small Perturbations of Regular Systems 59
5.4 Compositions of Perturbed Regular Systems 62
5.5 Minimal Number of Rounds for ATRAPOS 63

6 Experimental Results **65**

7 Related Work **69**

8 Conclusions **71**
8.1 Future Research 71

A Code **77**
A.1 Hilbert Series Computation (`hilbert_series.py`) 77
A.2 ATRAPOS Implementation (`atrapos.mag`) 77
A.3 Experiments (`experiments.mag`) 78

Chapter 1

Introduction

1.1 Background

The security of many of the classical asymmetric cryptographic algorithms used today (e.g. RSA and (elliptic curve) Diffie-Hellman) relies on problems which are believed to be intractable for classical computers, but can be efficiently solved by quantum computers. For example, recovering an RSA private key from an RSA public key is as hard as computing the prime factorization of a large integer [KL20, Section 9.2]. Performing RSA private key operations (such as decrypting a ciphertext) is therefore at most as hard as prime factorization. Similarly, the security of the Diffie-Hellman key exchange is based on the assumption that the discrete log problem is hard. While prime factorization and the discrete log problem are believed to be hard for classical computers, [Sho94] shows that quantum computers could solve these problems efficiently: on a sufficiently large quantum computer, these problems can be solved in time polynomial in $\log N$ (when factoring an integer N) or $\log p$ (in case of the discrete log problem on a group of order p).

In light of the threat of quantum computers to classical asymmetric cryptography, researchers have been developing and analyzing cryptographic system resistant to attacks by both classical and quantum computers, known as *post-quantum cryptography*. In 2016, NIST issued a call for proposals for post-quantum cryptography for two reasons [Nat16]. First, while no quantum computer has been built yet that is powerful enough to break practical cryptographic systems, there has been "noticeable progress in the development of quantum computers ... that have the potential to scale up to larger systems" [Nat16]. Second, the transition from classical cryptography to post-quantum cryptography will likely require significant effort and time. Additionally, private information that remains sensitive over a long period of time (e.g. medical data or state secrets) may be susceptible to "harvest now, decrypt later" attacks. In these attacks, information encrypted using keys established with classical asymmetric cryptography is harvested with the intention to decrypt it at a later time when

large-scale quantum computers have become accessible to the attacker. It is therefore prudent to develop and standardize post-quantum cryptography long before large-scale quantum computers are built.

Two of the post-quantum cryptography proposals that ended up being standardized by the U.S. National Institute of Standards and Technology (NIST) are the *key encapsulation mechanism* KYBER [Ava+21] (standardized as ML-KEM in FIPS 203 [Nat24b]) and the digital signature scheme DILITHIUM [Duc+21] (standardized as ML-DSA in FIPS 204 [Nat24a]).

For completeness, a key encapsulation mechanism allows a sender to generate a random secret key, which will be encrypted using the public key of the receiver (encapsulation). The receiver can decrypt this secret key using their private key (decapsulation). The secret key can then be used for secure communication using symmetric cryptography. A digital signature scheme allows a signer to sign a message using their private key. A verifier can then verify the signature using the public key of the signer.

Internally, KYBER and DILITHIUM use the SHA3 primitives SHA3-256, SHA3-512, SHAKE-128, and SHAKE-256, all of which are based on the extendable-output function (XOF) KECCAK [BDPV11b] and standardized by NIST in FIPS 202 [Nat15]. The primitives are used for several purposes, such as seed expansion and hashing. As shown by the pqm4 framework for benchmarking post-quantum cryptography on ARM Cortex-M4 CPUs [Kan+], KYBER and DILITHIUM spend a significant amount of CPU cycles on the SHA3 primitives. For example, during key generation (for both KYBER and DILITHIUM), typically 75% or more of the total CPU cycles are spent on SHA3. For other operations, the cycles spent on SHA3 range roughly between 60% and 80%. Thus, replacing the SHA3 primitives in KYBER and DILITHIUM by a more efficient alternative could result in a potentially significant speedup of KYBER and DILITHIUM. This insight prompted the development of ATRAPOS-SPONGE.

1.2 Problem Statement

ATRAPOS-SPONGE [DMMØ25] is a novel XOF, which is currently still in development and has not yet been published. It is designed to be an efficient replacement for SHA3 in KYBER and DILITHIUM on platforms where hardware acceleration for multiplication in $\mathbb{F}_p$ is available, where either $p = 3329$ (for KYBER) or $p = 8380417$ (for DILITHIUM). Internally, KYBER and DILITHIUM operate on elements of $\mathbb{F}_p$. ATRAPOS-SPONGE accomodates this by natively operating on elements of $\mathbb{F}_p$, as opposed to the SHA3 primitives, which operates on bits. Like the SHA3 primitives, ATRAPOS-SPONGE is based on the sponge construction. The permutation used inside ATRAPOS-SPONGE is called ATRAPOS.

[BDPV11a] shows that sponge constructions are computationally indistinguishable from random oracles, assuming that they are instantiated with a random permutation. It follows then, that the security of a real-world sponge con-

struction largely depends on the security of the permutation used within the sponge construction. In [BDPV11a], several *structural distinguishers* for permutations are listed, which could potentially be used in attacks against sponge constructions. In this thesis, we will be interested in a structural distinguisher called the constrained-input constrained-output (CICO) problem, which is strongly related to preimage attacks. The goal of this thesis is to analyze the complexity of solving a specific CICO problem (defined in Section 3.2).

Since the ATRAPOS permutations are multivariate polynomials (as a function of the digits of the input state), solving the CICO problem amounts to solving a system of polynomial equations. We will see that the complexity of solving these systems is determined by a quantity called the *ideal degree*. We will show that the ideal degrees related to ATRAPOS are maximal with respect to the number of multiplications performed. Moreover, we will show that solving the CICO problem requires an estimated $2^{2\ell R}$ field operations in $\mathbb{F}_p$, where ℓ is a fixed parameter related to the state size ($\ell = 17$ for KYBER and $\ell = 7$ for DILITHIUM) and R denotes the number of rounds. The complexity of the CICO problem can therefore be efficiently increased by increasing the number of rounds (up to the point where exhaustive search is the optimal attack).

1.3 Outline

The thesis is structured as follows. Chapter 2 covers the mathematical and cryptographic underpinnings of the cryptanalysis of ATRAPOS. Chapter 3 gives a specification of ATRAPOS-SPONGE, and defines the concrete CICO problem to be analyzed in this thesis. In Chapter 4, we derive the ideal degree corresponding to a single round of ATRAPOS. Chapter 5 extends these results to multiple rounds of ATRAPOS and, using these results, we derive a minimum number of rounds to achieve 128 bits of security (with respect to the CICO problem) when ATRAPOS-SPONGE is used in KYBER and DILITHIUM. In Chapter 6, we compare the theoretical complexity of the CICO problem to experimental results. Chapter 7 discusses related work. We conclude in Chapter 8.

Chapter 2

Preliminaries

2.1 Linear Algebra

We recall two concepts from linear algebra.

The first concept is the notion of direct sums. Decomposing a vector space as a direct sum of subspaces may help us understand the vector space through these subspaces.

Definition 1. Let K be a field, let V be a K-vector space, and let $W_1, W_2 \subseteq V$ be linear subspaces of V. The **sum** of W_1 and W_2 is the set

$$W_1 + W_2 = \{w_1 + w_2 \mid w_1 \in W_1, w_2 \in W_2\}.$$

We say that V is the **direct sum** of W_1 and W_2, denoted $V = W_1 \oplus W_2$, if $W_1 \cap W_2 = \{0\}$ and $W_1 + W_2 = V$.

The vector space V is finite dimensional if and only if both W_1 and W_2 are. In this case, $\dim_K V = \dim_K W_1 + \dim_K W_2$.

The next concept concerns quotient spaces. Intuitively, quotient spaces are obtained from a vector space by mapping similar vectors to the same element in the quotient space.

Definition 2. Let K be a field, let V be a K-vector space, and let $W \subseteq V$ be a linear subspace. We define the equivalence relation $\sim$ by $v \sim w$ if and only if $v - w \in W$. The **equivalence class** of v (with respect to this equivalence relation) is the set $[v] = \{u \in V | v \sim u\}$. We call the set $V/W = \{[v] \mid v \in V\}$ of equivalence classes the **quotient space** of V modulo W. This is again a K-vector space with the operations $\alpha[v] = [\alpha v]$ and $[v] + [w] = [v + w]$ for all $\alpha \in K$ and $v, w \in V/W$.

If V and W are finite-dimensional, then $\dim_K (V/W) = \dim_K V - \dim_K W$.

2.2 Polynomials

While polynomials are intuitively clear, there are some slight notational inconsistencies in the literature. To avoid confusion, we give an explicit definition of monomials and polynomials.

Definition 3. Given n variables $x_1, \dots, x_n$, a **monomial** in $x_1, \dots, x_n$ is any product of the form $x_1^{\alpha_1} \cdots x_n^{\alpha_n}$, where $\alpha_1, \dots, \alpha_n \in \mathbb{Z}_{\geq 0}$ We abbreviate this as $\mathbf{x}^\alpha$, where $\alpha = (\alpha_1, \dots, \alpha_n) \in \mathbb{Z}_{\geq 0}^n$.

If the number n of variables is small, we sometimes write $x, y, z, \dots$ instead of $x_1, x_2, x_3, \dots$.

Definition 4. Let K be a field. A **polynomial** in $x_1, \dots, x_n$ with coefficients in K is any finite K-linear combination of monomials $f = \sum_\alpha c_\alpha \mathbf{x}^\alpha$. The set of all such polynomials forms a ring, which we denote by $K[x_1, \dots, x_n]$.

We call f **univariate** if $n = 1$ and **multivariate** if $n > 1$.

Every $c_\alpha \mathbf{x}^\alpha$ in the finite sum above is called a **term** of f.

There are various ways to define degrees for monomials and polynomials. The most useful notion to us is that of total degrees.

Definition 5. For a monomial $\mathbf{x}^\alpha$ the sum $\sum_{k=1}^n \alpha_k$ is called the **total degree** of $\mathbf{x}^\alpha$, denoted by $|\alpha|$.

For a non-zero polynomial $f = \sum_\alpha c_\alpha \mathbf{x}^\alpha \in K[x_1, \dots, x_n]$, the **total degree** of f (or simply the **degree** of f) is defined as $\deg f = \max\{|\alpha| \mid c_\alpha \neq 0\}$. We define the total degree of the zero polynomial to be $\deg 0 = -\infty$.

The definition of the total degree allows us to define homogeneous polynomials.

Definition 6. We call a polynomial $f = \sum_\alpha c_\alpha \mathbf{x}^\alpha \in K[x_1, \dots, x_n]$ **homogeneous** of degree $d \in \mathbb{Z}_{\geq 0}$ if all monomials $\mathbf{x}^\alpha$ with $c_\alpha \neq 0$ have the same total degree d. We call f **inhomogeneous** if it is not homogeneous.

If $f = \sum_\alpha c_\alpha \mathbf{x}^\alpha \in K[x_1, \dots, x_n]$ is an arbitrary polynomial and $d \geq 0$, we call $f_d = \sum_{|\alpha|=d} c_\alpha \mathbf{x}^\alpha$ the **homogeneous part of degree** d of f. Using this notation, we have $f = \sum_{d \geq 0} f_d$.

Given a polynomial $f \in K[x_1, \dots, x_n]$, we define the **top homogeneous part** of f as $f_{\text{top}} = f_{\deg f}$ if $f \neq 0$ and $f_{\text{top}} = 0$ if $f = 0$.

We sometimes use the terms "linear" and "quadratic" to designate polynomials of degrees 1 and 2, respectively.

Example 7. Consider the polynomial ring $\mathcal{R} = K[x, y, z]$, and the polynomials $f = xyz + y^2z + 1$ and $g = x + y + z$.

The polynomial f is an inhomogeneous polynomial of degree 3. Its non-zero homogeneous parts are $f_{\text{top}} = f_3 = xyz + y^2z$ and $f_0 = 1$.

The polynomial g is a homogeneous polynomial of degree 1 (i.e. a linear polynomial). Its single non-zero homogeneous part is given by $g_{\text{top}} = g_1 = x + y + z$. ◊

Any monomial $\mathbf{x}^\alpha$ is also a polynomial. In this case, both definitions of the total degree agree, since $\deg \mathbf{x}^\alpha = |\alpha|$.

In this thesis, we will distinguish between monomials that contain exponents strictly larger than 1 (e.g. x^2y or x^3yz^5) and monomials whose exponents are at most 1 (e.g. xyz or 1).

Definition 8. We say that a monomial $\mathbf{x}^\alpha$ is **square-free** if $\max_i \alpha_i \leq 1$.

Given a monomial $\mathbf{x}^\alpha$, we define $\text{sqfree}(\mathbf{x}^\alpha) = \mathbf{x}^{\alpha'}$, where $\alpha'_i = \min\{\alpha_i, 1\}$ for all $1 \leq i \leq n$.

Given a polynomial ring $\mathcal{R} = K[x_1, \ldots, x_n]$, we define the K-linear subspace $\mathcal{SF} \subseteq \mathcal{R}$ by $\mathcal{SF} = \text{span}_K \{\mathbf{x}^\alpha \in R \mid \mathbf{x}^\alpha \text{ is square-free}\}$.

The notion of square-free monomials can be extended to terms $c\mathbf{x}^\alpha \in \mathcal{R}$ (where $c \neq 0$) by defining $c\mathbf{x}^\alpha$ to be square-free if $\mathbf{x}^\alpha$ is. Furthermore, we define $\text{sqfree}(c\mathbf{x}^\alpha) = c \cdot \text{sqfree}(\mathbf{x}^\alpha)$.

A polynomial $f = \sum_\alpha c_\alpha \mathbf{x}^\alpha \in K[x_1, \ldots, x_n]$ can be evaluated at any point $(a_1, \ldots, a_n) \in K^n$ by replacing each x_i by a_i. Since K is a field, this yields a well-defined expression. The next definition formalizes this concept.

Definition 9. Let K be a field. Every polynomial $f = \sum_\alpha c_\alpha \mathbf{x}^\alpha$ in $K[x_1, \ldots, x_n]$ induces a **polynomial function** $K^n \to K$ given by

$$f(a_1, \ldots, a_n) = \sum_\alpha c_\alpha a_1^{\alpha_1} \cdots a_n^{\alpha_n}.$$

We stress that the polynomial itself is a *formal* object. For example, let $f_1 = 0$ and $f_2 = x^2 + x$ be polynomials in $\mathbb{F}_2[x]$. Although f_1 and f_2 induce the same function $K^n \to K$, $a \mapsto 0$, they are distinct as polynomials, since the coefficients of their terms are different.

Next, we introduce the notion of polynomial sequences.

Definition 10. Let K be a field and let $\mathcal{R} = K[x_1, \ldots, x_n]$ be a polynomial ring. A vector $\mathbf{f} = (f_1, \ldots, f_s) \in \mathcal{R}^s$ (with $s \geq 1$) is called a **polynomial sequence**.

For every $\alpha = (\alpha_1, \ldots, \alpha_s) \in \mathbb{Z}^s_{\geq 0}$, we define $\mathbf{f}^\alpha = f_1^{\alpha_1} \cdots f_s^{\alpha_s}$.

Polynomials and polynomial sequences can be composed to form new polynomials or polynomial sequences.

Definition 11. Let K be a field, let $\mathcal{R} = K[x_1, \ldots, x_n]$, and let $\mathbf{f} = (f_1, \ldots, f_n) \in \mathcal{R}^n$. For a polynomial $g = \sum_\alpha c_\alpha \mathbf{x}^\alpha \in \mathcal{R}$ we define the **composition** $g \circ \mathbf{f}$ to be

$$g \circ \mathbf{f} = g(\mathbf{f}) = \sum_\alpha c_\alpha \mathbf{f}^\alpha \in \mathcal{R}.$$

Given $\mathbf{g} = (g_1, \ldots, g_s) \in \mathcal{R}^s$, we define

$$\mathbf{g} \circ \mathbf{f} = (g_1 \circ \mathbf{f}, \ldots, g_s \circ \mathbf{f}) \in \mathcal{R}^s.$$

For all $i \geq 0$, the **i-th iteration** of $\mathbf{f}$ is defined as the composition

$$\mathbf{f}^{(i)} = \underbrace{\mathbf{f} \circ \cdots \circ \mathbf{f}}_{i \text{ times}}.$$

We write $f_j^{(i)}$ for the j-th component of $\mathbf{f}^{(i)}$.

2.3 Varieties and Ideals

In this thesis, we will be interested in the systems of polynomial equations that correspond to ATRAPOS (defined in Chapter 3). This section and the following sections discuss techniques to solve polynomial systems.

In general, a polynomial system has the following form.

Definition 12. Let K be a field and $K[x_1, \ldots, x_n]$ a polynomial ring. A **polynomial system** is a system of equations

$$\begin{aligned} f_1(x_1, \ldots, x_n) &= 0 \\ f_2(x_1, \ldots, x_n) &= 0 \\ &\vdots \\ f_s(x_1, \ldots, x_n) &= 0 \end{aligned} \tag{2.1}$$

where $f_1, \ldots, f_s$ are polynomials in $K[x_1, \ldots, x_n]$.

We call $K^n = \{(a_1, \ldots, a_n) \mid a_1, \ldots, a_n \in K\}$ the **affine space** over K. A point $(a_1, \ldots, a_n)$ in this space is called a solution of the polynomial system (2.1) if $f_i(a_1, \ldots, a_n) = 0$ for all $1 \leq i \leq s$. The set of all solutions is often referred to as the affine variety, as formalized in the following definition.

Definition 13. Let K be a field and let $f_1, \ldots, f_s \in K[x_1, \ldots, x_n]$ be polynomials. We call the set

$$V(f_1, \ldots, f_s) = \{(a_1, \ldots, a_n) \in K^n \mid f_i(a_1, \ldots, a_n) = 0 \text{ for all } 1 \leq i \leq s\}$$

the **affine variety** of $f_1, \ldots, f_s$.

In general, a polynomial system may have zero solutions, finitely many solutions, or infinitely many solutions. For example, the system $f_1 = y - x$ in $\mathbb{R}[x, y]$ has infinitely many solutions of the form $(a, a) \in \mathbb{R}^2$, while the system $f_1 = 1$ has no solutions.

The following example demonstrates that the number of solutions may depend on the field over which the polynomials are defined.

Example 14. Consider the system in the variables x, y, z defined by

$$\begin{aligned} x^2 - y &= 0 \\ z^2 - 2 &= 0 \end{aligned} \tag{2.2}$$

If $K = \mathbb{R}$, the affine variety $V\left(x^2 - y, z^2 - 2\right)$ consists of the two parabola $\left\{\left(x, x^2, \sqrt{2}\right) \in \mathbb{R}^3 \mid x \in \mathbb{R}\right\}$ and $\left\{\left(x, x^2, -\sqrt{2}\right) \in \mathbb{R}^3 \mid x \in \mathbb{R}\right\}$. In this case, the system has infinitely many solutions.

If $K = \mathbb{Q}$, however, system (2.2) has no solutions, since there exists no $z \in \mathbb{Q}$ such that $z^2 = \sqrt{2}$.

Finally, if $K = \mathbb{F}_2$, it can be seen that

$$V\left(x^2 - y, z^2 - 2\right) = \{(0,0,0), (1,1,0)\},$$

showing that the polynomial system has finitely many solutions. ◊

Our next topic of study is that of ideals. As it will turn out, there is a close relation between affine varieties and ideals.

Definition 15. Let K be a field and let $\mathcal{I}$ be a subset of $\mathcal{R} = K[x_1, \ldots, x_n]$. We say that $\mathcal{I}$ is an **ideal** if the following properties hold:

1. $0 \in \mathcal{I}$.
2. If $f, g \in \mathcal{I}$, then $f + g \in \mathcal{I}$.
3. If $f \in \mathcal{I}$ and $h \in \mathcal{R}$, then $hf \in \mathcal{I}$.

We call $\mathcal{I} \subseteq \mathcal{R}$ a **proper** ideal of $\mathcal{R}$ if $\mathcal{I} \neq \mathcal{R}$.

Remark 16. It follows directly from Definition 15 that an ideal $\mathcal{I} \subseteq \mathcal{R}$ is proper if and only if $1 \notin \mathcal{I}$.

Trivial examples of ideals in $\mathcal{R} = K[x_1, \ldots, x_n]$ are the zero ideal $\mathcal{I} = \{0\}$ and the polynomial ring $\mathcal{R}$ itself. The latter is not a proper ideal of $\mathcal{R}$.

As another example, let f be an arbitrary polynomial in $\mathcal{R} = K[x_1, \ldots, x_n]$ and consider the set $\mathcal{I} = \{hf \mid h \in \mathcal{R}\}$. It is readily verified that this is an ideal, which we denote by $\mathcal{I} = \langle f \rangle$. This construction can be extended to multiple polynomials as follows.

Lemma 17. *Let K be a field and let $f_1, \ldots, f_s$ be polynomials in $\mathcal{R} = K[x_1, \ldots, x_n]$. The set*

$$\langle f_1, \ldots, f_s \rangle = \left\{ \sum_{k=1}^{s} h_k f_k \mid h_1, \ldots, h_s \in \mathcal{R} \right\}$$

is an ideal.

Proofs of Lemma 17 can be found in many introductory algebra books, see e.g. [CLO15, §1.4, Lemma 3]. Nevertheless, we include a proof here to familiarize the reader with the basic manipulations of ideals.

Proof of Lemma 17. We verify that $\langle f_1, \ldots, f_s \rangle$ satisfies the properties listed in Definition 15. By setting $h_1 = \cdots = h_s = 0$, we see that $0 \in \langle f_1, \ldots, f_s \rangle$. Next, suppose that $f, g \in \langle f_1, \ldots, f_s \rangle$. Then there exist $h_1, \ldots, h_s$ and $h'_1, \ldots, h'_s$, all in $\mathcal{R}$, such that $f = \sum_{k=1}^{s} h_k f_k$ and $g = \sum_{k=1}^{s} h'_k f_k$. It follows that

$$f + g = \sum_{k=1}^{s} h_k f_k + \sum_{k=1}^{s} h'_k f_k = \sum_{k=1}^{s} \left(h_k + h'_k\right) f_k \in \langle f_1, \ldots, f_s \rangle .$$

Finally, suppose that $f = \sum_{k=1}^{s} h_k f_k \in \langle f_1, \ldots, f_s \rangle$ and $g \in \mathcal{R}$. Then $gf = g \sum_{k=1}^{s} h_k f_k = \sum_{k=1}^{s} (g h_k) f_k \in \langle f_1, \ldots, f_s \rangle$. Since $\langle f_1, \ldots, f_s \rangle$ satisfies properties 1-3 of Definition 15, it is an ideal. □

The set $\langle f_1, \ldots, f_s \rangle$ is often called the **ideal generated by** $f_1, \ldots, f_s$. Alternatively, we say that $f_1, \ldots, f_s$ is a **basis** of $\langle f_1, \ldots, f_s \rangle$.

It is often possible to study an ideal by considering a basis that generates it. This is similar to how vector space bases in linear algebra often characterize vector spaces. Unlike the bases encountered in linear algebra, we do not require the elements of an ideal basis to be independent in some sense. Moreover, different ideal bases may have different cardinalities. This is illustrated by the next example.

Example 18. Let $f_1 = x^3 - y$, $f_2 = x^2 + y$, and $f_3 = y$ be polynomials in $K[x, y]$ and consider the ideal $\mathcal{I} = \langle f_1, f_2, f_3 \rangle$. It is easily verified that adding (a multiple of) a generator to another generator does not change the resulting ideal, so $\mathcal{I} = \langle f_1 + f_3, f_2 - f_3, f_3 \rangle = \left\langle x^3, x^2, y \right\rangle$. We now see that one of the generators, x^3, is a multiple of another, x^2. Thus, $\mathcal{I} = \left\langle x^2, y \right\rangle$. ◊

The following lemma is a first step in understanding the relation between the affine variety $V(f_1, \ldots, f_s)$ of $f_1, \ldots, f_s$ and the ideal $\langle f_1, \ldots, f_s \rangle$ generated by these polynomials.

Lemma 19. *Let K be a field and let $f_1, \ldots, f_s$ be polynomials in $K[x_1, \ldots, x_n]$. Then $(a_1, \ldots, a_n) \in V(f_1, \ldots, f_s)$ if and only if $f(a_1, \ldots, a_n) = 0$ for all $f \in \langle f_1, \ldots, f_s \rangle$.*

Proof. Let $(a_1, \ldots, a_n) \in V(f_1, \ldots, f_s)$. Any $f \in \langle f_1, \ldots, f_s \rangle$ is of the form $\sum_{k=1}^{s} h_k f_k$, so

$$f(a_1, \ldots, a_n) = \sum_{k=1}^{s} h_k(a_1, \ldots, a_n) \cdot f_k(a_1, \ldots, a_n) = 0.$$

Conversely, suppose that $f(a_1, \ldots, a_n) = 0$ for all $f \in \langle f_1, \ldots, f_s \rangle$. Then $f_1, \ldots, f_s \in \langle f_1, \ldots, f_s \rangle$ implies $(a_1, \ldots, a_n) \in V(f_1, \ldots, f_s)$. □

In Example 18 we saw that $f_1 = x^3 - y, f_2 = x^2 + y, f_3 = y$ and $g_1 = x^2, g_2 = y$ are two sets of polynomials that generate the same ideal. A simple computation shows that the affine varieties of these sets are the same as well, since both varieties are equal to $\{(0, 0)\} \subseteq K^2$. An immediate consequence of Lemma 19 is that this principle holds in general.

Lemma 20. *Let K be a field and let $f_1, \ldots, f_s$ and $g_1, \ldots, g_t$ be polynomials in $K[x_1, \ldots, x_n]$. If $f_1, \ldots, f_s$ and $g_1, \ldots, g_t$ generate the same ideal, i.e. $\langle f_1, \ldots, f_s \rangle = \langle g_1, \ldots, g_t \rangle$, then $V(f_1, \ldots, f_s) = V(g_1, \ldots, g_t)$.*

Proof. Let $\mathcal{I} = \langle f_1, \ldots, f_s \rangle = \langle g_1, \ldots, g_t \rangle$. Then

$$\begin{aligned}(a_1, \ldots, a_n) \in V(f_1, \ldots, f_s) &\iff f\,(a_1, \ldots, a_n) = 0 \text{ for all } f \in \mathcal{I} \\ &\iff (a_1, \ldots, a_n) \in V(g_1, \ldots, g_s)\end{aligned}$$

□

Lemma 20 shows that, given a variety $V(f_1, \ldots, f_s)$, we may replace the polynomial system $f_1(x_1, \ldots, x_n) = \cdots = f_s(x_1, \ldots, x_n) = 0$ by another (possibly easier to solve) system $g_1(x_1, \ldots, x_n) = \cdots = g_t(x_1, \ldots, x_n) = 0$, as long as the polynomials in both systems generate the same ideal. This prompts us to look for an "optimal" basis $g_1, \ldots, g_t$, i.e. a basis that allows us to easily solve the polynomial system. As we will see later, *Gröbner bases* are exactly the kinds of bases we are looking for.

We conclude this section by generalizing affine varieties of finite sets $f_1, \ldots, f_s$ to affine varieties of arbitrary ideals $\mathcal{I}$.

Definition 21. Let K be a field and let $\mathcal{I}$ be an ideal of the polynomial ring $K[x_1, \ldots, x_n]$. The **affine variety** of $\mathcal{I}$ is the set

$$V(\mathcal{I}) = \{(a_1, \ldots, a_n) \mid f\,(a_1, \ldots, a_n) = 0 \text{ for all } f \in \mathcal{I}\}\,.$$

If $\mathcal{I} = \langle f_1, \ldots, f_s \rangle$, then Lemma 19 implies $V(\mathcal{I}) = V(f_1, \ldots, f_s)$.

Remark 22. Although we will not prove it here, the *Hilbert Basis Theorem* states that every ideal $\mathcal{I} \subseteq \mathcal{R} = K[x_1, \ldots, x_n]$ is finitely generated [CLO15]. Consequently, for every ideal $\mathcal{I} \subseteq \mathcal{R}$ there exist $f_1, \ldots, f_s \in \mathcal{R}$ such that $V(\mathcal{I}) = V(f_1, \ldots, f_s)$.

The next definition plays a key role in reasoning about the complexity of solving polynomial systems.

Definition 23. Let K be a field. Given an ideal $\mathcal{I}$ in the polynomial ring $\mathcal{R} = K[x_1, \ldots, x_n]$, we define the **ideal degree** of $\mathcal{I}$ to be the vector space dimension $d_{\mathcal{I}} = \dim_K(\mathcal{R}/\mathcal{I})$.

We call $\mathcal{I}$ a **zero-dimensional** ideal if $d_{\mathcal{I}} < \infty$.

An example of a zero-dimensional ideal is $\mathcal{I} = \left\langle x^2, y \right\rangle \subseteq \mathbb{C}[x, y]$, since $\mathbb{C}[x, y]/\mathcal{I} = \{[1], [x]\}$ is finite dimensional over $\mathbb{C}$. The ideal $\mathcal{J} = \langle y \rangle \subseteq \mathbb{C}[x, y]$ is not zero-dimensional, since $\mathbb{C}[x, y]/\mathcal{J} = \left\{[1], [x], \left[x^2\right], \ldots\right\}$.

The ideal degree of an ideal $\mathcal{I}$ is related to the size of its affine variety $V(\mathcal{I})$.

Proposition 24. *Let K be a field and $\mathcal{I} \subseteq K[x_1, \ldots, x_n]$ an ideal. If $\mathcal{I}$ is zero-dimensional, then the affine variety $V(\mathcal{I})$ is finite and contains at most $d_{\mathcal{I}}$ points.*

Conversely, if $V(\mathcal{I})$ is finite and K is algebraically closed, then $\mathcal{I}$ is zero-dimensional.

Proof. Follows from [CLO15, §5.3, Theorem 6] and [CLO15, §5.3, Proposition 7]. □

Proposition 24 explains where zero-dimensional ideals derive their name from: if $\mathcal{I}$ is a zero-dimensional ideal, then $V(\mathcal{I})$ is finite and forms a so called "zero-dimensional" variety. (See [CLO15, Chapter 5] for a formal study on assigning dimensions to varieties.)

An upper bound for $d_{\mathcal{I}}$ is given by the Bézout bound.

Proposition 25 (Bézout bound, [KLR24, Theorem 1])**.** *Let K be a field and let $\mathcal{I} = \langle f_1, \dots, f_n \rangle$ be a zero-dimensional ideal of $K[x_1, \dots, x_n]$. Then,*

$$d_{\mathcal{I}} \leq \prod_{k=1}^{n} \deg f_k.$$

2.4 Algebras over Fields

Generally, rings allow for two operations: addition and multiplication. Certain rings, such as polynomial rings, have extra structure and allow for *scalar* multiplication by elements of a field K. For example, let $f = \sum_\alpha c_\alpha \mathbf{x}^\alpha$ be a polynomial in a polynomial ring $K[x_1, \dots, x_n]$. Since every $k \in K$ can be interpreted as a constant polynomial $k \in K[x_1, \dots, x_n]$, the multiplication kf of f by the scalar k is well-defined. The following definition generalizes this principle.

Definition 26. Let $\mathcal{A}$ be a set and K be a field. We call $\mathcal{A}$ a **K-algebra** if $\mathcal{A}$ is a ring which contains K as a subring.

For K-algebras $\mathcal{A}, \mathcal{B}$, we call $\varphi : \mathcal{A} \to \mathcal{B}$ a **homomorphism of K-algebras** (or a **K-algebra homomorphism**) if it is a ring homomorphism that maps every $k \in K$ to k. We call φ a **K-algebra isomorphism** if it is bijective. We say that φ is a **K-algebra automorphism** if $\mathcal{A} = \mathcal{B}$ and φ is a K-algebra isomorphism.

If a K-algebra isomorphism $\varphi : \mathcal{A} \to \mathcal{B}$ exists, we say that $\mathcal{A}$ and $\mathcal{B}$ are **isomorphic as K-algebras**. We denote this using "$\mathcal{A} \cong \mathcal{B}$ as K-algebras".

Examples of K-algebras are the polynomial ring $\mathcal{R} = K[x_1, \dots, x_n]$ and the quotient space $\mathcal{R}/\mathcal{I}$, where $\mathcal{I}$ is an ideal in $K[x_1, \dots, x_n]$. The canonical homomorphism can: $\mathcal{R} \to \mathcal{R}/\mathcal{I}$ which maps f to its equivalence class $[f]$ is an example of a surjective K-algebra homomorphism.

The next proposition is the K-algebra version of the First Isomorphism Theorem for groups, rings, etc. (see Figure 2.1).

Proposition 27 (The First Isomorphism Theorem for K-Algebras)**.** *Let $\phi : \mathcal{A} \to \mathcal{B}$ be a homomorphism of K-algebras. Then $\mathcal{A}/\operatorname{Ker}\phi \cong \operatorname{Im}\phi$ as K-algebras.*

An explicit isomorphism is given by $\varphi : \mathcal{A}/\operatorname{Ker}\phi \to \operatorname{Im}\phi$, defined by $[a] \mapsto \phi(a)$.

Proof. By [DF04, Chapter 7, Theorem 7], the mapping $[a] \mapsto \phi(a)$ is a well-defined ring isomorphism. Since it preserves scalar multiplication, it is also a K-algebra isomorphism. □

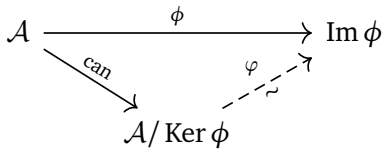

Figure 2.1: The First Isomorphism Theorem for K-Algebras. Here, can: $\mathcal{R} \to \mathcal{R}/\mathcal{I}$ denotes the canonical homomorphism mapping f to its equivalence class $[f]$.

2.5 Monomial Orderings

When computing with polynomials, it is often advantageous to define an ordering on its terms. Consider, for example, the process of dividing $f = x^3 + 1 \in K[x]$ by $g = x^2 + 2 \in K[x]$. Both polynomials are univariate, so there exist unique $q, r \in K[x]$ such that we can write $f = qg + r$, where either $r = 0$ or $\deg r < \deg f$. The "leading term" x^3 of f can be canceled by subtracting $x \cdot g$, since $f - x \cdot g = x^3 + 1 - \left(x^3 + 2x\right) = -2x + 1$. Note that $\deg(-2x + 1) = 1 < \deg g$, so $q = x$ and $r = -2x + 1$.

In this example, we implicitly used the ordering $1 < x < x^2 < \cdots$ and ordered the terms of polynomials in descending order. In terms of this ordering, the polynomial division algorithm systematically cancels the highest order term until a polynomial r is found such that the greatest term of g does not divide the greatest term of r. If instead of canceling highest order terms, we try to cancel arbitrary terms of f, we may not be able to find q and r. This shows that, even in the univariate case, ordering polynomial terms is of great importance.

This notion of ordering the terms of a polynomial can be generalized to the multivariate case. We make two remarks about these generalized orderings. First, note that after collecting terms, a polynomial $f \in K[x_1, \ldots, x_n]$ can be written as $f = \sum_\alpha c_\alpha \mathbf{x}^\alpha$, so an ordering $\mathbf{x}^{\alpha(1)} > \mathbf{x}^{\alpha(2)} > \cdots > \mathbf{x}^{\alpha(m)}$ on the monomials of f automatically yields an ordering $c_{\alpha(1)}\mathbf{x}^{\alpha(1)} > c_{\alpha(2)}\mathbf{x}^{\alpha(2)} > \cdots > c_{\alpha(m)}\mathbf{x}^{\alpha(m)}$ on the (non-zero) terms of f. We are therefore only concerned with *monomial* orderings.

Second, every monomial $\mathbf{x}^\alpha \in K[x_1, \ldots, x_n]$ is identified by its n-tuple of exponents $\alpha \in \mathbb{Z}^n_{\geq 0}$. Any ordering on $\mathbb{Z}^n_{\geq 0}$ therefore induces an ordering on the monomials of $K[x_1, \ldots, x_n]$ and vice versa.

Not every ordering on $\mathbb{Z}^n_{\geq 0}$ yields a useful ordering $>$ on the monomials of $K[x_1, \ldots, x_n]$. A first desirable property of $>$ is that for every two distinct monomials $\mathbf{x}^\alpha$ and $\mathbf{x}^\beta$ we have either $\mathbf{x}^\alpha > \mathbf{x}^\beta$ or $\mathbf{x}^\beta > \mathbf{x}^\alpha$, so that every polynomial

in $K[x_1, \ldots, x_n]$ can be written in a unique order. A second desirable property is that $\mathbf{x}^\alpha > \mathbf{x}^\beta$ implies $\mathbf{x}^\alpha \mathbf{x}^\gamma > \mathbf{x}^\beta \mathbf{x}^\gamma$. This ensures that the "leading term" of a polynomial f changes predictably if we multiply f by another polynomial g.

The requirements above are reformulated by Definition 28. Before presenting this definition, we address two points. First, recall that a partial ordering $\geq$ on an arbitrary set X is a relation $\geq$ which satisfies the following properties:

- Reflexivity: for all $x \in X$, $x \geq x$.
- Transitivity: for all $x, y, z \in X$, $x \geq y$ and $y \geq z$ implies $x \geq z$.
- Anti-symmetry: for all $x, y \in X$, $x \geq y$ and $y \geq x$ implies $x = y$.

Second, given a relation $>$ on an arbitrary set X, we define the relation $\geq$ on X by $x \geq y$ if and only if $x > y$ or $x = y$ for all $x, y \in X$.

Definition 28. A monomial ordering on $K[x_1, \ldots, x_n]$ is an ordering $>$ on $\mathbb{Z}^n_{\geq 0}$ such that:

1. The relation $\geq$ on $K[x_1, \ldots, x_n]$ induced by $>$ is a total ordering. That is, $\geq$ is a partial ordering and for all $\alpha, \beta \in \mathbb{Z}^n_{\geq 0}$, exactly one of $\alpha > \beta$, $\alpha = \beta$, and $\beta > \alpha$ holds.
2. If $\alpha, \beta, \gamma \in \mathbb{Z}^n_{\geq 0}$ with $\alpha > \beta$, then $\alpha + \gamma > \beta + \gamma$.
3. $>$ is a well-ordering. That is, for every non-empty subset $S \subseteq \mathbb{Z}^n_{\geq 0}$ there exists $\alpha \in S$ such that $\beta > \alpha$ for all $\beta \in \mathbb{Z}^n_{\geq 0} \setminus \{\alpha\}$.

If $>$ is a monomial ordering on $\mathbb{Z}^n_{\geq 0}$, we write $\mathbf{x}^\alpha > \mathbf{x}^\beta$ if $\alpha > \beta$.

Remark 29. Property 3 in Definition 28 is often necessary to guarantee termination of algorithms using monomial orderings.

Next, we define two common monomial orderings.

Definition 30 (Lexicographic Ordering)**.** Let $\alpha = (\alpha_1, \ldots, \alpha_n)$ and $\beta = (\beta_1, \ldots, \beta_n)$ be elements of $\mathbb{Z}^n_{\geq 0}$. The lexicographic ordering on $\mathbb{Z}^n_{\geq 0}$ is defined by $\alpha >_{\text{lex}} \beta$ if the leftmost non-zero entry of $\alpha - \beta$ is strictly positive.

Definition 31 (Degree-Reverse-Lexicographic Ordering)**.** Let $\alpha = (\alpha_1, \ldots, \alpha_n)$ and $\beta = (\beta_1, \ldots, \beta_n)$ be elements of $\mathbb{Z}^n_{\geq 0}$. The degree-reverse-lexicographic (DRL) ordering on $\mathbb{Z}^n_{\geq 0}$ is defined by $\alpha >_{\text{drl}} \beta$ if $\deg \alpha > \deg \beta$ or $\deg \alpha = \deg \beta$ and the rightmost non-zero entry of $\alpha - \beta$ is strictly negative.

The following example compares the lexicographic and DRL orderings.

Example 32. Let K be a field and consider the polynomial ring $K[x, y, z, w]$. For the lexicographic ordering we find:

- $x >_{\text{lex}} y >_{\text{lex}} z >_{\text{lex}} w >_{\text{lex}} 1$,

- $x >_{\text{lex}} y^2 >_{\text{lex}} z^3 >_{\text{lex}} w^4$,
- $x^3 >_{\text{lex}} x^2 >_{\text{lex}} x$,
- $xw >_{\text{lex}} yz$,

while for the DRL ordering we find

- $x >_{\text{drl}} y >_{\text{drl}} z >_{\text{drl}} w >_{\text{drl}} 1$,
- $w^4 >_{\text{drl}} z^3 >_{\text{drl}} y^2 >_{\text{drl}} x$,
- $x^3 >_{\text{drl}} x^2 >_{\text{drl}} x$,
- $yz >_{\text{drl}} xw$.

$\Diamond$

Remark 33. The DRL ordering may seem somewhat artificial at first, but often leads to more efficient computations.

Some monomial orderings "preserve degrees" in the following sense.

Definition 34. Let K be a field. We call a monomial ordering $>$ on $K[x_1, \ldots, x_n]$ a **graded ordering** if, for all $\alpha, \beta \in \mathbb{Z}^n_{\geq 0}$, we have $\alpha > \beta$ whenever $|\alpha| > |\beta|$.

It follows directly from Definition 31 and Example 32 that the DRL ordering is a graded ordering, while the lexicographic ordering is not.

Given a monomial ordering, we can finally properly define the notion of leading monomials and leading terms.

Definition 35. Let K be a field and fix a monomial ordering $>$ on $K[x_1, \ldots, x_n]$. Let $f = \sum_\alpha c_\alpha \mathbf{x}^\alpha$ be a non-zero polynomial in $K[x_1, \ldots, x_n]$. The **leading monomial** of f is the monomial $\mathbf{x}^\alpha$ such that α is maximal (with respect to $>$) among all $\alpha' \in \mathbb{Z}^n_{\geq 0}$ with $c_{\alpha'} \neq 0$. The leading monomial is denoted by $\mathrm{LM}(f) = \mathbf{x}^\alpha$.

If $\mathrm{LM}(f) = \mathbf{x}^\alpha$, we call $c_\alpha \mathbf{x}^\alpha$ the **leading term** of f, denoted by $\mathrm{LT}(f) = c_\alpha \mathbf{x}^\alpha$.

Remark 36. Some authors write $\mathrm{LM}_>(f)$ and $\mathrm{LT}_>(f)$ for the leading monomial and leading term of f to emphasize the dependence on the monomial ordering $>$. In practice, the monomial ordering is clear from the context, and we choose to drop the subscript $>$ from the notation.

As we saw in the polynomial division example at the beginning of this section, leading terms play an important role in polynomial computations. Likewise, the leading terms of all polynomials in an ideal $\mathcal{I} \subseteq K[x_1, \ldots, x_n]$ play an important role when trying to understand the ideal.

Definition 37. Let K be a field and fix a monomial ordering $>$ on the polynomial ring $\mathcal{R} = K[x_1, \ldots, x_n]$. For every ideal $\mathcal{I} \subseteq \mathcal{R}$, we define $\mathrm{LT}(\mathcal{I})$ to be the set

$$\mathrm{LT}(\mathcal{I}) = \{\mathrm{LT}(f) \mid f \in \mathcal{I} \setminus \{0\}\}.$$

The **leading term ideal** of $\mathcal{I}$ is the ideal $\langle \mathrm{LT}(\mathcal{I}) \rangle$ generated by $\mathrm{LT}(\mathcal{I})$. Elements of this ideal are finite sums of the form $\sum_{i=1}^{m} h_i \, \mathrm{LT}(f_i)$, where $h_1, \ldots, h_m \in \mathcal{R}$ and $f_1, \ldots, f_m \in \mathcal{I} \setminus \{0\}$.

There is a close relation between the quotient space $\mathcal{R}/\mathcal{I}$ and the leading term ideal $\langle \mathrm{LT}(\mathcal{I}) \rangle$.

Lemma 38 ([CLO15, §5.3, Proposition 4])**.** *Let K be a field and define $\mathcal{R} = K[x_1, \ldots, x_n]$. Fix a monomial ordering on $\mathcal{R}$. For every ideal $\mathcal{I} \subseteq K[x_1, \ldots, x_n]$, $\mathcal{R}/\mathcal{I}$ is isomorphic as a K-vector space to $S = \mathrm{span}_K \{\boldsymbol{x}^\alpha \in \mathcal{R} \mid \boldsymbol{x}^\alpha \notin \langle \mathrm{LT}(\mathcal{I}) \rangle\}$.*

If $\mathcal{I} = \langle f_1, \ldots, f_s \rangle$, then $\mathrm{LT}(f_i) \in \mathrm{LT}(\mathcal{I}) \subseteq \langle \mathrm{LT}(\mathcal{I}) \rangle$ for all $1 \leq i \leq s$. Since $\langle \mathrm{LT}(\mathcal{I}) \rangle$ is an ideal, it follows that

$$\langle \mathrm{LT}(f_1), \ldots, \mathrm{LT}(f_s) \rangle \subseteq \langle \mathrm{LT}(\mathcal{I}) \rangle. \tag{2.3}$$

Perhaps surprisingly, the reverse inclusion does not hold in general.

Example 39. Let K be a field and let $>$ be a monomial ordering on $K[x, y]$ such that $x > y$. Define $\mathcal{I} = \langle f_1, f_2 \rangle \subseteq K[x, y]$, where $f_1 = x + y$ and $f_2 = -x$. Then $f_1 + f_2 = y$, so $y = \mathrm{LT}(f_1 + f_2) \in \langle \mathrm{LT}(\mathcal{I}) \rangle$. However, $\langle \mathrm{LT}(f_1), \mathrm{LT}(f_2) \rangle = \langle x \rangle$, so $y \notin \langle \mathrm{LT}(f_1), \mathrm{LT}(f_2) \rangle$. ◊

2.6 Gröbner Bases

Sets for which the inclusion in Equation (2.3) is in fact an equality are called Gröbner bases.

Definition 40. Let K be a field and fix a monomial ordering $>$ on $K[x_1, \ldots, x_n]$. Let $\mathcal{I} \subseteq K[x_1, \ldots, x_n]$ be an ideal. We say that a subset $G = \{g_1, \ldots, g_t\}$ of $\mathcal{I}$ is a **Gröbner basis** of $\mathcal{I}$ (with respect to $>$) if $\langle \mathrm{LT}(g_1), \ldots, \mathrm{LT}(g_t) \rangle = \langle \mathrm{LT}(\mathcal{I}) \rangle$.

Note that Definition 40 does not require that G is a basis of $\mathcal{I}$. It turns out that this is already encoded by the fact that it is a Gröbner basis.

Proposition 41 ([CLO15, §2.5, Corollary 6])**.** *Let K be a field and fix a monomial ordering $>$ on $K[x_1, \ldots, x_n]$. Every ideal $\mathcal{I} \subseteq K[x_1, \ldots, x_n]$ has a Gröbner basis with respect to $>$. Moreover, every Gröbner basis of $\mathcal{I}$ is a basis of $\mathcal{I}$.*

Gröbner bases can be used to solve several problems related to ideals. One of these problems is the *ideal membership problem*: given polynomials $f, f_1, \ldots, f_s \in K[x_1, \ldots, x_n]$, is $f \in \mathcal{I} = \langle f_1, \ldots, f_s \rangle$? This problem can be solved by computing

a Gröbner basis $G = \{g_1, \ldots, g_t\}$ of $\mathcal{I}$ (with respect to an arbitrary monomial ordering). The remainder on division of f by G using a generalization of the usual polynomial division algorithm is zero if and only if $f \in \mathcal{I}$, see e.g. [CLO15, Chapter 2].

Another application of Gröbner bases is solving polynomial systems. This will be the topic of the next sections.

Remark 42. We note that various algorithms exist to compute a Gröbner basis of $\mathcal{I} = \langle f_1, \ldots, f_s \rangle$ with respect to a monomial ordering $>$, such as Buchberger's algorithm [CLO15, §2.7], or the more performant F4 [Fau99] and F5 [Fau02] algorithms.

Additionally, the *basis conversion* algorithm presented in [FGLM93], known as the FGLM algorithm, converts a Gröbner basis for a zero-dimensional ideal with respect to a monomial ordering $>_1$ to a Gröbner basis with respect to another monomial ordering $>_2$.

2.7 Elimination Theory

Let $\mathcal{I} = \langle f_1, \ldots, f_s \rangle$ be an ideal in $K[x_1, \ldots, x_n]$. By Proposition 41, there exists a Gröbner basis $G = \{g_1, \ldots, g_t\}$ with respect to any monomial ordering $>$. We observed before that $V(f_1, \ldots, f_s) = V(g_1, \ldots, g_t)$ (Lemma 20). If G is a Gröbner basis with respect to the lexicographic ordering, we can compute $V(g_1, \ldots, g_t)$ using *elimination theory*.

We start by defining the elimination ideals of an ideal $\mathcal{I} \subseteq K[x_1, \ldots, x_n]$. Intuitively, these elimination ideals correspond to polynomial systems whose affine variety is a superset of $V(g_1, \ldots, g_t)$.

Definition 43. Let K be a field and let $\mathcal{I} \subseteq K[x_1, \ldots, x_n]$ be an ideal. For all $0 \leq l \leq n$, the ideal $\mathcal{I}_l \subseteq K[x_{l+1}, \ldots, x_n]$ defined by

$$\mathcal{I}_l = \mathcal{I} \cap K[x_{l+1}, \ldots, x_n]$$

is called the l-th **elimination ideal** of $\mathcal{I}$.

We leave it as an exercise to the reader to verify that for all $0 \leq l \leq n$, $\mathcal{I}_l$ is indeed an ideal.

For Gröbner bases with respect to the lexicographic ordering, we can explicitly compute a basis of $\mathcal{I}_l$.

Theorem 44 (The Elimination Theorem, [CLO15, §3.1, Theorem 2])**.** *Let K be a field, let $\mathcal{I}$ be an ideal of $K[x_1, \ldots, x_n]$, and let $G = \{g_1, \ldots, g_t\}$ be a Gröbner basis of $\mathcal{I}$ with respect to the lexicographic ordering. Then, for all $0 \leq l \leq n$, $G_l = G \cap K[x_{l+1}, \ldots, x_n]$ is a Gröbner basis of the l-th elimination ideal $\mathcal{I}_l$.*

We now present a crude algorithm to compute $V(g_1, \ldots, g_t)$ given a Gröbner basis of $\mathcal{I} \subseteq K[x_1, \ldots, x_n]$ with respect to the lexicographic ordering. The

algorithm uses the Elimination Theorem to reduce the problem of computing $V(g_1, \ldots, g_t)$ to the problem of univariate root finding.

For simplicity, we assume that the field K is finite.

Observe that the elimination ideals (viewed as ideals of $K[x_1, \ldots, x_n]$) form the following ascending chain:

$$\emptyset = \mathcal{I}_n \subseteq \mathcal{I}_{n-1} \subseteq \cdots \subseteq \mathcal{I}_0 = \mathcal{I}.$$

Their affine varieties therefore form a descending chain:

$$K^n = V(\mathcal{I}_n) \supseteq V(\mathcal{I}_{n-1}) \supseteq \cdots \supseteq V(\mathcal{I}_0) = V(\mathcal{I}).$$

We know that $V(\mathcal{I}_n) = K^n$ and compute $V(\mathcal{I}_{n-1}), \ldots, V(\mathcal{I}_0)$ inductively.

Having computed $V(\mathcal{I}_l)$ for some $0 < l < n$, we observe that $G_{l-1} = G \cap K[x_l, \ldots, x_n]$ is a basis of $V(\mathcal{I}_{l-1})$ consisting of polynomials in $K[x_l, \ldots, x_n]$. For every $(a_{l+1}, \ldots, a_n) \in V(\mathcal{I}_l)$, the substitution $(x_{l+1}, \ldots, x_n) = (a_{l+1}, \ldots, a_n)$ in G_{l-1} yields a system of univariate polynomials in $K[x_l]$. The affine variety $V(\mathcal{I}_l)$ is then simply the union of the solution sets corresponding to every possible substitution.

At some point, the algorithm has computed $V(\mathcal{I}_0) = V(\mathcal{I})$. By assumption, $\mathcal{I} = \langle g_1, \ldots, g_t \rangle$, so $V(g_1, \ldots, g_t) = V(\mathcal{I}_0)$.

Remark 45. The algorithm presented here can be adapted to the case where K is infinite as well. Some care should be taken, since the affine varieties are not guaranteed to be finite.

2.8 Solving Polynomial Systems

The algorithm presented in the previous section provides a method to compute the affine variety $V(f_1, \ldots, f_s)$ of arbitrary polynomial systems $f_1, \ldots, f_s$: first compute a Gröbner basis $G = \{g_1, \ldots, g_t\}$ of $\mathcal{I} = \langle f_1, \ldots, f_s \rangle$ with respect to the lexicographic ordering and then use the algorithm above to compute $V(f_1, \ldots, f_s) = V(g_1, \ldots, g_t)$.

In practice, directly computing a Gröbner basis with respect to the lexicographic ordering is computationally expensive (both in time and memory). Polynomial systems corresponding to cryptographic problems often yield a zero-dimensional ideal $\mathcal{I} \subseteq K[x_1, \ldots, x_n]$. For these ideals it is usually more efficient to compute a DRL Gröbner basis first and then use a *basis conversion* algorithm to obtain a lexicographic Gröbner basis. State-of-the-art algorithms therefore use the following approach to compute $V(\mathcal{I})$ [KLR24]:

1. Compute a Gröbner basis of $\mathcal{I}$ with respect to the DRL ordering, using e.g. the F4 [Fau99] or F5 [Fau02] algorithm.

2. Convert the DRL Gröbner basis to a Gröbner basis with respect to the lexicographic ordering, using a basis conversion algorithm such as the FGLM algorithm [FGLM93].

3. Solve for one or more solutions in $V(\mathcal{I})$ using univariate polynomial solving, as described in Section 2.7.

While indirectly computing a lexicographic Gröbner basis from a DRL Gröbner basis is often more efficient than directly computing a lexicographic Gröbner basis, it should be noted that the indirect computation is still computationally expensive. This is what ultimately protects cryptographic primitives against algebraic attacks.

In Chapter 6 we will experimentally determine that, for the polynomial systems corresponding to ATRAPOS, step 1 is negligible compared to the FGLM step. We will therefore not discuss the time complexity of this step. (The interested reader is referred to e.g. [KLR24] for a discussion on the time complexity of step 1.) Moreover, Remark 47 at the end of this section shows that the asymptotic time complexity of step 2 dominates the asymptotic time complexity of step 3. Hence, we will analyze the security of ATRAPOS against algebraic attacks by analyzing the complexity of the FGLM algorithm in step 2. We will also briefly discuss this choice in Subsection 3.2.1.

An upper bound on the time complexity of the FGLM algorithm is given by the following proposition.

Proposition 46 ([FGLM93, Proposition 5.1])**.** *Let K be a field and let $\mathcal{I}$ be a zero-dimensional ideal of $\mathcal{R} = K[x_1, \ldots, x_n]$. If $\mathcal{I}$ has ideal degree $d_\mathcal{I} = \mathcal{R}/\mathcal{I}$, then the FGLM algorithm has a worst-case time complexity of $\mathcal{O}\left(nd_\mathcal{I}^\omega\right)$ field operations in K (i.e. addition and multiplication), where $2 \leq \omega \leq 3$ is the* matrix multiplication exponent *(see below).*

In this thesis, we define the matrix multiplication exponent $2 \leq \omega \leq 3$ in Proposition 46 such that an attacker can multiply two dense $n \times n$ matrices using $\mathcal{O}(n^\omega)$ field operations in K. Common choices include $\omega = 3$ (naïve algorithm), $\omega = \log_2 7 \approx 2.81$ (Strassen algorithm) [Str69], $\omega \approx 2.37$ [WXXZ24], or $\omega = 2$ (lower bound) [KLR24].

While the algorithm in [WXXZ24] has a lower asymptotic time complexity than Strassen multiplication [Str69], it is mainly of theoretical interest, since the hidden multiplicative constant in the $\mathcal{O}$ notation is prohibitively large for practical implementations. In contrast, Strassen multiplication has a hidden multiplicative constant of 4.7 [Str69]. Nevertheless, in many cases, the conservative choice $\omega = 2$ is more reasonable to account for algorithms exploiting *sparsity* in the involved matrices [FGHR14; FM11].

Note that the bound in Proposition 46 is an *upper* bound for the time complexity of the FGLM step. Since the security of ATRAPOS against algebraic attacks depends on this step being hard, we would ideally have a *lower* bound for the time complexity instead. However, no such practical lower bounds are known. In many practical realizations of the FGLM algorithm (or variants thereof), the hidden multiplicative constant in $\mathcal{O}\left(nd_\mathcal{I}^\omega\right)$ is reasonably close to unity, compared to $d_\mathcal{I}^\omega$. It is common to estimate the time complexity of FGLM (with respect to

field operations) using $\mathcal{C}_{\text{FGLM}} = nd_{\mathcal{I}}^{\omega}$ [Alb+19], even when taking $\omega = 2$. This is not a realistic assumption if the involved matrices are dense, but this estimate serves as a security margin to account for attacks that exploit sparsity of the matrices [KLR24].

The lexicographic Gröbner basis obtained by FGLM can be used to find all solutions in the variety of $\mathcal{I}$. In cases where finding a single solution suffices, it may not always be necessary to compute the entire lexicographic Gröbner basis. In these cases, the conservative estimate $\mathcal{C}_{\text{FGLM}} = d_{\mathcal{I}}^{\omega}$ may be more appropriate.

Remark 47. There exist practical algorithms to perform univariate polynomial solving (step 3 of the algorithm outlined above) using $\mathcal{O}\left(d^{1.815}\right)$ field operations, where $d \leq d_{\mathcal{I}}$ [KLR24]. This bound is (asymptotically) lower than the bound for the FGLM algorithm (Proposition 46). It is therefore justified to argue only about the time complexity of steps 1 and 2, when analyzing the complexity of the algorithm above.

2.9 Homogeneous Ideals

In this section, we study ideals that are generated by homogeneous polynomials.

Definition 48. Let K be a field. We call an ideal $\mathcal{I}$ of $K[x_1, \ldots, x_n]$ **homogeneous** if it is generated by a basis of homogeneous polynomials. That is, $\mathcal{I}$ is homogeneous if and only if there exist homogeneous polynomials $f_1, \ldots, f_s \in K[x_1, \ldots, x_n]$ such that $\mathcal{I} = \langle f_1, \ldots, f_s \rangle$.

A trivial example of a homogeneous ideal in $\mathcal{R} = K[x_1, \ldots, x_n]$ is $\mathcal{R}$ itself, since $\mathcal{R} = \langle 1 \rangle$.

Another example of a homogeneous ideal is $\mathcal{I} = \left\langle x + y, y^2 \right\rangle \subseteq \mathcal{R}$. Although $\mathcal{I}$ is a homogeneous ideal, it contains inhomogeneous polynomials, such as $f = xy^2 + x + y$. In this case, the homogeneous components of f are xy^2 and $x + y$, which both belong to $\mathcal{I}$.

The next lemma shows that every polynomial in a homogeneous ideal $\mathcal{I}$ can be written as a sum of homogeneous polynomials in $\mathcal{I}$.

Lemma 49. *Let K be a field and let $\mathcal{I} \subseteq K[x_1, \ldots, x_n]$ be a homogeneous ideal. For each $f \in K[x_1, \ldots, x_n]$, $\mathcal{I}$ contains f if and only if $\mathcal{I}$ contains all homogeneous components of f.*

Proof. Follows from [CLO15, §8.3, Definition 1] and [CLO15, §8.3, Theorem 2]. □

Before we can fully appreciate the significance of Lemma 49, we have to make a definition.

Definition 50. Let K be a field and let $\mathcal{R} = K[x_1, \ldots, x_n]$. For all $m \geq 0$, we define $\mathcal{I}_m$ to be the set of homogeneous polynomials in $\mathcal{I}$ of degree m, together with the zero vector.

Every $\mathcal{I}_m$ is a linear subspace of $\mathcal{R}$. (For all $m \geq 0$, $\mathcal{I}_m$ is not an ideal, unless $\mathcal{I}_m = \{0\}$.)

Remark 51. In Section 2.7, the notation $\mathcal{I}_m$ was used to denote the m-th elimination ideal of $\mathcal{I}$. In the remainder of this thesis, $\mathcal{I}_m$ will refer to the vector space defined in Definition 50 instead. This will also be clear from the context.

To verify that $\mathcal{I}_m$ is a linear subspace of $\mathcal{R}$, let $m \geq 0$. The sum of two polynomials in $\mathcal{R}_m$ is either zero or has degree $m \geq 0$. Likewise, multiplying a polynomial in $\mathcal{R}_m$ by a scalar in K results in a polynomial which is either zero or has degree $m \geq 0$. Since every $\mathcal{I}_m$ contains the zero polynomial, and is closed under addition and scalar multiplication, it is a linear subspace of $\mathcal{R}$.

It follows from Lemma 49 that $\mathcal{I} = \mathcal{I}_0 + \mathcal{I}_1 + \cdots$. Moreover, $\mathcal{I}_i \cap \mathcal{I}_j = \{0\}$ for all distinct $i, j \geq 0$. Thus, $\mathcal{I}$ can be written as the infinite direct sum $\mathcal{I} = \bigoplus_{m \geq 0} \mathcal{I}_m$. As a special case, we can write $\mathcal{R} = \bigoplus_{m \geq 0} \mathcal{R}_m$.

The decomposition of $\mathcal{R}$ and $\mathcal{I}$ into the subspaces $\mathcal{R}_0, \mathcal{R}_1, \ldots$ and $\mathcal{I}_0, \mathcal{I}_1, \ldots$ will help us better understand the relations between $\mathcal{I}$ and $\mathcal{R}$. For example, we have

$$\mathcal{R}/\mathcal{I} = \left(\bigoplus_{m \geq 0} \mathcal{R}_m\right) \Bigg/ \left(\bigoplus_{m \geq 0} \mathcal{I}_m\right) = \bigoplus_{m \geq 0} (\mathcal{R}_m/\mathcal{I}_m).$$

It now follows that $\dim_K(\mathcal{R}_m/\mathcal{I}_m) = \sum_{m \geq 0} \dim_K(\mathcal{R}_m/\mathcal{I}_m)$. A tool which helps us understand the right-hand side of this equation is the Hilbert series.

Definition 52. Let K be a field, let $\mathcal{R} = K[x_1, \ldots, x_n]$, and let $\mathcal{I}$ be a homogeneous ideal of $\mathcal{R}$. The **Hilbert series** of $\mathcal{R}/\mathcal{I}$ is the formal power series

$$\mathrm{HS}_{\mathcal{R}/\mathcal{I}}(t) = \sum_{m=0}^{\infty} \dim_K(\mathcal{R}_m/\mathcal{I}_m) \cdot t^m.$$

From the discussion preceding Definition 52, it is clear that evaluating the Hilbert series of $\mathcal{R}/\mathcal{I}$ at $t = 1$ yields the ideal degree of $\mathcal{I} \subseteq \mathcal{R}$:

$$\mathrm{HS}_{\mathcal{R}/\mathcal{I}}(1) = \sum_{m=0}^{\infty} \dim_K(\mathcal{R}_m/\mathcal{I}_m) = \dim_K(\mathcal{R}/\mathcal{I}) = d_{\mathcal{I}}.$$

The subspaces $\mathcal{I}_m$ can be explicitly described if we have a basis $f_1, \ldots, f_s$ for $\mathcal{I}$ consisting of homogeneous polynomials of the same degree.

Lemma 53. *Let K be a field and let $f_1, \ldots, f_s \in K[x_1, \ldots, x_n]$ be homogeneous polynomials of degree d. If $\mathcal{I} \subseteq K[x_1, \ldots, x_n]$ is the ideal generated by $f_1, \ldots, f_s$, then*

$$\mathcal{I}_m = \mathrm{span}_K\{\boldsymbol{x}^\alpha f_i \mid 1 \leq i \leq s \text{ and } \deg \alpha = m - d\}$$

for all $m \geq 0$. In particular,

$$\mathcal{I}_0 = \mathcal{I}_1 = \cdots = \mathcal{I}_{d-1} = \{0\}.$$

Proof. Let $S_m = \operatorname{span}_K \{\mathbf{x}^\alpha f_i \mid 1 \le i \le s \text{ and } \deg \alpha = m - d\}$ for all $m \ge 0$. By induction on m, we show that $\mathcal{I}_m = S_m$ for all $m \ge 0$. The base case $m = 0$ is trivial, so assume that $\mathcal{I}_m = S_m$ for some $m \ge 0$.

Since $\mathcal{I}$ is an ideal, every $\mathbf{x}^\alpha f_i$ with $\deg \alpha = m + 1 - d$ is an element of $\mathcal{I}$ of degree $m+1$ and therefore belongs to $\mathcal{I}_{m+1}$. It follows that $S_m \subseteq \mathcal{I}_{m+1}$.

To prove the reverse inclusion, let $f \in \mathcal{I}_{m+1}$. If $f = 0$, it is clear that $f \in S_{m+1}$, so suppose that $f \neq 0$. Then, there exist $h_1, \ldots, h_s \in K[x_1, \ldots, x_n]$ such that

$$f = \sum_{i=1}^{s} h_i f_i = \sum_{i=1}^{s} \sum_{j \ge 0} (h_i)_j f_i = \sum_{j \ge 0} \sum_{i=1}^{s} (h_i)_j f_i,$$

where $(h_i)_j$ denotes the j-th degree homogeneous part of h_i. For all $j \ge 0$, $\sum_{i=1}^{s} (h_i)_j f_i$ is either the zero polynomial or a polynomial of degree $j + d$. We know that $\deg f = m+1$, so all terms with $j+d \neq m+1$ must vanish. Therefore, $f = \sum_{i=1}^{s} (h_i)_{m+1-d} f_i$. For all $1 \le i \le s$, $(h_i)_{m+1-d}$ is a (possibly empty) sum of terms with degree $(m + 1 - d)$, so $f = (h_i)_{m+1-d} f_i \in S_{m+1}$. We conclude that $\mathcal{I}_{m+1} \subseteq S_m$. □

2.10 Regular Sequences

In this section, we review the notion of regular sequences of polynomials. Regular sequences will play a key role when we extend the results for a single round of ATRAPOS to multiple rounds of ATRAPOS.

Definition 54. Let K be a field, let $\mathcal{R} = K[x_1, \ldots, x_n]$ be a polynomial ring, and let $\mathcal{I} \subseteq \mathcal{R}$ be an ideal. We call $[f] \in \mathcal{R}/\mathcal{I}$ a **non-zero-divisor** for $\mathcal{R}/\mathcal{I}$ if $[f] \cdot [g] = 0$ in $\mathcal{R}/\mathcal{I}$ implies $[g] = 0$ for all $[g] \in \mathcal{R}/\mathcal{I}$.

We say that the sequence $f_1, \ldots, f_s$ of polynomial in $\mathcal{R}$ is a **regular sequence** if $\langle f_1, \ldots, f_s \rangle \neq \mathcal{R}$ and if for all $1 \le i \le s$, $[f_i]$ is a non-zero-divisor for $\mathcal{R}/\langle f_1, \ldots, f_{i-1} \rangle$.

Generally, whether the polynomials $f_1, \ldots, f_s \in \mathcal{R}$ form a regular sequence may depend on their order. The next lemma shows that the order is not important if the polynomials are homogeneous.

Lemma 55. *Let K be a field and let $f_1, \ldots, f_s$ in $\mathcal{R} = K[x_1, \ldots, x_n]$ be homogeneous polynomials of degrees $d_1, \ldots, d_s$. Then $f_1, \ldots, f_s$ is a regular sequence if and only if for every permutation σ on $\{1, \ldots, s\}$, $f_{\sigma(1)}, \ldots, f_{\sigma(s)}$ is a regular sequence.*

Proof. Follows from [KR05, Corollary 5.2.17]. □

The next lemma allows us to determine whether a sequence $f_1, \ldots, f_s$ of homogeneous polynomials is a regular sequence, if we know the Hilbert series of $\mathcal{R}/\langle f_1, \ldots, f_s \rangle$.

Lemma 56 ([Sta78, Corollary 3.2])**.** *Let K be a field and let $f_1, \ldots, f_s$ in $\mathcal{R} = K[x_1, \ldots, x_n]$ be homogeneous polynomials of degrees $d_1, \ldots, d_s$. Then*

$$\mathrm{HS}_{\mathcal{R}/\langle f_1,\ldots,f_s\rangle}(t) \geq \frac{1}{(1-t)^n} \cdot \prod_{k=1}^{s} \left(1 - t^{d_k}\right),$$

where equality holds if and only if $f_1, \ldots, f_s$ is a regular sequence. (We define $\sum_{m=0}^{\infty} a_m t^m \geq \sum_{m=0}^{\infty} b_m t^m$ if $a_m \geq b_m$ for all $m \geq 0$.)

Lemma 56 lets us easily compute the ideal degree of an ideal generated by a regular sequence of homogeneous polynomials.

Corollary 57. *Let K be a field and let $f_1, \ldots, f_s$ be a regular sequence in $\mathcal{R} = K[x_1, \ldots, x_n]$ of homogeneous polynomials of degrees $d_1, \ldots, d_s$. Then the ideal $\mathcal{I} = \langle f_1, \ldots, f_s \rangle$ has ideal degree $d_1 \cdots d_s$.*

Proof. By Lemma 56, the Hilbert series of $\mathcal{R}/\mathcal{I}$ is $\mathrm{HS}_{\mathcal{R}/\mathcal{I}}(t) = \prod_{i=1}^{n} \frac{1-t^{d_i}}{1-t}$. Recognizing that, for all $1 \leq i \leq s$, $\frac{1-t^{d_i}}{1-t}$ equals the geometric sum $\sum_{j=0}^{d_i-1} t^j$, we find that $\mathcal{I}$ has ideal degree $\mathrm{HS}_{\mathcal{R}/\mathcal{I}}(1) = \prod_{i=1}^{n} \left(\sum_{j=0}^{d_i-1} 1\right) = d_1 \cdots d_s$. □

The following corollary shows how Lemma 56 implies that the variables $x_1, \ldots, x_n \in K[x_1, \ldots, x_n]$ form a regular sequence. It is not hard to derive the same result using Definition 54 directly, and we leave it as an exercise for the reader to do so.

Corollary 58. *Let K be a field. The variables $x_1, \ldots, x_n \in \mathcal{R} = K[x_1, \ldots, x_n]$ form a regular sequence.*

Proof. Let $\mathcal{I} = \langle x_1, \ldots, x_n \rangle$. For all $m \geq 1$ we have $\mathcal{I}_m = \mathcal{R}_m$, while $\mathcal{I}_0 = \{0\}$. It follows that $\mathrm{HS}_{\mathcal{R}/\mathcal{I}}(t) = 1$. Since $\frac{(1-t)^n}{(1-t)^n}$ is also equal to 1, Lemma 56 implies that $x_1, \ldots, x_n$ form a regular sequence. □

Remark 59. The regular sequence $x_1, \ldots, x_n$ from Corollary 58 is the simplest kind of regular sequence one can have, in the sense that the degrees of the polynomials are minimal. To see this, let $f_1, \ldots, f_s$ be polynomials in $\mathcal{R} = K[x_1, \ldots, x_n]$ and suppose that one of the polynomials, say f_j, has degree ≤ 0. If $f_j = 0$, then $\left[f_j\right]$ is not a non-zero-divisor for $\left\langle f_1, \ldots, f_{j-1} \right\rangle$. On the other hand, if $f_j = k \in K \setminus \{0\}$, then $k^{-1} f_j = 1$, so $\langle f_1, \ldots, f_s \rangle = \langle 1 \rangle$ is not a proper ideal. Therefore, if $f_1, \ldots, f_s$ forms a regular sequence in $\mathcal{R}$, we must have $\deg f_i \geq 1$ for all $1 \leq i \leq s$.

The following result shows that a regular sequence remains regular if we extend its base polynomial ring by new variables $y_1, \ldots, y_m$.

Lemma 60. *Let K be a field and suppose that $f_1, \ldots, f_s$ is a regular sequence in $\mathcal{R} = K[x_1, \ldots, x_n]$. Then $f_1, \ldots, f_s$ is also a regular sequence in the extended polynomial ring $\mathcal{S} = K[x_1, \ldots, x_n, y_1, \ldots, y_m]$.*

Proof. We make two observations. First, we can interpret $\mathcal{S}$ as $\mathcal{S} = \mathcal{R}[y_1, \ldots, y_m]$. Second, if $\langle g_1, \ldots, g_t \rangle$ is an ideal in $\mathcal{R}$, its generators are elements of $\mathcal{S}$ as well. We can interpret the ideal in $\mathcal{S}$ having generators $g_1, \ldots, g_t \in \mathcal{S}$ as $\langle g_1, \ldots, g_t \rangle [y_1, \ldots, y_m]$.

We want to show that for all $1 \leq i \leq s$, $[f_i]$ is a non-zero-divisor for the quotient space $\mathcal{R}[y_1, \ldots, y_m] / \langle f_1, \ldots, f_{i-1} \rangle [y_1, \ldots, y_m]$. Fix $1 \leq i \leq s$ and let $\mathcal{I} = \langle f_1, \ldots, f_{i-1} \rangle \subseteq \mathcal{R}$. By repeated application of [DF04, Chapter 9, Proposition 2], there exists a ring isomorphism

$$\mathcal{R}[y_1, \ldots, y_m] / \mathcal{I}[y_1, \ldots, y_m] \to (\mathcal{R}/\mathcal{I})[y_1, \ldots, y_{i-1}]$$

defined by $\left[\sum_\alpha c_\alpha \mathbf{y}^\alpha\right] \mapsto \sum_\alpha [c_\alpha] \mathbf{y}^\alpha$. (The c_α are polynomials in $\mathcal{R}$.) Now, suppose that $[u f_i] = 0$ in $\mathcal{R}[y_1, \ldots, y_m] / \mathcal{I}[y_1, \ldots, y_m]$ for some $u = \sum_\alpha c_\alpha \mathbf{y}^\alpha \in \mathcal{R}[y_1, \ldots, y_m]$. (Again, the c_α are polynomials in $\mathcal{R}$.) Then $\left[\sum_\alpha (c_\alpha f_i) \mathbf{y}^\alpha\right] = 0$ in $\mathcal{R}[y_1, \ldots, y_m] / \mathcal{I}[y_1, \ldots, y_m]$ and by the isomorphism above, we have that $\sum_\alpha [c_\alpha f_i] \mathbf{y}^\alpha = 0$. It follows that all $[c_\alpha f_i]$ are zero in $\mathcal{R}/\mathcal{I}$. But $[f_i]$ is a non-zero-divisor for $\mathcal{R}/\mathcal{I}$, so the $[c_\alpha]$ are all zero in $\mathcal{R}/\mathcal{I}$. By the isomorphism, it follows that $[u] = \left[\sum_\alpha c_\alpha \mathbf{y}^\alpha\right] = 0$ in $\mathcal{R}[y_1, \ldots, y_m] / \mathcal{I}[y_1, \ldots, y_m]$, and we conclude that $[f_i]$ is a non-zero-divisor for $\mathcal{R}[y_1, \ldots, y_m] / \mathcal{I}[y_1, \ldots, y_m]$. □

In this thesis, we're mainly interested in regular sequences because we can easily characterize their *syzygies*, as defined next.

Definition 61. Let K be a field and let $f_1, \ldots, f_s$ and $u_1, \ldots, u_s$ be polynomials in $K[x_1, \ldots, x_n]$. We call $(u_1, \ldots, u_s)$ a **syzygy** of $(f_1, \ldots, f_s)$ if $\sum_{k=1}^s u_k f_k = 0$.

As an example, let $f_1 = x^2, f_2 = xy \in K[x, y]$. Then $y f_1 - x f_2 = 0$, so $(y, -x)$ is a syzygy of (f_1, f_2).

More generally, let $f_1, \ldots, f_s$ be arbitrary polynomials in $K[x_1, \ldots, x_n]$ and let $\mathbf{e}_1, \ldots, \mathbf{e}_s \in \mathcal{R}^s$ denote the standard basis vectors of $\mathcal{R}^s$. (That is, $\mathbf{e}_i$ equals 1 in its i-th component and zero elsewhere.) For all $1 \leq i < j \leq s$ we have $f_j f_i - f_i f_j = 0$, which means that $f_j \mathbf{e}_i - f_i \mathbf{e}_j$ is a syzygy of $(f_1, \ldots, f_s)$. A syzygy of this form is called a **trivial syzygy**.

The next lemma shows that all syzygies of a regular sequence are generated by trivial ones.

Lemma 62. *Let K be a field and let $f_1, \ldots, f_s$ in $K[x_1, \ldots, x_n]$ be a regular sequence. If $(u_1, \ldots, u_s)$ is a syzygy of $(f_1, \ldots, f_s)$, then there exist polynomials $v_{ij} \in \mathcal{R}$ with $1 \leq i < j \leq s$ such that $(u_1, \ldots, u_s) = \sum_{1 \leq i < j \leq s} v_{ij} \left(f_j \mathbf{e}_i - f_i \mathbf{e}_j\right)$.*

Proof. Follows from [Eis95, Corollary 17.5]. □

The following result is sometimes easier to work with, since it describes the polynomials $u_1, \ldots, u_s$ of the syzygy individually.

Corollary 63. *Let K be a field and let $f_1, \ldots, f_s$ in $K[x_1, \ldots, x_n]$ be a regular sequence. If $(u_1, \ldots, u_s)$ is a syzygy of $(f_1, \ldots, f_s)$, then there exist polynomials $w_{ij} \in \mathcal{R}$ with $1 \leq i, j \leq s$ such that $u_i = \sum_{j=1}^s w_{ij} f_j$ for all $1 \leq i \leq s$. These polynomials satisfy $w_{ij} = -w_{ji}$ and $w_{ii} = 0$ for all $1 \leq i, j \leq s$.*

Proof. Define

$$w_{ij} = \begin{cases} v_{ij} & \text{if } i < j \leq s, \\ 0 & \text{if } i = j, \\ -v_{ji} & \text{if } 1 \leq j < i. \end{cases}$$

Considering the k-th components of both sides of the equality $(u_1, \ldots, u_s) = \sum_{1 \leq i < j \leq s} v_{ij} \left(f_j \mathbf{e}_i - f_i \mathbf{e}_j \right)$ from Lemma 62, we find that, for all $1 \leq k \leq s$, we have $u_k = \sum_{k<j\leq s} v_{kj} f_j - \sum_{1 \leq i < k} v_{ik} f_i = \sum_{j=1}^{s} w_{kj} f_j$. □

2.11 Security Level

Designers of cryptographic primitives usually claim that the primitive has a security level of e.g. 128 or 256 bits against certain attacks. For the attacks considered in this thesis, the following notion of security level suffices.

Definition 64. Suppose that an attack against a cryptographic primitive has complexity $\mathcal{C}$ and success probability P. We define the **security level** λ, measured in bits, of the cryptographic primitive against this attack as $\lambda = \log_2 \left(\frac{\mathcal{C}}{\mathrm{P}} \right)$.

The unit of complexity depends on the context. For a hash function H (Section 2.12), $\mathcal{C}$ usually denotes the number of evaluations of H. For the algebraic attack discussed in this thesis, $\mathcal{C}$ will denote the number of field operations in $\mathbb{F}_p$ (addition and multiplication).

Example 65. In later chapters, we estimate the complexity of an algebraic attack against ATRAPOS-SPONGE to be $\mathcal{C} = 2^{2\ell R}$ field operations in $\mathbb{F}_p$, where R is the number of rounds of the ATRAPOS permutation and ℓ is a parameter that depends on p. The success probability of such an attack is $\mathrm{P} = 1$. Assuming that this is the best possible attack, we say that ATRAPOS-SPONGE has a security level of $\lambda = \log_2 \left(\frac{2^{2\ell R}}{1} \right) = 2\ell R$ bits (with respect to field operations in $\mathbb{F}_p$) against algebraic attacks. We will use this in Section 5.5 to derive the minimal number of rounds needed to obtain at least 128 bits of security against algebraic attacks. ◊

2.12 Hash Functions

Recall that a hash function can be defined as a function $H \colon \mathbb{F}_p^* \to \mathbb{F}_p^\ell$ which takes an input $x \in \mathbb{F}_p^*$ of arbitrary length and maps it to an output $H(x)$ of length ℓ. A common value for p is $p = 2$, in which case H operates on bits.

Three desirable security properties of a cryptographic hash function include:

1. Preimage resistance: given a random output $y \in H\left(\mathbb{F}_p^*\right)$, it should be computationally infeasible to find an $x \in \mathbb{F}_p^*$ such that $H(x) = y$.

2. Second-preimage resistance: given a random $x \in \mathbb{F}_p^*$, it should be computationally infeasible to find $x' \in \mathbb{F}_p^* \setminus \{x\}$ such that $H(x') = H(x)$.

3. Collision resistance: it should be computationally infeasible to find two distinct $x, x' \in \mathbb{F}_p^*$ with $H(x') = H(x)$.

Note that the security properties are listed in increasing order of strength; collision resistance implies second-preimage resistance, which, in turn, implies preimage resistance (assuming that the definition of "infeasible" is kept fixed).

We can derive upper bounds for the security level of a hash function H with respect to preimage resistance, second-preimage resistance, and collision resistance by considering random oracles. For a random oracle, the probability of finding a preimage x corresponding to some given $y \in \mathbb{F}_p^\ell$ after $\mathcal{C} \leq p^\ell$ queries is $\mathrm{P} \approx \mathcal{C}/p^\ell$. For any $1 \leq \mathcal{C} \leq p^\ell$, we find that $\mathcal{C}/\mathrm{P} \approx p^\ell$. Thus, $\lambda \approx \log_2 p^\ell = \ell \log_2 p$ gives an upper bound on the security level in bits of H against preimage attacks. The same upper bound holds for second-preimage attacks.

Due to the birthday problem, the probability of finding a collision after $\mathcal{C}$ queries gets close to 1 if $\mathcal{C}$ approaches $\sqrt{p^\ell} = p^{\ell/2}$. It follows that the security level in bits of H against collision attacks is bounded from above by $\lambda \approx \log_2 p^{\ell/2} = \ell/2 \cdot \log_2 p$.

For a secure hash function, the actual security level should be as close to these derived upper bounds as possible.

Remark 66. The bounds derived here assume that the adversary uses classical computers. In the presence of quantum computers, the security level against preimage and second preimage attacks is bounded from above by $\approx \ell/2 \cdot \log_2 p$, while the security level against collision attacks is bounded from above by $\approx \ell/3 \cdot \log_2 p$. See e.g. [KL20, §14.1] for a discussion.

2.13 The Sponge Construction

Extendable-output functions (XOFs) extend the notion of hash functions by allowing for arbitrary length outputs. That is, a XOF is a function $H : \mathbb{F}_p^* \times \mathbb{N} \to \mathbb{F}_p^*$, which maps pair (M, ℓ), consisting of a message M of arbitrary length and a requested output length ℓ, to an output of length ℓ. The usual security notions of preimage resistance, second-preimage resistance, and collision resistance for hash functions carry over to XOFs. The corresponding security levels, however, are different.

A practical method to build extendable-output functions is using the sponge construction (Figure 2.2) [BDPV11a]. The sponge construction operates on a state in $\mathbb{F}_p^b$. The state is partitioned into a c-digit inner part and an r-digit outer part [BDPV11a], so that $r + c = b$. We refer to r and c as the rate and capacity, respectively. The state is initially zero.

[1] Thanks to Bart Mennink for providing the TikZ code for this figure.

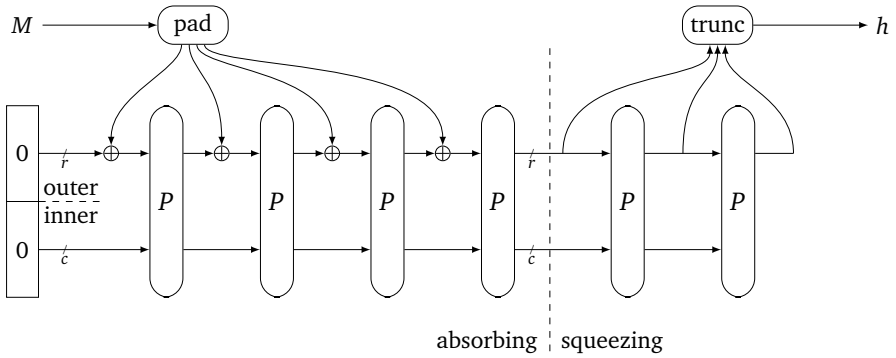

Figure 2.2: The sponge construction.[1]

After injectively padding and cutting the input message M into r-digit blocks, the sponge construction proceeds in two phases. During the absorbing phase, the first block of r-digit from the padded message is added (entrywise) to the outer part of the state. The state is then mapped to a new state by a permutation $P\colon \mathbb{F}_p^b \to \mathbb{F}_p^b$. This process repeats until all r-digit blocks are consumed. The squeezing phase follows the absorbing phase. During this phase, r-digit blocks from the outer part of the state are returned by calling P as often as required. Finally, the r-digit output blocks are combined and truncated to the first ℓ digits.

As shown in [BDPV11a], sponge constructions are computationally indistinguishable from random oracles, assuming *generic attacks*. We call an attack generic if only exploits general properties of sponge constructions, but not of the specific permutation used within the sponge construction. In this setting, the sponge construction, where the output is truncated to ℓ digits, has a classical security level of $\approx \min(c/2, n) \cdot \log_2 p$ bits against (second) preimage attacks and a classical security level of $\approx \min(c/2, n/2) \cdot \log_2 p$ bits against collision attacks.

Of course, practical implementations of sponge constructions are trivially distinguishable from random oracles, since the permutation P used in the sponge construction has a compact description (which we assume is known publically).

However, the mere fact that the permutation P has a trivial distinguisher, due to its compact description, provides no practical benefits in an attack. In [BDPV11a], examples of structural (non-trivial) distinguishers for P are listed, whose existence may be of practical use in attacks. Well-known structural distinguishers are differential and linear cryptanalysis. In this thesis, however, we will be interested in a structural distinguisher called the constrained-input constrained-output (CICO) problem. This is not a single problem, but rather a family of problems.

Definition 67 (CICO problem)**.** Let $P\colon \mathbb{F}_p^b \to \mathbb{F}_p^b$ be a permutation. Every CICO problem has the following form: given a set of possible inputs, $\mathcal{X} \subseteq \mathbb{F}_p^b$, and a

set of possible outputs, $\mathcal{Y} \subseteq \mathbb{F}_p^b$, find a pair $(x, y) \in \mathcal{X} \times \mathcal{Y}$ such that $y = P(x)$.

The general notion of the CICO problem subsumes the specific notions of preimage resistance, second-preimage resistance, collision resistance, etc. For example, consider the preimage resistance property which states that, given a random output $\overline{y}$, it should be computationally infeasible to find an input $\overline{x}$ that maps to $\overline{y}$. For an r-bit input and an r-bit output, this corresponds to the following CICO problem (where we ignore padding for the sake of the argument): given $\mathcal{X} = \mathbb{F}_p^b$ and $\mathcal{Y} = \{\overline{y}\}$, find a pair $(x, y) \in \mathcal{X} \times \mathcal{Y}$ with $y = P(x)$. Note that we can account for padding by constricting $\mathcal{X}$.

We leave it as an exercise to the reader to formulate other security properties in terms of CICO problems.

Chapter 3

ATRAPOS-SPONGE

In this section, we define ATRAPOS-SPONGE, the ATRAPOS permutations, and the CICO problem for ATRAPOS that we will analyze in this thesis.

3.1 Specification

ATRAPOS-SPONGE is an extendable-output function (XOF) based on the sponge construction [DMMØ25]. It is designed to be an efficient alternative for SHA3 in KYBER and DILITHIUM on platforms where hardware acceleration for multiplication in $\mathbb{F}_p$ is available, where either $p = 3329$ (for KYBER) or $p = 8380417$ (for DILITHIUM). ATRAPOS, the permutation used in ATRAPOS-SPONGE, operates on states of digits in $\mathbb{F}_p$. The states are represented by two-dimensional arrays consisting of 3 rows and ℓ columns:

$$\mathbf{a} = \begin{pmatrix} a_{0,2} & a_{1,2} & \cdots & a_{\ell-1,2} \\ a_{0,1} & a_{1,1} & \cdots & a_{\ell-1,1} \\ a_{0,0} & a_{1,0} & \cdots & a_{\ell-1,0} \end{pmatrix},$$

where $a_{x,y} \in \mathbb{F}_p$ for all $0 \leq x \leq \ell - 1$ and $0 \leq y \leq 2$. Equivalently, states can be interpreted as one-dimensional vectors in $\mathbb{F}_p^{3\ell}$. Given a one-dimensional representation $\mathbf{a} = (a_0, \ldots, a_{3\ell-1}) \in \mathbb{F}_p^{3\ell}$, its two-dimensional representation can be computed by converting the one-dimensional index i to the two-dimensional index $(i \bmod \ell, i \bmod 3)$. Since the ATRAPOS permutations are defined using two-dimensional representations, we will represent states in $\mathbb{F}_p^{3\ell}$ by their two-dimensional representation from now on.

The ATRAPOS-SPONGE specification is still in development and a number of parameters have not been fixed yet. For example, the current ATRAPOS-SPONGE specification allows either the bottom row (the $a_{x,0}$) or the bottom two rows (the $a_{x,0}$ and $a_{x,1}$) to be used for the outer part in the sponge construction. In this thesis, we will confine to the case where only the bottom row is used for the outer part, since it greatly simplifies the analysis. In this case, the rate and

capacity of ATRAPOS-SPONGE are $r = \ell$ and $c = 2\ell$, respectively. The specification prescribes $\ell = 17$ for KYBER and $\ell = 7$ for DILITHIUM in order to achieve 128 bits (with respect to evaluations of the ATRAPOS permutation) of generic collision resistance in the presence of adversaries having access to a quantum computer.

3.1.1 The ATRAPOS Permutations

ATRAPOS is a family of permutations. Every ATRAPOS permutation is a composition of R round functions. The j-th round function is itself a composition $\gamma \circ \iota_j \circ \rho \circ \theta$, where $\gamma, \iota_j, \rho, \theta : \mathbb{F}_p^{3\ell} \to \mathbb{F}_p^{3\ell}$ are the following (polynomial) permutations:

$$\begin{aligned}
&\theta : a_{x,y} \leftarrow a_{x,y} + a_{x+1,y+1} + a_{x+4,y+4} + a_{x+5,y+5} && \forall x, y \\
&\rho : a_{x,y} \leftarrow a_{x+r_y,y} && \forall x, y \\
&\iota_j : a_{0,0} \leftarrow a_{0,0} + c_j && \\
&\gamma : a_{x,0} \leftarrow a_{x,0} + a_{x,1} a_{x,2} && \forall x
\end{aligned}$$

Note that x ranges over $\{0, \ldots, \ell - 1\}$ and y ranges over $\{0, 1, 2\}$. (The x-components of indices above are implicitly taken modulo ℓ and the y-components are implicitly taken modulo 3.) Here, $c_j \in \mathbb{F}_p$ is a round-dependent constant. The current ATRAPOS specification has not fixed the c_j yet, but our analysis of ATRAPOS will be independent of their specific values.

The $0 < r_y < \ell$ are row-dependent column shifts. The shift r_0 is set to 0. Specific values for r_1 and r_2 have not yet been chosen, but it has been determined that they have to satisfy $r_2 = r_1 + 1$ or $r_2 = r_1 + 4$. These choices of r_1 and r_2 ensure that the ideal degree related to ATRAPOS are maximal, as we will see in Chapter 4. Again, the specific values of r_1 and r_2 do not change our analysis.

We define $\text{ATRAPOS}[1] = \gamma \circ \iota_1 \circ \rho \circ \theta$ to be the first round of the ATRAPOS permutation. Generally, having defined $\text{ATRAPOS}[1], \ldots, \text{ATRAPOS}[R-1]$ we define $\text{ATRAPOS}[R] = \gamma \circ \iota_R \circ \rho \circ \theta \circ \text{ATRAPOS}[R-1]$ to be the first R rounds of ATRAPOS.

Remark 68. The ATRAPOS-SPONGE specification only allows $\ell = 17$ (for KYBER) and $\ell = 7$ (for DILITHIUM). However, the ATRAPOS permutations defined in this section are permutations for all odd $\ell \geq 3$. We will therefore analyze ATRAPOS-SPONGE for arbitrary odd $\ell \geq 3$.

3.2 CICO problem

As discussed in Section 2.13, permutations used in sponge constructions need to be secure with respect to structural distinguishers. In this section, we motivate

and define a CICO problem for ATRAPOS by studying a preimage attack. For simplicity, we only consider a target digest consisting of a single r-digit block.

Consider ATRAPOS-SPONGE with an output length of r digits, as depicted in Figure 3.1. Let $M' = \mathrm{pad}(M)$ denote the padded version of M. We can write this as the concatenation $M' = M_1 \,\|\, M_2 \,\|\, \cdots \,\|\, M_n$, where $M_1, \ldots, M_n$ are r-digit blocks. We study the following problem.

Definition 69 (Relaxed preimage attack)**.** Given an r-digit target digest h, find r-digit blocks $M_1, \ldots, M_n$ such that the *padded* message $M' = M_1 \,\|\, M_2 \,\|\, \cdots \,\|\, M_n$ results in the (truncated) digest h.

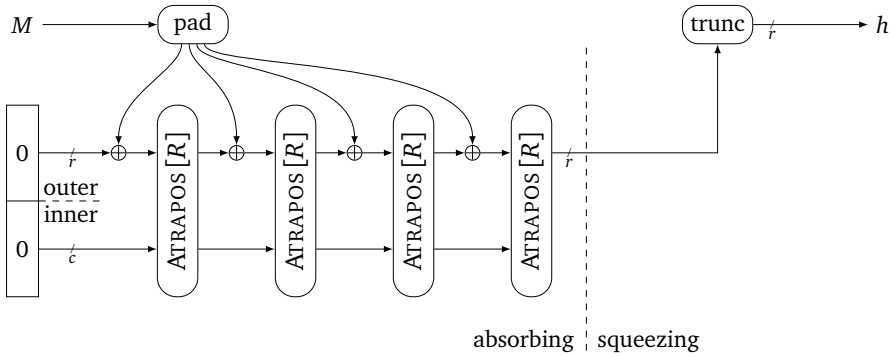

Figure 3.1: The sponge construction for ATRAPOS, where the output consists of a single r-digit block.

Note that the adversary may assume that M' is already a padded message, i.e. there exists some message M such that $M' = \mathrm{pad}(M)$. This is to the advantage of the adversary, since M' may not actually be in the image of the padding function. By showing that the relaxed preimage attack is computationally infeasible, we therefore also show that the regular preimage attack, where we need to find a message M before the padding step, is computationally infeasible.

To formalize this, let $M' = M_1 \| M_2 \| \cdots \| M_n$ be any padded message. Let $\mathbf{a}$ denote the state just before the last call to ATRAPOS $[R]$ in Figure 3.1 (after the entrywise addition of M_n to the outer part of the state) and let $\mathbf{a}' = \text{ATRAPOS}\,[R](\mathbf{a})$ denote the state after the last call to ATRAPOS $[R]$. We suggestively write

$$\mathbf{a} = \begin{pmatrix} a_{0,2} & a_{1,2} & \cdots & a_{\ell-1,2} \\ a_{0,1} & a_{1,1} & \cdots & a_{\ell-1,1} \\ x_1 & x_2 & \cdots & x_\ell \end{pmatrix}$$

and

$$\mathbf{a}' = \begin{pmatrix} a'_{0,2} & a'_{1,2} & \cdots & a'_{\ell-1,2} \\ a'_{0,1} & a'_{1,1} & \cdots & a'_{\ell-1,1} \\ y_1 & y_2 & \cdots & y_\ell \end{pmatrix},$$

to indicate that an adversary can easily set the bottom row of $\mathbf{a}$ as desired, and to indicate that the bottom row of $\mathbf{a}'$ corresponds to the output digest h.

Solving the relaxed preimage attack now corresponds to finding $\mathbf{a}_1, \mathbf{a}_2, \mathbf{x} \in \mathbb{F}_p^\ell$ such that the state $\mathbf{a} = (\mathbf{a}_1, \mathbf{a}_2, \mathbf{x})^\mathsf{T}$ is mapped to the state $\mathbf{a}' = \left(\mathbf{a}'_1, \mathbf{a}'_2, \mathbf{y}\right)^\mathsf{T}$, where $\mathbf{a}'_1, \mathbf{a}'_2, \mathbf{y} \in \mathbb{F}_p^\ell$ and $\mathbf{y}$ is the target digest. This motivates the following CICO problem.

Definition 70 (CICO problem for ATRAPOS-SPONGE)**.** Given $\mathbf{a}_1, \mathbf{a}_2, \mathbf{y} \in \mathbb{F}_p^\ell$, find $\mathbf{x} \in \mathbb{F}_p^\ell$ such that there exist $\mathbf{a}'_1, \mathbf{a}'_2 \in \mathbb{F}_p^\ell$ with

$$\text{ATRAPOS}\,[R](\mathbf{a}_2, \mathbf{a}_1, \mathbf{x}) = \left(\mathbf{a}'_2, \mathbf{a}'_1, \mathbf{y}\right).$$

Note that in this CICO problem, $\mathbf{a}_1$ and $\mathbf{a}_2$ are fixed. We will see that the CICO problem in Definition 70 is computationally infeasible (for R sufficiently large) for all choices of $\mathbf{a}_1$ and $\mathbf{a}_2$. Consequently, the relaxed preimage attack in Definition 69 is also computationally infeasible (for R sufficiently large).

3.2.1 Polynomial Modeling

In the CICO problem in Definition 70, $\mathbf{a}_1, \mathbf{a}_2 \in \mathbb{F}_p^\ell$ are fixed, while $\mathbf{x} \in \mathbb{F}_p^\ell$ is a variable. Thus, we may consider $x_1, \ldots, x_\ell$ to be variables in a polynomial ring $\mathcal{R} = \mathbb{F}_p[x_1, \ldots, x_\ell]$. We can then take $a_{x,1}$ and $a_{x,2}$ to be constants in this ring. From this point of view, ATRAPOS is a mapping

$$\begin{pmatrix} a_{0,2} & a_{1,2} & \cdots & a_{\ell-1,2} \\ a_{0,1} & a_{1,1} & \cdots & a_{\ell-1,1} \\ x_1 & x_2 & \cdots & x_\ell \end{pmatrix} \mapsto \begin{pmatrix} a'_{0,2} & a'_{1,2} & \cdots & a'_{\ell-1,2} \\ a'_{0,1} & a'_{1,1} & \cdots & a'_{\ell-1,1} \\ g_1 & g_2 & \cdots & g_\ell \end{pmatrix}.$$

Since the round function of ATRAPOS is a composition of polynomials, ATRAPOS itself is a polynomial mapping. Therefore, g_i, $a'_{x,1}$, and $a'_{x,2}$ are polynomials in $\mathcal{R}$. The CICO problem in Definition 70 is equivalent to finding a solution for the polynomial system

$$\begin{aligned} g_1(x_1, \ldots, x_\ell) &= y_1 \\ g_2(x_1, \ldots, x_\ell) &= y_2 \\ &\vdots \\ g_\ell(x_1, \ldots, x_\ell) &= y_\ell \end{aligned} \tag{3.1}$$

Recall from Section 2.8 that state-of-the-art algorithms for solving polynomial systems consist of three steps (DRL Gröbner basis computation using F4/F5, lexicographic Gröbner basis computation using FGLM, univariate polynomial solving), where the running time is dominated by the first two steps (Remark 47). The experimental results in Chapter 6 show that, for small systems, the first step (DRL Gröbner basis computation) is negligible compared to the second step (lexicographic Gröbner basis computation using FGLM). In this

thesis, we will formally show that the second step (lexicographic Gröbner basis computation using FGLM) is computationally infeasible. Thus, even if an adversary is able to perform the other two steps efficiently, the overall three-step algorithm is still infeasible due to the second step.

In line with the discussion in Section 2.8, we estimate the complexity of the FGLM step by $\mathcal{C}_{\text{FGLM}} = d_{\mathcal{I}}^{\omega}$, where we conservatively set $\omega = 2$ to account for attacks that exploit sparsity of the polynomial system in Equation (3.1). In Chapter 4 and Chapter 5 we will show that, for all odd $\ell > 3$ and for all $R \geq 1$, the ideal corresponding to Equation (3.1) has ideal degree $2^{\ell R}$. It then follows that the complexity of the FGLM step for an R-round ATRAPOS permutation is given by $\mathcal{C}_{\text{FGLM}} = 2^{2\ell R}$.

Chapter 4

Single-Round Analysis

In this chapter, we determine the complexity of solving the CICO problem in Definition 70 for a single round of ATRAPOS. For simplicity, we restrict to the case where $\ell > 3$ is an odd number not divisible by 3, which is the relevant case for KYBER ($\ell = 17$) and DILITHIUM ($\ell = 7$).

In Section 4.1, we will work out a polynomial system $\mathcal{F}_{\text{inh}}$, which corresponds to a single round of ATRAPOS whenever $\ell > 3$ is not divisible by 3. From $\mathcal{F}_{\text{inh}}$ we obtain a system $\mathcal{F}_{\text{hom}}$ consisting of homogeneous polynomials, whose ideal $\mathcal{I}_{\text{hom}} = \langle \mathcal{F}_{\text{hom}} \rangle$ has the same ideal degree as $\mathcal{I}_{\text{inh}} = \langle \mathcal{F}_{\text{inh}} \rangle$, but is easier to work with. In Section 4.2, we consider the Hilbert series of $\mathbb{F}_p[x_1, \ldots, x_\ell]/\mathcal{I}_{\text{hom}}$ and make a claim (the "Direct Sum Claim"), which is equivalent to this Hilbert series being equal to $\sum_{m=0}^{\ell} \binom{\ell}{m} \cdot t^m$. In Section 4.3 and Section 4.4 we give a proof sketch and a formal proof, respectively, for this claim. As a by-product, we learn that $\mathcal{I}_{\text{hom}}$ is generated by a regular sequence of polynomials.

Throughout this chapter, $p \geq 3$ will denote a prime, $\ell \geq 3$ an odd number denoting the width of the two-dimensional ATRAPOS state, and $\mathcal{R}$ will denote the polynomial ring $\mathbb{F}_p[x_1, \ldots, x_\ell]$.

4.1 Polynomial System

Let

$$\mathbf{a} = \begin{pmatrix} a_{0,2} & a_{1,2} & \cdots & a_{\ell-1,2} \\ a_{0,1} & a_{1,1} & \cdots & a_{\ell-1,1} \\ a_{0,0} & a_{1,0} & \cdots & a_{\ell-1,0} \end{pmatrix} = \begin{pmatrix} a_{0,2} & a_{1,2} & \cdots & a_{\ell-1,2} \\ a_{0,1} & a_{1,1} & \cdots & a_{\ell-1,1} \\ x_1 & x_2 & \cdots & x_\ell \end{pmatrix}$$

be the input state and let

$$\mathbf{a}' = \begin{pmatrix} a'_{0,2} & a'_{1,2} & \cdots & a'_{\ell-1,2} \\ a'_{0,1} & a'_{1,1} & \cdots & a'_{\ell-1,1} \\ a'_{0,0} & a'_{1,0} & \cdots & a'_{\ell-1,0} \end{pmatrix} = \begin{pmatrix} a'_{0,2} & a'_{1,2} & \cdots & a'_{\ell-1,2} \\ a'_{0,1} & a'_{1,1} & \cdots & a'_{\ell-1,1} \\ g_1 & g_2 & \cdots & g_\ell \end{pmatrix}$$

be the output state after a single round of ATRAPOS, say the j-th round. As noted in Section 3.2, every element of $\mathbf{a}'$ can be viewed as a polynomial in

the polynomial ring $\mathcal{R} = \mathbb{F}_p[x_1, \ldots, x_\ell]$. The complexity of solving the CICO problem in Definition 70 is given by $\dim_{\mathbb{F}}(\mathcal{R}/\mathcal{I})$, where $\mathcal{I}$ is the ideal generated by $g_1, \ldots, g_\ell$. We explicitly compute expressions for $g_1, \ldots, g_\ell$ and use these to determine $\dim_{\mathbb{F}}(\mathcal{R}/\mathcal{I})$.

We introduce the variables $\mathbf{a}_1 = \theta(\mathbf{a})$, $\mathbf{a}_2 = \rho(\mathbf{a}_1)$, $\mathbf{a}_3 = \iota(\mathbf{a}_2)$ for the intermediate states. (Hence, $\mathbf{a}' = \gamma(\mathbf{a}_3)$.) After applying θ to $\mathbf{a}$, the value of the state at (x, y) is given by

$$(\mathbf{a}_1)_{x,y} = a_{x,y} + a_{x+1,y+1} + a_{x+4,y+4} + a_{x+5,y+5}$$

for all $0 \le x \le \ell - 1$ and $0 \le y \le 2$. (Recall that the x-components of the indices are taken modulo ℓ, while the y-components are taken modulo 3.) Applying ρ to $\mathbf{a}_1$, we have

$$\begin{aligned}(\mathbf{a}_2)_{x,y} &= (\rho(\mathbf{a}_1))_{x+r_y,y} \\ &= a_{x+r_y,y} + a_{x+1+r_y,y+1} + a_{x+4+r_y,y+4} + a_{x+5+r_y,y+5}\end{aligned}$$

for all $0 \le x \le \ell - 1$ and $0 \le y \le 2$. After the ι step, we have

$$\begin{pmatrix}(\mathbf{a}_3)_{x,2}\\(\mathbf{a}_3)_{x,1}\\(\mathbf{a}_3)_{x,0}\end{pmatrix} = \begin{pmatrix}(\iota(\mathbf{a}_2))_{x,2}\\(\iota(\mathbf{a}_2))_{x,1}\\(\iota(\mathbf{a}_2))_{x,0}\end{pmatrix} = \begin{pmatrix}a_{x+r_2,2} + a_{x+1+r_2,0} + a_{x+4+r_2,0} + a_{x+5+r_2,1}\\a_{x+r_1,1} + a_{x+1+r_1,2} + a_{x+4+r_1,2} + a_{x+5+r_1,0}\\a_{x+r_0,0} + a_{x+1+r_0,1} + a_{x+4+r_0,1} + a_{x+5+r_0,2} + c_j\end{pmatrix}$$

for all $0 \le x \le \ell - 1$. Finally, after applying γ to $\mathbf{a}_3$, we find that for all $0 \le x \le \ell - 1$,

$$\begin{aligned}a'_{x,0} &= (\mathbf{a}_3)_{x,0} + (\mathbf{a}_3)_{x,1}(\mathbf{a}_3)_{x,2} \\ &= a_{x+5+r_1,0}\left(a_{x+1+r_2,0} + a_{x+4+r_2,0}\right) + \epsilon_{x+1}(\mathbf{x}),\end{aligned}$$

where $\epsilon_1(\mathbf{x}), \ldots, \epsilon_\ell(\mathbf{x})$ are polynomials in $\mathcal{R}$ of degree < 2.

Write $\overline{n}$ for the unique integer $m \in \{1, \ldots, \ell\}$ such that $n - m \in \ell\mathbb{Z}$. Then, for all $i \in \mathbb{Z}$ we have $a_{i-1,0} = x_{\overline{i}}$. It follows that

$$g_i = a'_{i-1,0} = x_{\overline{i+5+r_1}}\left(x_{\overline{i+1+r_2}} + x_{\overline{i+4+r_2}}\right) + \epsilon_i(\mathbf{x})$$

for all $1 \le i \le \ell$.

In the ATRAPOS specification (Section 3.1), we saw that either $r_2 = r_1 + 1$ or $r_2 = r_1 + 4$. If $r_2 = r_1 + 1$, then

$$g_i = x^2_{\overline{i+5+r_1}} + x_{\overline{i+5+r_1}} \cdot x_{\overline{i+2+r_1}} + \epsilon_i(\mathbf{x})$$

and if $r_2 = r_1 + 4$, we have

$$g_i = x^2_{\overline{i+5+r_1}} + x_{\overline{i+5+r_1}} \cdot x_{\overline{i+8+r_1}} + \epsilon_i(\mathbf{x}).$$

Without loss of generality, we can relabel the x_i using $i \mapsto \overline{i-5}$ to obtain equations of the form

$$g_i = x^2_{\overline{i+r_1}} + x_{\overline{i+r_1}} \cdot x_{\overline{i-3+r_1}} + \epsilon_i(\mathbf{x})$$

and

$$g_i = x^2_{\overline{i+r_1}} + x_{\overline{i+r_1}} \cdot x_{\overline{i+3+r_1}} + \epsilon_i(\mathbf{x}),$$

respectively.

As long as $\gcd(\ell, 3) = 1$, we can relabel $x_1, \ldots, x_\ell$ and $g_1, \ldots, g_\ell$ in such a way that

$$\begin{aligned} g_1 &= x_1^2 + x_1 x_2 + \epsilon_1 \\ g_2 &= x_2^2 + x_2 x_3 + \epsilon_2 \\ &\vdots \\ g_\ell &= x_\ell^2 + x_1 x_\ell + \epsilon_\ell \end{aligned}$$

Since 3 is prime, the condition $\gcd(\ell, 3) = 1$ is equivalent to ℓ not being divisible by 3. This condition is certainly satisfied when $\ell = 17$ (for KYBER) or $\ell = 7$ (for DILITHIUM).

Remark 71. To delimit the scope of this thesis, we will not analyze the ideals corresponding to the ATRAPOS permutations when ℓ is divisible by 3.

If $\ell > 3$ is odd and not divisible by 3, solving the CICO problem in Definition 70 for a single round of ATRAPOS amounts to solving the polynomial system

$$\begin{aligned} g_1 &= y_1 \\ g_2 &= y_2 \\ &\vdots \\ g_\ell &= y_\ell \end{aligned}$$

Since $y_1, \ldots, y_\ell$ are constants in $\mathbb{F}[x_1, \ldots, x_n]$, we may absorb them in the ϵ_i. Thus, without loss of generality, solving the CICO problem for a single round of ATRAPOS (where, again, $\ell > 3$ is not divisible by 3) is equivalent to solving the polynomial system

$$\begin{aligned} g_1 &= f_1 + \epsilon_1 = 0 \\ g_2 &= f_2 + \epsilon_2 = 0 \\ &\vdots \\ g_\ell &= f_\ell + \epsilon_\ell = 0 \end{aligned} \tag{4.1}$$

where

$$f_i \coloneqq x_i^2 + x_i x_{\overline{i+1}} \in \mathbb{F}[x_1, \ldots, x_n]$$

contains the terms of g_i of degree 2 and $\epsilon_i \in \mathbb{F}[x_1, \ldots, x_n]$ contains the lower degree terms.

We call $\mathcal{F}_{\text{inh}} = (g_1, \ldots, g_\ell) \in \mathcal{R}^\ell$ the **inhomogeneous system** and refer to $\mathcal{I}_{\text{inh}} = \langle \mathcal{F}_{\text{inh}} \rangle \subseteq \mathcal{R}$ as the ideal corresponding to the inhomogeneous system.

The notation ϵ_i for the lower degree terms in Equation (4.1) is suggestive: as we will see in Chapter 5, these terms do not contribute to the ideal degree.

It therefore makes sense to study a simplified version of the system in Equation (4.1) where the ϵ_i are removed:

$$\begin{aligned} f_1 &= x_1^2 + x_1 x_2 = 0 \\ f_2 &= x_2^2 + x_2 x_3 = 0 \\ &\vdots \\ f_\ell &= x_\ell^2 + x_\ell x_1 = 0 \end{aligned}$$

We will refer to $\mathcal{F}_{\text{hom}} = (f_1, \ldots, f_\ell) \in \mathcal{R}^\ell$ as the **homogeneous system** corresponding to $\mathcal{F}_{\text{inh}}$. The ideal $\mathcal{I}_{\text{hom}} = \langle \mathcal{F}_{\text{hom}} \rangle \subseteq \mathcal{R}$ is the ideal corresponding the homogeneous system. For the sake of brevity, in this section we will often drop the subscript and write $\mathcal{I}$ instead of $\mathcal{I}_{\text{hom}}$.

The Bézout bound shows that the ideal degree $d_{\mathcal{I}} = \dim_{\mathbb{F}}(\mathcal{R}/\mathcal{I})$ is at most $\prod_{k=1}^{\ell} \deg f_i = 2^\ell$. Since the hardness of the CICO problem relies on the ideal degree being large, the optimal case is when $d_{\mathcal{I}_{\text{inh}}} = 2^\ell$. In the remainder of this chapter, we will see that $d_{\mathcal{I}}$ is indeed equal to 2^ℓ.

Remark 72. The ideal $\mathcal{I} = \langle \mathcal{F}_{\text{hom}} \rangle$ is well-defined even if ℓ is divisible by 3. Although in these cases, $\mathcal{I}$ does not correspond to an ideal induced by ATRAPOS, it is still useful to study them. For example, the Hilbert series of $\mathcal{R}/\mathcal{I}$ for $\ell = 3$ and $\ell = 9$ (see Table 4.1 in the next section), will help us understand the Hilbert series for arbitrary ℓ.

We will therefore study $\mathcal{I}$ for all odd $\ell \geq 3$. At the end of Chapter 5, we will again restrict to the cases where ℓ is not a multiple of 3.

For even ℓ, the ideal $\mathcal{I}$ does not have ideal degree 2^ℓ anymore. We will not discuss this case, since ATRAPOS is not a permutation for even ℓ.

4.2 Direct Sum Claim

For small odd ℓ, the ideal degree $d_{\mathcal{I}}$ can be computed directly by evaluating the Hilbert series $\text{HS}_{R/\mathcal{I}}(t)$ at $t = 1$ (see Section 2.9). Table 4.1 illustrates this for small odd values of $\ell \geq 3$. The code can be found in Section A.1.

ℓ	Hilbert series $\text{HS}_{R/\mathcal{I}}(t)$	$d_{\mathcal{I}}$
3	$t^3 + 3t^2 + 3t + 1$	2^3
5	$t^5 + 5t^4 + 10t^3 + 10t^2 + 5t + 1$	2^5
7	$t^7 + 7t^6 + 21t^5 + 35t^4 + 35t^3 + 21t^2 + 7t + 1$	2^7
9	$t^9 + 9t^8 + 36t^7 + 84t^6 + 126t^5 + 126t^4 + 84t^3 + 36t^2 + 9t + 1$	2^9

Table 4.1: Hilbert series of $\mathcal{R}/\mathcal{I}$ and ideal degree $d_{\mathcal{I}}$ for small values of ℓ.

It may not be immediately obvious, but the coefficients of the Hilbert series listed in Table 4.1 are binomial coefficients. More precisely, it seems that

$\dim_{\mathbb{F}}(\mathcal{R}_m/\mathcal{I}_m) = \binom{\ell}{m}$, where $\mathcal{R}_m$ and $\mathcal{I}_m$ denote the m-th degree homogeneous subspaces of $\mathcal{R}$ and $\mathcal{I}$, respectively, defined in Definition 50. If this holds in general, then indeed

$$d_{\mathcal{I}} = \mathrm{HS}_{\mathcal{R}/\mathcal{I}}(1) = \sum_{m=0}^{\infty} \dim_{\mathbb{F}}(\mathcal{R}_m/\mathcal{I}_m) = \sum_{m=0}^{\ell} \binom{\ell}{m} = 2^{\ell}.$$

(The equality $\sum_{m=0}^{\ell} \binom{\ell}{m} = 2^{\ell}$ is a well-known equality and can be proven using induction or using a combinatorial argument: both sides of the equation count the number of subsets of $\{1, 2, \ldots, \ell\}$.) Thus, it suffices to show the equality

$$\dim_{\mathbb{F}}(\mathcal{R}_m/\mathcal{I}_m) = \binom{\ell}{m} \tag{4.2}$$

for all $m \in \mathbb{Z}_{\geq 0}$.

For all $m \in \mathbb{Z}_{\geq 0}$, the vector spaces $\mathcal{R}_m$ and $\mathcal{I}_m$ are finite-dimensional, so we may use Definition 2 to see that $\dim_{\mathbb{F}}(\mathcal{R}_m/\mathcal{I}_m) = \dim_{\mathbb{F}} \mathcal{R}_m - \dim_{\mathbb{F}} \mathcal{I}_m$. Equation (4.2) is therefore equivalent to $\dim_{\mathbb{F}} \mathcal{I}_m + \binom{\ell}{m} = \dim_{\mathbb{F}} \mathcal{R}_m$. In other words, we want to show that every $\mathcal{R}_m$ can be written as the direct sum $\mathcal{R}_m = \mathcal{I}_m \oplus V_m$ for some $\binom{\ell}{m}$-dimensional linear subspace $V_m \subseteq \mathcal{R}_m$.

The next example verifies Equation (4.2) for $\ell = 3$ and $0 \leq m \leq 3$ by exhibiting linear subspaces $V_m \subseteq \mathcal{R}_m$ with the stated property.

Example 73. Let $\ell = 3$ and $\mathcal{R} = \mathbb{F}_p[x, y, z]$. Then $\mathcal{F}_{\mathrm{hom}} = (f_1, f_2, f_3)$, where

$$f_1 = x^2 + xy$$
$$f_2 = y^2 + yz$$
$$f_3 = z^2 + zx$$

Recall from Lemma 53 that $\mathcal{I}_m$ can be written as

$$\mathcal{I}_m = \mathrm{span}_{\mathbb{F}}\{\mathbf{x}^{\alpha} f_i \mid 1 \leq i \leq n \text{ and } \deg \alpha = m - 2\}$$

for all $m \in \mathbb{Z}_{\geq 0}$. For $0 \leq m \leq 3$ we find the following:

$$\mathcal{I}_0 = \mathrm{span}_{\mathbb{F}}\{0\}$$
$$\mathcal{I}_1 = \mathrm{span}_{\mathbb{F}}\{0\}$$
$$\mathcal{I}_2 = \mathrm{span}_{\mathbb{F}}\{f_1, f_2, f_3\}$$
$$\mathcal{I}_3 = \mathrm{span}_{\mathbb{F}}\{xf_1, yf_1, zf_1, xf_2, yf_2, zf_2, xf_3, yf_3, zf_3\}$$

We verify that Equation (4.2) holds for each of these homogeneous subspaces by explicitly computing an $\binom{\ell}{m}$-dimensional linear subspace $V_m \subseteq \mathcal{R}_m$ such that $\mathcal{R}_m = \mathcal{I}_m \oplus V_m$.

- For $m = 0$ we have $\mathcal{I}_0 = \{0\}$ and $\mathcal{R}_0 = \mathbb{F}_p$, since the 0-degree polynomials in $\mathcal{R}$ are exactly the constant polynomials. The space $V_0 = \mathbb{F}_p = \mathrm{span}_{\mathbb{F}}\{1\} \subseteq \mathcal{R}_0$ is a linear space of dimension $\binom{3}{0} = 1$ such that $\mathcal{R}_0 = \mathcal{I}_0 \oplus V_0$.

- For $m = 1$, $\mathcal{I}_1$ is again the zero space and $\mathcal{R}_1 = \text{span}_{\mathbb{F}}\{x, y, z\}$. Thus, $V_1 = \text{span}_{\mathbb{F}}\{x, y, z\}$ satisfies $\mathcal{R}_m = \mathcal{I}_m \oplus V_m$ and $\dim_{\mathbb{F}} V_1 = 3 = \binom{3}{1}$.

- For $m = 2$, let $V_2 = \text{span}_{\mathbb{F}}\{xy, xz, yz\}$. It is clear that $\mathcal{I}_2 \cap V_2 = \{0\}$ and $\dim_{\mathbb{F}} V_2 = 3 = \binom{3}{2}$. Moreover, every non-zero monomial in $f \in \mathcal{R}$ is either of the form $f = x_i^2$ or $f = x_i x_j$, where $x_i, x_j \in \{x, y, z\}$. If $f = x_i x_j$, then it is clear that $f \in V_2$. If $f = x_i^2$, then $f = f_i - x_i x_{\overline{i+1}}$. In both cases, $f \in \mathcal{I}_2 \oplus V_2$ and it follows that $\mathcal{R}_m \subseteq \mathcal{I}_2 \oplus V_2$. The reverse inclusion holds trivially, so $\mathcal{R}_m = \mathcal{I}_2 \oplus V_2$.

- For $m = 3$, $\mathcal{R}_3 = \text{span}_{\mathbb{F}}\left\{x^3, x^2y, x^2z, xy^2, \ldots, z^3\right\}$. It can be verified that every monomial $f \in \mathcal{R}$ can be written as $f = g + h$, where $g \in \mathcal{I}_3$ and h is either 0 or $\pm xyz$. For example, $x^2y = g + h$, where $g = yf_1 - xf_2 \in \mathcal{I}_3$ and $h = xyz$. We leave it as an exercise to the reader to verify that this property holds for the other monomials as well. (In the next sections, we discuss a structural method to find such "decompositions".) The element xyz is not in $\mathcal{I}_3$, so for $V_3 = \text{span}_{\mathbb{F}}\{xyz\}$ we have $\dim_{\mathbb{F}} V_3 = 1 = \binom{3}{3}$ and $\mathcal{R}_m = \mathcal{I}_m \oplus V_m$.

In all cases above, $V_m = \mathcal{SF}_m$, the linear $\mathbb{F}$-span of all m-th degree square-free monomials, satisfies $\mathcal{R}_m = \mathcal{I}_m \oplus V_m$. ◊

Remark 74. The relation $\mathcal{R}_m = \mathcal{I}_m \oplus V_m$ does not uniquely define V_m (when $m \geq 2$). For example, both $V_2 = \text{span}_{\mathbb{F}}\{xy, xz, yz\}$ and $W_2 = \text{span}_{\mathbb{F}}\left\{x^2, y^2, z^2\right\}$ satisfy $\mathcal{I}_2 \oplus V_2 = \mathcal{R}_2 = \mathcal{I}_2 \oplus W_2$, but $V_2 \neq W_2$. Of course, both V_2 and W_2 have vector space dimension $\dim_{\mathbb{F}} \mathcal{R}_2 - \dim_{\mathbb{F}} \mathcal{I}_2 = 3$.

The following claim generalizes the result we found in Example 73.

Claim 75 (Direct Sum Claim)**.** *Let $p \geq 3$ be prime, let $\ell \geq 3$ be odd, and let $\mathcal{I} = \langle \mathcal{F}_{\text{hom}} \rangle$ be the ideal defined in Section 4.1. For all $m \in \mathbb{Z}_{\geq 0}$, $\mathcal{R}_m$ is the direct sum of $\mathcal{I}_m$ and $\mathcal{SF}_m$.*

Note that $\mathcal{SF}_m$ is an $\binom{\ell}{m}$-dimensional subspace of $\mathcal{R}_m$. By our discussion above, Claim 75 is equivalent to Equation (4.2).

The next sections are dedicated to proving the claim.

4.3 Proof Sketch for the Direct Sum Claim

Let p, ℓ, m be as in Claim 75. By the definition of direct sums (Definition 1), the Direct Sum Claim (Claim 75) is equivalent to $\mathcal{I}_m \cap \mathcal{SF}_m = \{0\}$ and $\mathcal{R}_m = \mathcal{I}_m + \mathcal{SF}_m$. Most of the difficulty of showing the claim comes from the latter statement. In this subsection, we discuss a systematic method to decompose a monomial f in $\mathcal{R}_m$ as $f = g + h$, where g and h are polynomials in $\mathcal{I}_m$ and $\mathcal{SF}_m$, respectively. They key insight here is that decompositions in $\mathcal{R}_m$ can be used to find decompositions in $\mathcal{R}_{m+1}$. This will allow us to prove the equality

$\mathcal{R}_m = \mathcal{I}_m + \mathcal{SF}_m$ using induction on $m \geq 0$. We will see that it is not immediately obvious that the decomposition method presented in this section always produces a decomposition, and we postpone the proof of this to Section 4.4.

Since polynomials in $\mathcal{R}_m$ are $\mathbb{F}_p$-linear combinations of monomials in $\mathcal{R}_m$, a decomposition of an arbitrary polynomial $f \in \mathcal{R}_m$ can be found by decomposing its terms. We therefore focus on decomposing terms in $\mathcal{R}_m$.

The following example shows the decompositions of all quadratic monomials in $\mathcal{R} = \mathbb{F}_p[x, y, z]$.

Example 76. Let $\ell = 3$ and $\mathcal{R} = \mathbb{F}_p[x, y, z]$. Then $\mathcal{F}_{\text{hom}} = (f_1, f_2, f_3)$, where

$$f_1 = x^2 + xy$$
$$f_2 = y^2 + yz$$
$$f_3 = z^2 + xz$$

The second-degree monomials in $\mathcal{R}$ are $\{x^2, y^2, z^2, xy, yz, xz\}$. For every $f \in \{xy, yz, xz\}$, we have the trivial decomposition $f = g + h$, where $g = 0 \in \mathcal{I}_2$ and $h = f \in \mathcal{SF}_2$, since these monomials are already square-free.

The monomial $f = x^2 \in \mathcal{R}_2$ can be written as

$$f = \left(x^2 + xy\right) - xy = g + h, \tag{4.3}$$

where $g = x^2 + xy \in \mathcal{I}_2$ and $h = -xy \in \mathcal{SF}_2$. Similarly, we have

$$y^2 = \left(y^2 + yz\right) - yz \in \mathcal{I}_2 + \mathcal{SF}_2 \tag{4.4}$$

and

$$z^2 = \left(z^2 + xz\right) - xz \in \mathcal{I}_2 + \mathcal{SF}_2,$$

since $y^2 + yz$ and $z^2 + xz$ are elements of $\mathcal{I}_2$ and $yz, xz \in \mathcal{SF}_2$. ◊

The next example shows how a decomposition in $\mathcal{R}_2$ from Example 76 can be used to decompose a monomial in $\mathcal{R}_3$.

Example 77. As in Example 76, let $\ell = 3$ and $\mathcal{R} = \mathbb{F}_p[x, y, z]$. We want to decompose $f = x^2 z \in \mathcal{R}_3$.

Multiplying both sides of Equation (4.3) by z yields

$$x^2 z = \left(x^2 + xy\right) z - xyz. \tag{4.5}$$

We already know that $x^2 + xy \in \mathcal{I}_2$, so Lemma 53 implies $\left(x^2 + xy\right) z \in \mathcal{I}_3$. Moreover, xyz is square-free. It follows that the sum in Equation (4.5) is a decomposition of $x^2 z \in \mathcal{R}_3$ into its $\mathcal{I}_3$ and $\mathcal{SF}_3$ components. ◊

In this example we saw that multiplying the decomposition of $x^2 \in \mathcal{R}_2$ by z yields a decomposition of $x^2 z$. This is possible, because the square-free part h of the decomposition of x^2 is not divisible by z. In general, the decomposition of a monomial f / x_k may have a square-free part which is divisible by x_k. In this case, we have to take some extra steps to find a decomposition of f, as illustrated by the following example.

Example 78. Again, let $\ell = 3$ and $\mathcal{R} = \mathbb{F}_p[x,y,z]$. We want to decompose $f = x^3 \in \mathcal{R}_3$.

Similarly to Example 77, we start by multiplying both sides of Equation (4.3) by x:

$$x^3 = \big(x^2 + xy\big)x - x^2y. \tag{4.6}$$

Lemma 53 implies that $\big(x^2+xy\big)x \in \mathcal{I}_3$, but the term x^2y is certainly not square-free. This does not mean all hope is lost; if $-x^2y$ can be decomposed as $-x^2y = g' + h'$, then we obtain a decomposition $x^3 = g + h \in \mathcal{I}_3 + \mathcal{SF}_3$, where $g = \big(x^2+xy\big)x + g' \in \mathcal{I}_3$ and $h = h' \in \mathcal{SF}_3$.

We multiply both sides of Equation (4.3) by $-y$ to obtain

$$-x^2y = -\big(x^2+xy\big)y + xy^2. \tag{4.7}$$

The term xy^2 is still not square-free, so we try to decompose xy^2. Multiplying both sides of Equation (4.4) by $-x$ yields

$$xy^2 = \big(y^2+yz\big)x - xyz. \tag{4.8}$$

The term xyz is square-free. Combining equations (4.6) to (4.8), we obtain

$$\begin{aligned} x^3 &= \big(x^2+xy\big)x - x^2y \\ &= \big(x^2+xy\big)x - \big(x^2+xy\big)y + xy^2 \\ &= \underbrace{\big(x^2+xy\big)x - \big(x^2+xy\big)y + \big(y^2+yz\big)x}_{:=g} + \underbrace{(-xyz)}_{:=h} \end{aligned}$$

Note that g is a sum of polynomials in $\mathcal{I}_3$ and is therefore itself a polynomial in $\mathcal{I}_3$. We also see that $h \in \mathcal{SF}_3$, so $f = g + h$ is a decomposition of f into its $\mathcal{I}_3$ and $\mathcal{SF}_3$ components. ◊

Assume, for a moment, that we know how to decompose non-zero terms in $\mathcal{R}_m$ for some integer $m \geq 2$. (The cases $m = 0$ and $m = 1$ are trivial.) We now discuss a method to find decompositions of non-zero terms in $\mathcal{R}_{m+1}$ as well. To this end, let $f_{(1)} = c\mathbf{x}^\alpha$ be a non-zero term in $\mathcal{R}_{m+1}$ that is not already square-free. Then there exists at least one variable x_k dividing $f_{(1)}$ such that $\tilde{f}_{(1)} \coloneqq f_{(1)}/x_k \in \mathcal{R}_m$ is still not square-free. By assumption, there exist $\tilde{g}_{(1)} \in \mathcal{I}_m$ and $\tilde{h}_{(1)} \in \mathcal{SF}_m$ such that $\tilde{f} = \tilde{g}_{(1)} + \tilde{h}_{(1)}$. Define $g_{(1)} = x_k\tilde{g}_{(1)}$ and $h_{(1)} = x_k\tilde{h}_{(1)}$, so that $f_{(1)} = g_{(1)} + h_{(1)}$. (If we allowed x_k such that $f_{(1)}/x_k$ is square-free, we would have $g_{(1)} = 0$ and $h_{(1)}$, which does not help us.)

As in the examples, it follows from Lemma 53 that $g_{(1)}$ is in $\mathcal{I}_{m+1}$. If $\tilde{h}_{(1)}$ is not divisible by x_k, then $h_{(1)}$ is square-free and we have found the desired decomposition. If $h_{(1)} = \tilde{h}_{(1)} = 0$, we have a decomposition as well. Otherwise, if $\tilde{h}_{(1)}$ is a non-zero term divisible by x_k, we set $f_{(2)} = h_{(1)}$ and repeat. Continuing,

we obtain the following system:

$$\begin{aligned} f_{(1)} &= g_{(1)} + h_{(1)} \\ f_{(2)} &= g_{(2)} + h_{(2)} \\ &\vdots \\ f_{(j)} &= g_{(j)} + h_{(j)} \\ &\vdots \end{aligned} \tag{4.9}$$

In Section 4.4, we will see that every $h_{(i)}$ is again a term.

The system above is finite if and only if some $h_{(j)}$ is square-free or zero. In this case, it follows from Equation (4.9) that $f_{(1)} = \sum_{k=1}^{j} g_{(k)} + h_{(j)}$. By construction, we have $\sum_{k=1}^{j} g_{(k)} \in \mathcal{I}_{m+1}$ and $h_{(j)} \in \mathcal{SF}_{m+1}$, from which it follows that $f_{(1)} \in \mathcal{I}_{m+1} + \mathcal{SF}_{m+1}$.

Now consider the case where the system is infinite, i.e. none of the $h_{(k)}$ are square-free. Observe that every $f_{(k)}$ is a term: $f_{(1)}$ is a term and, by construction, $f_{(i)} = h_{(i-1)}$ is a term for all $i > 1$. There are only finitely many monomials, so there must be $1 \leq i < j$ such that $f_{(i)} = d f_{(j)}$ for some $d \in \mathbb{F}_p \setminus \{0\}$. It can be shown that $d \in \{-1, 1\}$, but we will not prove this fact until the formal proof in Section 4.4. Suppose first that $1 \leq i < j$ are such that $f_{(i)} = -f_{(j)}$. Then,

$$f_{(i)} = \sum_{k=i}^{j-1} g_{(k)} + h_{(j-1)} = \sum_{k=i}^{j-1} g_{(k)} + f_{(j)} = \sum_{k=i}^{j-1} g_{(k)} - f_{(i)}.$$

Adding $f_{(i)}$ to both sides gives $2f_{(i)} = \sum_{k=i}^{j-1} g_{(k)} \in \mathcal{I}_{m+1}$. Since $p > 2$, 2 is invertible in $\mathbb{F}_p$ and we see that $f_{(i)} \in \mathcal{I}_{m+1}$. We conclude that $f_{(1)} = \sum_{k=1}^{i-1} g_{(k)} + f_{(i)} \in \mathcal{I}_{m+1} \subseteq \mathcal{I}_{m+1} + \mathcal{SF}_{m+1}$.

If there exist no $1 \leq i < j$ such that $f_{(i)} = -f_{(j)}$, then at least there exist $1 \leq i < j$ with $f_{(i)} = f_{(j)}$. This yields the equation

$$f_{(i)} = \sum_{k=i}^{j-1} g_{(k)} + h_{(j-1)} = \sum_{k=i}^{j-1} g_{(k)} + f_{(j)} = \sum_{k=i}^{j-1} g_{(k)} + f_{(i)}.$$

From this equation it follows that $\sum_{k=i}^{j-1} g_{(k)} = 0$, but we don't learn anything new about $f_{(i)}$.

In the following example, the system in Equation (4.9) is infinite.

Example 79. Let $\ell = 3$ and $\mathcal{R} = \mathbb{F}_p[x, y, z]$. We want to decompose $f_{(1)} = x^2yz \in \mathcal{R}_4$. Since there exist no square-free monomials of degree 4 in $\mathcal{R}$, a decomposition exists if and only if $f_{(1)} \in \mathcal{I}_4$. Necessarily, the system in Equation (4.9) is infinite.

In the first step we divide $f_{(1)}$ by one of its divisors x_k such that $\tilde{f}_{(1)} = f_{(1)}/x_k$ is still not square-free. In previous examples, there was only one possibility for x_k, but in this example there are multiple possibilities.

We first consider the system obtained by choosing x_k such that k is minimal:

$$\begin{aligned} x^2yz &= y \cdot \left(\tilde{g}_{(1)} - xyz\right) = g_{(1)} - xy^2z \\ -xy^2z &= x \cdot \left(\tilde{g}_{(2)} - xyz\right) = g_{(2)} - x^2yz \\ &\vdots \end{aligned}$$

It follows that $2x^2yz = g_{(1)} + g_{(2)} \in \mathcal{I}$, so $f = x^2yz \in \mathcal{I}$, as predicted.

Next, we choose the x_k such that k is maximal:

$$\begin{aligned} x^2yz &= z \cdot \left(\tilde{g}_{(1)} + xyz\right) = g_{(1)} + xyz^2 \\ xyz^2 &= y \cdot \left(\tilde{g}_{(2)} + xyz\right) = g_{(2)} + xy^2z \\ &\vdots \end{aligned}$$

This system does not yield a decomposition for f. ◊

For every term $f \in \mathcal{R}_{m+1}$, we may determine a system as in Equation (4.9), given that we know how to decompose monomials $\tilde{f} \in \mathcal{R}_m$. The last example shows that, depending on how we choose the x_k, we may or may not derive a decomposition from this system. In the example, the strategy of choosing x_k such that k is minimal yielded a decomposition of f. In the next section, we will see that this strategy *always* produces decompositions. Some additional information on the structure of decompositions is needed to prove this.

4.4 Proving the Direct Sum Claim

In the decomposition system in Equation (4.9), every non-zero $h_{(i)}$ is obtained by multiplying the square-free term $\tilde{h}_{(i)}$ by some variable x_k. If x_k does not divide $\tilde{h}_{(i)}$, then $h_{(i)}$ is square-free. Otherwise, $h_{(i)}$ is "almost" square-free, except for x_k^2. This motivates the following definition.

Definition 80. Let f be a monomial in $\mathcal{R}_m$. We say that f is **almost square-free** if there exist $m-1$ indices $1 \leq i_1 < \ldots < i_{m-1} \leq \ell$ and some $1 \leq r \leq m-1$ such that $f = x_{i_r} \cdot x_{i_1} x_{i_2} \cdots x_{i_{m-1}}$. We call i_r and x_{i_r} the **repeated index** and **repeated factor** of f, respectively.

We define $\mathcal{ASF}$ to be the $\mathbb{F}_p$-vector space spanned by the almost square-free monomials of $\mathcal{R}$.

We extend the notion of almost square-free monomials to terms $c\mathbf{x}^\alpha \in \mathcal{R}$ (where $c \neq 0$) by calling $c\mathbf{x}^\alpha$ almost square-free if $\mathbf{x}^\alpha$ is.

By induction on $m \geq 0$, we will prove that $\mathcal{R}_m = \mathcal{I}_m + \mathcal{SF}_m$ using the decomposition method presented in the previous section. As indicated in the proof sketch, the induction hypothesis does not only need to include the statement $\mathcal{R}_m = \mathcal{I}_m + \mathcal{SF}_m$, but also needs to include some information on the structure

of decompositions $f = g + h \in \mathcal{I}_m + \mathcal{SF}_m$ to show that we can always find a system as in Equation (4.9) from which a decomposition of f can be derived. (That is, we do not end up with equations of the form $f_{(i)} = \sum_{k=i}^{j-1} g_{(k)} + f_{(i)}$.) The following theorem captures these requirements. Recall from Section 4.1 that $\overline{n}$ denotes the unique integer $m \in \{1, \dots, \ell\}$ such that $n - m \in \ell\mathbb{Z}$.

Theorem 81. *Let $p \geq 3$ be prime, let $\ell \geq 3$ be odd, and let $\mathcal{I} = \langle \mathcal{F}_{\text{hom}} \rangle$ be the ideal defined in Section 4.1. For all $m \in \mathbb{Z}_{\geq 0}$, the following properties hold:*

1. *$\mathcal{R}_m = \mathcal{I}_m + \mathcal{SF}_m$.*
2. *Suppose that $2 \leq m \leq \ell$ and that $f = c\boldsymbol{x}^\alpha$ is an almost square-free term in $\mathcal{ASF}_m$ with repeated factor x_{i_r}. Then there exists a decomposition $f = g + h$, with $g \in \mathcal{I}_m$ and $h = (-1)^k \cdot x_{\overline{i_r+k}} \cdot \text{sqfree}(f) \in \mathcal{SF}_m$, where $k \geq 1$ is the minimal positive integer such that $x_{\overline{i_r+k}}$ does not divide f.*

Before proceeding with the proof of Theorem 81, we prove a number of lemmas for special cases. These lemmas are of great importance for the final proof, since they describe the structure of h for certain decompositions $f = g + h \in \mathcal{I}_m + \mathcal{SF}_m$.

Lemma 82. *Suppose that Theorem 81 holds for a fixed $2 \leq m < \ell$. Let $f = c \cdot x_{i_1}^2 x_{i_2} \cdots x_{i_m}$ be an almost square-free term in $\mathcal{ASF}_{m+1}$. Then there exists a decomposition $f = g + h$, where $g \in \mathcal{I}_{m+1}$ and $h \in \mathcal{SF}_{m+1}$ are as described in property 2 of Theorem 81.*

Proof. Let $f_{(1)} = f$. Then $\tilde{f}_{(1)} = f_{(1)}/x_{i_2} = c \cdot x_{i_1}^2 x_{i_3} \cdots x_{i_m} \in \mathcal{ASF}_m$. By assumption, Theorem 81 holds for m, so we can write $\tilde{f}_{(1)} = \tilde{g}_{(1)} + \tilde{h}_{(1)}$ such that $\tilde{g}_{(1)} \in \mathcal{I}_m$ and

$$\tilde{h}_{(1)} = (-1)^{k_1} c \cdot x_{\overline{i_1+k_1}} \cdot x_{i_1} x_{i_3} \cdots x_{i_m} \in \mathcal{SF}_m,$$

where $k_1 \geq 1$ is the minimal positive integer such that $x_{\overline{i_1+k_1}}$ does not divide $\tilde{f}_{(1)}$. Thus, $\overline{i_1 + k_1}$ is not equal to any of $i_1, i_3, \dots, i_m$.

If, in addition, it holds that $\overline{i_1 + k_1} \neq i_2$, then

$$h_{(1)} = x_{i_2}\tilde{h}_{(1)} = (-1)^{k_1} c \cdot x_{\overline{i_1+k_1}} \cdot x_{i_1} x_{i_2} x_{i_3} \cdots x_{i_m}$$

is a square-free term of degree $m + 1$. Moreover, $g_{(1)} = x_{i_2}\tilde{g}_{(1)}$ is an element of $\mathcal{I}_{m+1}$ by Lemma 53. We can therefore write $f_{(1)} = g_{(1)} + h_{(1)} \in \mathcal{I}_{m+1} + \mathcal{SF}_{m+1}$. To show minimality of k_1, suppose that there exists some $1 \leq k < k_1$ such that $x_{\overline{i_1+k}}$ does not divide f. Then $x_{\overline{i_1+k}}$ does not divide $\tilde{f}_{(1)}$ either, contradicting the fact that k_1 is minimal among the positive integers k such that $x_{\overline{i_1+k}}$ does not divide $\tilde{f}_{(1)}$. We conclude that $f = g_{(1)} + h_{(1)}$ is a decomposition in the form of property 2 of Theorem 81.

Now, suppose that $\overline{i_1+k_1}=i_2$. Then

$$h_{(1)}=(-1)^{k_1}\,c\cdot x_{i_1}x_{i_2}^2x_{i_3}\cdots x_{i_m}\in\mathcal{ASF}_{m+1}.$$

Let $f_{(2)}=h_{(1)}$ and write

$$\tilde{f}_{(2)}=f_{(2)}/x_{i_1}=(-1)^{k_1}\,c\cdot x_{i_2}^2x_{i_3}\cdots x_{i_m}\in\mathcal{ASF}_m.$$

By assumption, there exists a decomposition $\tilde{f}_{(2)}=\tilde{g}_{(2)}+\tilde{h}_{(2)}$, where $\tilde{g}_{(2)}\in\mathcal{I}_m$ and

$$\tilde{h}_{(2)}=(-1)^{k_1+k_2}\,c\cdot x_{\overline{i_2+k_2}}\cdot x_{i_2}x_{i_3}\cdots x_{i_m}\in\mathcal{SF}_m,$$

where $k_2\geq 1$ is minimal among the positive integers k for which $x_{\overline{i_2+k}}$ does not divide $\tilde{f}_{(2)}$. We will show that $\overline{i_2+k_2}\neq i_1$.

By definition, $k_1\geq 1$ is the smallest positive integer such that $x_{\overline{i_1+k_1}}$ does not divide $\tilde{f}_{(1)}$. If $k_1\geq 2$, then x_{i_1+1} is a monomial strictly between x_{i_1} and x_{i_2} dividing $\tilde{f}_{(1)}$ and therefore divides $f_{(1)}$. This is not possible, since there is no divisor of f between x_{i_1} and x_{i_2}, so k_1 must be 1. Thus, $i_2=i_1+1$. The condition $\overline{i_2+k_2}=i_1$ is therefore equivalent to $k_2=\ell-1$. Consequently, we can show that $\overline{i_2+k_2}\neq i_1$ by showing that $k_2\neq\ell-1$.

If $k_2=\ell-1$, then every variable in

$$M=\left\{x_{\overline{i_2}},x_{\overline{i_2+1}},\ldots,x_{\overline{i_2+\ell-2}}\right\}$$

must divide $\tilde{f}_{(2)}$. Now, M contains $\ell-1$ distinct variables, while $\tilde{f}_{(2)}$ contains $m-1<\ell-1$ distinct variables. This is a contradiction and we conclude that $k_2<\ell-1$.

It follows that $\overline{i_2+k_2}\neq i_1$, so

$$h_{(2)}=x_{i_1}\tilde{h}_{(2)}=(-1)^{k_1+k_2}\,c\cdot x_{\overline{i_2+k_2}}\cdot x_{i_1}x_{i_2}\cdots x_{i_m}$$

is a square-free term in $\mathcal{SF}_{m+1}$. Notice that $\overline{i_2+k_2}=\overline{i_1+k_1+k_2}$. Defining $k=k_1+k_2$, we have

$$h_{(2)}=(-1)^{k}\,c\cdot x_{\overline{i_1+k}}\cdot x_{i_1}x_{i_2}\cdots x_{i_m}\in\mathcal{SF}_{m+1}.$$

From minimality of k_1 and k_2 it follows that k is the smallest positive integer such that $x_{\overline{i_1+k}}$ does not divide f. We conclude that, for $g=g_{(1)}+g_{(2)}\in\mathcal{I}_{m+1}$ and $h=h_{(2)}\in\mathcal{SF}_{m+1}$, $f=g+h$ is a decomposition as described in property 2 of Theorem 81. □

The next lemma generalizes the result of Lemma 82 to the case where the repeated index is any of $i_1,\ldots,i_m$.

Lemma 83. *Suppose that Theorem 81 holds for some fixed $2\leq m<\ell$. Let $f=c\cdot x_{i_r}\cdot x_{i_1}x_{i_2}\cdots x_{i_m}$ be an almost square-free term in $\mathcal{ASF}_{m+1}$ with repeated factor x_{i_r}. Then there exists a decomposition $f=g+h$, where $g\in\mathcal{I}_{m+1}$ and $h\in\mathcal{SF}_{m+1}$ are as described in property 2 of Theorem 81.*

Proof. For all $k \in \mathbb{Z}$, let $S^k \colon \mathcal{R} \to \mathcal{R}$ be the $\mathbb{F}_p$-algebra automorphism defined by $x_i \mapsto x_{\overline{i+k}}$. Let $\mathcal{F}_{\text{hom}} = (f_1, \ldots, f_s)$ be as defined in Section 4.1. We first note that $\mathcal{F}_{\text{hom}}$ is S^k-invariant in the sense that $S^k(\{f_1, \ldots, f_s\}) = \{f_1, \ldots, f_s\}$ for all $k \in \mathbb{Z}$. (Note that S^k does not necessarily map f_i to f_i.) From Lemma 53 we see that $\mathcal{I}_{m+1}$ is S^k-invariant as well.

Let $f' = S^{1-i_r}(f)$. Then $f' = c \cdot x_{j_1}^2 x_{j_2} \cdots x_{j_m}$, where we set $j_k = \overline{i_{\overline{k-1+r}} + 1 - i_r}$ for all $1 \leq k \leq m$. (The inner $\overline{k+1-r}$ wraps around to 1 at m.) Observe that $j_1 < \cdots < j_m$, so we can now apply Lemma 82 to write $f' = g' + h'$, where $g' \in \mathcal{I}_{m+1}$ and

$$h' = (-1)^k c \cdot x_{\overline{j_1+k}} \cdot x_{j_1} x_{j_2} \cdots x_{j_m} \in \mathcal{SF}_{m+1}.$$

Here, k is the smallest positive integer such that $x_{\overline{j_1+k}}$ does not divide f'. Let $g = S^{i_r-1}(g')$ and $h = S^{i_r-1}(h')$. By construction, we have $f = S^{i_r-1}(f') = g + h$. Since $\mathcal{I}_{m+1}$ is S^k-invariant, we have $g \in \mathcal{I}_{m+1}$. Moreover,

$$h = (-1)^k c \cdot x_{\overline{i_r+k}} \cdot x_{i_1} x_{i_2} \cdots x_{i_m} \in \mathcal{SF}_{m+1}.$$

Note that $x_{\overline{i_r+k'}}$ divides f if and only if $x_{\overline{j_1+k'}}$ divides f', so k is also the smallest positive integer such that $x_{\overline{i_r+k}}$ does not divide f. We conclude that $f = g + h$ is a decomposition as described in property 2 of Theorem 81. □

The following lemma is the analogon of Lemma 82 for the case $m = \ell$.

Lemma 84. *Suppose that Theorem 81 holds for $m = \ell$. Let $f = c \cdot x_r \cdot x_1 x_2 \cdots x_\ell$ be an almost square-free term in $\mathcal{ASF}_{\ell+1}$ with repeated factor x_r, where $r \in \{1, 2\}$. Then $f \in \mathcal{I}_{\ell+1}$.*

Proof. Let $f_{(1)} = f$ and define $r' = 2$ if $r = 1$ and $r' = 1$ if $r = 2$. Then

$$\tilde{f}_{(1)} = f / x_{r'} = c \cdot x_r \cdot x_r x_3 \cdots x_\ell \in \mathcal{ASF}_\ell.$$

By assumption, we can write $\tilde{f}_{(1)} = \tilde{g}_{(1)} + \tilde{h}_{(1)}$ such that $\tilde{g}_{(1)} \in \mathcal{I}_\ell$ and

$$\tilde{h}_{(1)} = (-1)^{k_1} c \cdot x_{\overline{r+k_1}} \cdot x_r x_3 \cdots x_\ell \in \mathcal{SF}_\ell,$$

where k_1 is the minimal positive integer such that $x_{\overline{r+k_1}}$ does not divide $\tilde{f}_{(1)}$. Since $x_r x_3 \cdots x_\ell$ is divisible by every x_i except for $x_{r'}$, $\tilde{h}_{(1)}$ can only be square-free if $\overline{r + k_1} = r'$. Let $g_{(1)} = x_{r'} \tilde{g}_{(1)} \in \mathcal{I}_{\ell+1}$ and let

$$h_{(1)} = x_{r'} \tilde{h}_{(1)} = (-1)^{k_1} c \cdot x_{r'}^2 \cdot x_3 \cdots x_\ell \in \mathcal{ASF}_{\ell+1}.$$

Let $f_{(2)} = h_{(1)}$ and $\tilde{f}_{(2)} = f_{(2)} / x_r \in \mathcal{ASF}_\ell$. By assumption, we can write $\tilde{f}_{(2)} = \tilde{g}_{(2)} + \tilde{h}_{(2)}$, where $\tilde{g}_{(2)} \in \mathcal{I}_{\ell+1}$ and

$$\tilde{h}_{(2)} = (-1)^{k_1+k_2} c \cdot x_{\overline{r'+k_2}} \cdot x_{r'} x_3 \cdots x_\ell \in \mathcal{SF}_\ell,$$

where $k_2 \geq 1$ is the minimal positive integer such that $x_{\overline{r'+k_2}}$ does not divide $\tilde{f}_{(2)}$. The only possibility for $\tilde{h}_{(2)}$ to be square-free is if $\overline{r'+k_2} = r$. Defining $h_{(2)} = x_r \tilde{h}_{(2)}$, we find

$$\begin{aligned} h_{(2)} &= (-1)^{k_1+k_2}\, c \cdot x_r^2 \cdot x_{r'} x_3 \cdots x_\ell \\ &= (-1)^{k_1+k_2}\, c \cdot x_r \cdot x_1 x_2 \cdots x_\ell \,. \\ &= (-1)^{k_1+k_2}\, f_{(1)} \end{aligned}$$

Now, $\overline{r+k_1} = r' = \overline{r'}$ and $\overline{r'+k_2} = r = \overline{r}$ imply

$$\overline{k_1+k_2} = \overline{(r'-r)+(r-r')} = 0.$$

Since k_1 and k_2 are positive and both strictly smaller than ℓ, it follows that $k_1 + k_2 = \ell$. But ℓ is odd, so $h_{(2)} = -f_{(1)}$. We therefore have

$$\begin{aligned} f_{(1)} &= g_{(1)} + h_{(1)} \\ &= g_{(1)} + f_{(2)} \\ &= g_{(1)} + g_{(2)} + h_{(2)} \\ &= g_{(1)} + g_{(2)} - f_{(1)} \end{aligned}$$

From these equations, we obtain $2f_{(1)} = g_{(1)} + g_{(2)} \in \mathcal{I}_{\ell+1}$. Notice that 2 is invertible in $\mathbb{F}_p$, so $f = f_{(1)} \in \mathcal{I}_{\ell+1}$. □

Similarly to how Lemma 83 extends Lemma 82, the next lemma extends Lemma 84 to the case where the repeated index is any of $1, \ldots, \ell$.

Lemma 85. *Suppose that Theorem 81 holds for $m = \ell$. Let $f = c \cdot x_r \cdot x_1 x_2 \cdots x_\ell$ be an almost square-free term in $\mathcal{ASF}_{\ell+1}$ with repeated factor x_r. Then $f \in \mathcal{I}_{\ell+1}$.*

Proof. Let $S^k \colon \mathcal{R} \to \mathcal{R}$ be the $\mathbb{F}_p$-algebra automorphism defined in the proof of Lemma 83. Similar to the proof of Lemma 83, we can write $S^{1-i_r}(f) = c \cdot x_{j_1}^2 x_{j_2} \cdots x_{j_m}$ for indices $j_1 < \cdots < j_m$. By Lemma 84, $S^{1-i_r}(f)$ is in $\mathcal{I}_{\ell+1}$. Since $\mathcal{I}_{\ell+1}$ is S^k-invariant, we conclude that $f \in \mathcal{I}_{\ell+1}$. □

Lemma 86. *Suppose that Theorem 81 holds for some fixed $m > \ell$. Then any term $f = c\boldsymbol{x}^\alpha \in \mathcal{R}_{m+1}$ is in $\mathcal{I}_{m+1}$.*

Proof. Let x_i be a monomial dividing f and let $\tilde{f} = f/x_i \in \mathcal{R}_m$. By assumption, we can write $\tilde{f} = \tilde{g} + \tilde{h}$ such that $\tilde{g} \in \mathcal{I}_m$ and $\tilde{h} \in \mathcal{SF}_m$. But $\mathcal{SF}_m = \{0\}$, since there exist no square-free terms of degree $> \ell$. Thus, $f = x_i \tilde{f} = x_i \tilde{g} \in \mathcal{I}_{m+1}$. □

We are now in a position to prove the theorem.

Proof of Theorem 81. The proof is by induction on $m \geq 0$. Note that in order to prove $\mathcal{I}_m + \mathcal{SF}_m = \mathcal{R}_m$, it suffices to show that every monomial $\mathbf{x}^\alpha \in \mathcal{R}_m$ can be written as a sum of polynomials in $\mathcal{I}_m$ and $\mathcal{SF}_m$. The decomposition of a polynomial in $\mathcal{R}_m$ is then simply an $\mathbb{F}_p$-linear combination of the decomposition of its monomials.

Base case $m = 0$**.** For $m = 0$ we have $\mathcal{R}_0 = \mathcal{SF}_0 = \mathbb{F}_p$ and $\mathcal{I}_0 = \{0\}$, so $\mathcal{R}_0 = \mathcal{I}_0 + \mathcal{SF}_0$.

Base case $m = 1$**.** If $m = 1$, then $\mathcal{I}_1 = \{0\}$. All monomials in $\mathcal{R}_1$ are necessarily square-free, so $\mathcal{R}_1 = \mathcal{I}_1 + \mathcal{SF}_1$.

Base case $m = 2$**.** Let $f = x_i x_j$ be a monomial in $\mathcal{R}_2$. If $i \neq j$, then f is already square-free and can be written as $f = g + h$ for $g = 0 \in \mathcal{I}_2$ and $h = f \in \mathcal{SF}_2$. If $i = j$, then f is almost square-free and can be written as $f = g + h$ for $g = x_i^2 + x_i x_{\overline{i+1}} \in \mathcal{I}_2$ and $h = -x_i x_{\overline{i+1}} \in \mathcal{SF}_2$. Note that this decomposition satisfies property 2 of the theorem.

Induction step. Assume that the theorem holds for some fixed $m \geq 2$. We want to show that properties 1 and 2 of the theorem hold for $m+1$ as well. To this end, let $f = \mathbf{x}^\alpha$ be a monomial in $\mathcal{R}_{m+1}$. We proceed by case analysis on m.

- Suppose $2 \leq m < \ell$. If f is almost square-free, then both $f \in \mathcal{I}_{m+1} + \mathcal{SF}_{m+1}$ and property 2 follow from Lemma 83. Otherwise, let x_i be a monomial that divides f and define $\tilde{f} / x_i \in \mathcal{R}_m$. By the induction hypothesis, we can write $\tilde{f} = \tilde{g} + \tilde{h}$, where $\tilde{g} \in \mathcal{I}_m$ and $\tilde{h} \in \mathcal{SF}_m$. Let $g = x_i \tilde{g} \in \mathcal{I}_{m+1}$ and $h = x_i \tilde{h}$. Then either $h \in \mathcal{SF}_{m+1}$ or $h \in \mathcal{ASF}_{m+1}$. In the latter case, Lemma 83 implies $h \in \mathcal{I}_{m+1} + \mathcal{SF}_{m+1}$, so in both cases, we have $f = g + h \in \mathcal{I}_{m+1} + \mathcal{SF}_{m+1}$.

- Suppose $m = \ell$. If f is almost square-free, then it follows from Lemma 85 that $f \in \mathcal{I}_{m+1} \subseteq \mathcal{I}_{m+1} + \mathcal{SF}_{m+1}$. Otherwise, let $\tilde{f} = f / x_i \in \mathcal{I}_m$ for some monomial x_i dividing f. By the induction hypothesis, there exist $\tilde{g} \in \mathcal{I}_m$ and $\tilde{h} \in \mathcal{SF}_m$ such that $\tilde{f} = \tilde{g} + \tilde{h}$. The only square-free monomial of degree $m = \ell$ is $x_1 \cdots x_\ell$, so $\tilde{h}$ must be divisible by $x_1, \ldots, x_\ell$. Define $g = x_i \tilde{g} \in \mathcal{I}_{m+1}$ and $h = x_i \tilde{h} \in \mathcal{ASF}_{m+1}$. Since h is almost square-free, it follows from Lemma 85 that $h \in \mathcal{I}_{m+1}$. Therefore, $f = g + h \in \mathcal{I}_{m+1} \subseteq \mathcal{I}_{m+1} + \mathcal{SF}_{m+1}$.

- Suppose $m > \ell$. It follows immediately from Lemma 86 that $f \in \mathcal{I}_{m+1} \subseteq \mathcal{I}_{m+1} + \mathcal{SF}_{m+1}$.

In all cases, we find that properties 1 and 2 hold, so the theorem holds for $m+1$ as well. □

The next result shows that the sums $\mathcal{R}_m = \mathcal{I}_m + \mathcal{SF}_m$ from Theorem 81 are actually direct sums, in the sense that $\mathcal{I}_m \cap \mathcal{SF}_m = \{0\}$ for all $m \geq 0$. As a by-product of the proof, we also learn that the generators $\mathcal{F}_{\text{hom}} = (f_1, \ldots, f_\ell)$ of $\mathcal{I}$ form a regular sequence.

Lemma 87. *Let $p \geq 3$ be prime, let $\ell \geq 3$ be odd, and let $\mathcal{I} = \langle \mathcal{F}_{\text{hom}} \rangle$ be the ideal defined in Section 4.1. The Direct Sum Claim (Claim 75) holds for all $m \geq 0$. That is, $\mathcal{R}_m = \mathcal{I}_m \oplus \mathcal{SF}_m$ for all $m \geq 0$. Additionally, $\mathcal{F}_{\text{hom}} = (f_1, \ldots, f_\ell)$ forms a regular sequence.*

Proof. Observe that $f_1, \ldots, f_\ell$ is a sequence of homogeneous polynomials, all of degree 2. Therefore, Lemma 56 implies that for all $m \geq 0$, $\dim_{\mathbb{F}}(\mathcal{R}_m/\mathcal{I}_m) = \dim_{\mathbb{F}} \mathcal{R}_m - \dim_{\mathbb{F}} \mathcal{I}_m$ is bounded from below by the m-th coefficient of the power series $\frac{1}{(1-t)^\ell} \cdot \prod_{k=1}^{\ell} \left(1 - t^2\right)$. Working out this power series, we find that

$$\frac{\prod_{k=1}^{\ell} \left(1 - t^2\right)}{(1-t)^\ell} = \frac{(1+t)^\ell (1-t)^\ell}{(1-t)^\ell} = (1+t)^\ell = \sum_{m=0}^{\infty} \binom{\ell}{m} \cdot t^m,$$

so $\dim_{\mathbb{F}} \mathcal{R}_m - \dim_{\mathbb{F}} \mathcal{I}_m \geq \binom{\ell}{m}$. Simultaneously, the equality $\mathcal{R}_m = \mathcal{I}_m + \mathcal{SF}_m$ from Theorem 81 implies $\dim_{\mathbb{F}} \mathcal{R}_m \leq \dim_{\mathbb{F}} \mathcal{I}_m + \dim_{\mathbb{F}} \mathcal{SF}_m = \dim_{\mathbb{F}} \mathcal{I}_m + \binom{\ell}{m}$. Combining the two inequalities yields $\dim_{\mathbb{F}} \mathcal{R}_m = \dim_{\mathbb{F}} \mathcal{I}_m + \dim_{\mathbb{F}} \mathcal{SF}_m$.

We can now conclude two things. First, $\mathcal{I}_m$ and $\mathcal{SF}_m$ must be disjoint, except for the zero vector, from which we obtain that the sum $\mathcal{R}_m = \mathcal{I}_m + \mathcal{SF}_m$ is a direct sum. Second, the inequality in Lemma 56 is in fact an equality for $f_1, \ldots, f_\ell$, so this sequence of polynomials is regular. □

The following corollary follows immediately from Corollary 57.

Corollary 88. *Let $p \geq 3$ be prime and let $\ell \geq 3$ be odd. The ideal $\mathcal{I} = \langle \mathcal{F}_{\text{hom}} \rangle$ defined in Section 4.1 has ideal degree 2^ℓ.*

A consequence of Corollary 88 is that the ideal degrees related to a single round of ATRAPOS is 2^ℓ when $\ell > 3$ is not divisible by 3.

Remark 89. Whenever $\ell \geq 3$ is divisible by 3, one can consider the homogeneous variant of the ideal induced by a single round of ATRAPOS. (This ideal is obtained by taking the top homogeneous parts of the polynomials $g_1, \ldots, g_\ell$ that arise from the "natural" modeling of ATRAPOS in Section 4.1.) It turns out that, just like when ℓ is not divisible by 3, these homogeneous ideals are generated by regular sequences. Hence, the ideal degree corresponding to a single round of ATRAPOS is 2^ℓ for all $\ell \geq 3$. A formal proof of this fact is beyond the scope of the thesis, but the reader is encouraged to adapt the techniques presented in this chapter to the case where $\ell \geq 3$ is a multiple of 3.

Chapter 5

Multi-Round Analysis

This chapter extends the results of Chapter 4 to multiple rounds of ATRAPOS. Let $\mathcal{F}_{\text{hom}} = (f_1, \ldots, f_\ell)$ be the homogeneous system corresponding to a single round of ATRAPOS, assuming that $\ell > 3$ is odd and not divisible by 3. A key property of the polynomials in $\mathcal{F}_{\text{hom}}$ is that they form a regular sequence of homogeneous polynomials of the same degree. Many of the results in this chapter only rely on this property and do not use further structure of $\mathcal{F}_{\text{hom}}$. We will therefore state many of the results in terms of regular sequences of homogeneous polynomials of the same degree.

Section 5.1 discusses the polynomial system corresponding to R rounds of ATRAPOS. Section 5.2 is concerned with compositions of regular systems of polynomials having the same degree. Next, Section 5.3 shows that "small perturbations" of these regular systems, obtained by adding lower degree terms, have the same ideal degree as the original regular system. Section 5.4 then combines these results to show that the ideal corresponding to R rounds of ATRAPOS has ideal degree $2^{\ell R}$.

5.1 Introduction

In Section 4.1, we worked out the polynomials $\mathcal{F}_{\text{inh}} = (g_1, \ldots, g_\ell)$ in the outer part of the state after applying the j-th round of ATRAPOS on the input state

$$\mathbf{a} = \begin{pmatrix} a_{0,2} & a_{1,2} & \cdots & a_{\ell-1,2} \\ a_{0,1} & a_{1,1} & \cdots & a_{\ell-1,1} \\ x_1 & x_2 & \cdots & x_\ell \end{pmatrix}.$$

(We assume that $\ell > 3$ is odd and not divisible by 3.) Note that $g_1, \ldots, g_\ell$ implicitly depend on the round constant c_j. We therefore write $\mathcal{F}_{j,\text{inh}} = \left(g_{j1}, \ldots, g_{j\ell}\right)$ to emphasize this dependence. Using this notation, applying ATRAPOS$[R]$ to $\mathbf{a}$ yields a state whose outer part is given by the composition $\mathcal{F}_{R,\text{inh}} \circ \cdots \circ \mathcal{F}_{1,\text{inh}}$. We could similarly write $\mathcal{F}_{j,\text{hom}} = \left(f_{j1}, \ldots, f_{j\ell}\right)$, for the top homogeneous parts of $\mathcal{F}_{j,\text{inh}}$. However, in the previous chapter, we saw that the top homogeneous

parts of $g_{j1}, \ldots, g_{j\ell}$ are independent of j, since they do not involve the round coefficient c_j. The top homogeneous parts of $\mathcal{F}_{j,\text{inh}}$ can therefore be unambiguously denoted by $\mathcal{F}_{\text{hom}} = (f_1, \ldots, f_\ell)$.

One may expect that the top homogeneous parts of the polynomials in $\mathcal{F}_{R,\text{inh}} \circ \cdots \circ \mathcal{F}_{1,\text{inh}}$ can be obtained by iterating $\mathcal{F}_{\text{hom}}$ R times. That is, one may expect that the top homogeneous parts of $\mathcal{F}_{R,\text{inh}} \circ \cdots \circ \mathcal{F}_{1,\text{inh}}$ are given by $\mathcal{F}_{\text{hom}}^{(R)}$ (see Definition 11). It turns out that this is indeed the case, but this is not a trivial fact. The following example shows that, for arbitrary $\mathbf{g} \in (K[x_1, \ldots, x_n])^n$ with top homogeneous parts $\mathbf{f} \in (K[x_1, \ldots, x_n])^n$, the top homogeneous parts of $\mathbf{g} \circ \mathbf{g}$ are not generally equal to $\mathbf{f} \circ \mathbf{f}$.

Example 90. Let K be an arbitrary field and let

$$\mathbf{g} = (x - y + 1, x - y) \in (K[x, y])^2 .$$

The top homogeneous parts of $\mathbf{g}$ are equal to

$$\mathbf{f} = (x - y, x - y) \in (K[x, y])^2 .$$

We have $\mathbf{g} \circ \mathbf{g} = (2, 1)$, but $\mathbf{f} \circ \mathbf{f} = (0, 0)$, so the top homogeneous parts of $\mathbf{g} \circ \mathbf{g}$ are certainly not equal to $\mathbf{f} \circ \mathbf{f}$. ◊

The reason that the top homogeneous parts of $\mathbf{g} \circ \mathbf{g}$ in Example 90 are not equal to $\mathbf{f} \circ \mathbf{f}$ is that the polynomials $\mathbf{f} = (x - y, x - y)$ do not form a regular sequence. In the next sections, we will see that if $\mathbf{f} = (x - y, x - y)$ is a regular sequence of polynomials of the same degree, the top homogeneous parts of $\mathbf{g} \circ \mathbf{g}$ are, in fact, equal to $\mathbf{f} \circ \mathbf{f}$.

Before continuing, we introduce some notation. First, in the remainder of this chapter, we let K be an arbitrary field and let $\mathcal{R} = K[x_1, \ldots, x_n]$ be a polynomial ring. For the next part, suppose that $\mathbf{g} = (g_1, \ldots, g_s) \in \mathcal{R}^s$ is a sequence of polynomials with top homogeneous parts $\mathbf{f} = (f_1, \ldots, f_s) \in \mathcal{R}^s$. (That is, $f_i = (g_i)_{\text{top}}$ for all $1 \leq i \leq s$.) If $\mathbf{f}$ is a regular sequence of homogeneous polynomials of the same degree, then many properties of $\mathbf{f}$ (e.g. the leading term ideal or ideal degree) are the same for $\mathbf{g}$. Thus, writing $\epsilon_i = g_i - f_i$ for all $1, \ldots, s$, we can view $\mathbf{g} = (f_1 + \epsilon_1, \ldots, f_s + \epsilon_s)$ as a "small perturbation" of $\mathbf{f}$.

Definition 91. Let $\mathbf{f} = (f_1, \ldots, f_s) \in \mathcal{R}^s$ be a regular sequence of d-th degree homogeneous polynomials. Given polynomials $\epsilon_1, \ldots, \epsilon_s \in \mathcal{R}$, we call $\mathbf{g} = (f_1 + \epsilon_1, \ldots, f_s + \epsilon_s) \in \mathcal{R}^s$ a **small perturbation** of $\mathbf{f}$ if $\deg \epsilon_i < d$ for all $1 \leq i \leq s$.

If a property of $\mathbf{f}$ in Definition 91 is the same for any small perturbation $\mathbf{g}$, we say that the property is invariant under small perturbations. For example, in the next sections we will see that the leading term ideal or ideal degree of regular sequences of homogeneous polynomials of the same degree are invariant under small perturbations.

5.2 Compositions of Homogeneous Regular Systems

In this section, we prove that, given a regular sequence $\mathbf{f} = (f_1, \ldots, f_n)$ of d_f-th degree homogeneous polynomials in $\mathcal{R}$, and another regular sequence $\mathbf{h} = (h_1, \ldots, h_n)$ of d_h-th degree homogeneous polynomials in $\mathcal{R}$, the sequence $h_1(\mathbf{f}), \ldots, h_n(\mathbf{f})$ again forms a regular sequence of $d_h d_f$-th degree homogeneous polynomials.

We first show that the top homogeneous part of compositions involving regular sequences behave predictably.

Lemma 92. *Let $f_1, \ldots, f_n$ be a regular sequence of polynomials in $\mathcal{R}$. Let $0 \leq k < n$. If $h \in K[x_{n-k}, \ldots, x_n]$ is a polynomial such that $h\big([f_{n-k}]_{\mathcal{Q}}, \ldots, [f_n]_{\mathcal{Q}}\big) = 0$ in $\mathcal{Q} = \mathcal{R}/\langle f_1, \ldots, f_{n-k-1}\rangle$, then $h = 0$.*

Proof. Let $m = \deg h$. By induction on (m, k) (with respect to the lexicographical ordering), we show that $h = 0$ holds for all $m \in \{-\infty\} \cup \mathbb{Z}_{\geq 0}$ and $0 \leq k \leq n$.

Base case. If $m \in \{-\infty, 0\}$ and $0 \leq k < n$, then $h = c$ for some $c \in K$. It follows that $[c]_{\mathcal{Q}} = h\big([f_{n-k}]_{\mathcal{Q}}, \ldots, [f_n]_{\mathcal{Q}}\big) = 0$ and we conclude that $c = 0$.

Induction Step. Now, let $m > 0$ and $0 \leq k < n$. Suppose that the lemma holds for all $\big(m', k'\big) < (m, k)$. Write $h = x_{n-k}q + r$ with $\deg q \leq \deg h - 1$ and $r \in K[x_{n-k+1}, \ldots, x_n]$. Then $h\big([f_{n-k}]_{\mathcal{Q}}, \ldots, [f_n]_{\mathcal{Q}}\big) = 0$ implies the equality $[f_{n-k}]_{\mathcal{Q}}\, q\big([f_1]_{\mathcal{Q}}, \ldots, [f_n]_{\mathcal{Q}}\big) + r\big([f_{n-k+1}]_{\mathcal{Q}}, \ldots, [f_n]_{\mathcal{Q}}\big) = 0$. Passing to the quotient space $\mathcal{Q}' = \mathcal{R}/\langle f_1, \ldots, f_{n-k}\rangle$, we find that $r\big([f_{n-k+1}]_{\mathcal{Q}'}, \ldots, [f_n]_{\mathcal{Q}'}\big) = 0$. If $k = 0$, then $r \in K$, so $\deg r \leq 0$. If $k > 0$, then $r \in K\big[x_{n-(k-1)}, \ldots, x_n\big]$. In both cases, we may apply the induction hypothesis to obtain $r = 0$. This yields $[f_k q(f_1, \ldots, f_n)]_{\mathcal{Q}} = 0$. Since $[f_k]_{\mathcal{Q}}$ is a non-zero-divisor, we must have $[q(f_1, \ldots, f_n)]_{\mathcal{Q}} = 0$. The induction hypothesis gives $q = 0$, and we conclude that $h = 0$. □

Proposition 93. *Let $\mathbf{f} = (f_1, \ldots, f_n) \in \mathcal{R}^n$ be a regular sequence of d-th degree homogeneous polynomials and let $\mathbf{g} = (f_1 + \epsilon_1, \ldots, f_n + \epsilon_n) \in \mathcal{R}^n$ be a small perturbation of $\mathbf{f}$. Let $h \in \mathcal{R}$ be an arbitrary polynomial. Then, the top homogeneous component of $h \circ \mathbf{g}$ is given by $(h(f_1 + \epsilon_1, \ldots, f_n + \epsilon_n))_{\text{top}} = h_{\text{top}}(f_1, \ldots, f_n)$. Moreover, the composition $h(f_1 + \epsilon_1, \ldots, f_n + \epsilon_n)$ has degree $d \cdot \deg h$.*

Proof. The statement is trivial for $h = 0$, so we assume $h \neq 0$. Write $h = \sum_\alpha c_\alpha \mathbf{x}^\alpha$ so that $h(\mathbf{g}) = \sum_\alpha c_\alpha \mathbf{g}^\alpha$. Observe that $(h(\mathbf{g}))_{d'} = 0$ for all $d' > d \cdot \deg h$, and $(h(\mathbf{g}))_{d \cdot \deg h} = \big(h_{\text{top}}(\mathbf{g})\big)_{d \cdot \deg h} = h_{\text{top}}(f_1, \ldots, f_n)$. Thus, it suffices to show that $h_{\text{top}}(f_1, \ldots, f_n)$ is non-zero. If $h_{\text{top}}(f_1, \ldots, f_n) = 0$, then Lemma 92 (with $k = n - 1$ and $\mathcal{Q} = \mathcal{R}/\langle 0\rangle$) implies $h_{\text{top}} = 0$. This is a contradiction ($h \neq 0$), so $h_{\text{top}}(f_1, \ldots, f_n)$ must be non-zero. This proves the proposition. □

The main trick in this section is to introduce auxiliary variables $y_1, \ldots, y_n$ to reduce the degrees of the involved polynomials.

Solving the polynomial system

$$\begin{aligned} h_1(\mathbf{f}(x_1,\ldots,x_n)) &= 0 \\ &\vdots \\ h_n(\mathbf{f}(x_1,\ldots,x_n)) &= 0 \end{aligned}$$

in $K[x_1,\ldots,x_n]$ is equivalent (in a precise sense that we will discuss shortly) to first solving

$$\begin{aligned} h_1(y_1,\ldots,y_n) &= 0 \\ &\vdots \\ h_n(y_1,\ldots,y_n) &= 0 \end{aligned}$$

in $K[y_1,\ldots,y_n]$ (or $K[x_1,\ldots,x_n,y_1,\ldots,y_n]$) and then solving

$$\begin{aligned} f_1(x_1,\ldots,x_n) &= y_1 \\ &\vdots \\ f_n(x_1,\ldots,x_n) &= y_n \end{aligned}$$

in $K[x_1,\ldots,x_n,y_1,\ldots,y_n]$.

To formalize the introduction of the auxiliary variables, we define $\mathcal{S} = K[x_1,\ldots,x_n,y_1,\ldots,y_n]$. Solving the system $\mathbf{h}(\mathbf{f}(\mathbf{x})) = 0$ in $\mathcal{R}$ is equivalent to solving $\mathbf{h}(\mathbf{y}) = \mathbf{f}(\mathbf{x}) - \mathbf{y} = 0$ in $\mathcal{S}$, in the sense that their affine varieties are equal if we only consider the x-coordinates of the affine variety of the second system. Formally, let $\mathcal{I} = \langle \mathbf{h}(\mathbf{f}) \rangle \subseteq \mathcal{R}$ and $\mathcal{J} = \langle \mathbf{h}(\mathbf{y}), \mathbf{f} - \mathbf{y} \rangle \subseteq \mathcal{S}$. Then $V(\mathcal{I}) = \pi_n(V(\mathcal{J}))$, where $\pi_n \colon K^{2n} \to K^n$ is the projection mapping $(a_1,\ldots,a_n,b_1,\ldots,b_n)$ to the n-tuple $(a_1,\ldots,a_n)$. In fact, the restriction $\pi_n \colon V(\mathcal{J}) \to V(\mathcal{I})$ is a bijection, since every $(a_1,\ldots,a_n) \in V(\mathcal{I})$ uniquely determines $(b_1,\ldots,b_n) \in K^n$ such that $(a_1,\ldots,a_n,b_1,\ldots,b_n) \in V(\mathcal{J})$.

However, we're ultimately interested in the ideal degree of $\mathcal{I}$ and the mere fact that $V(\mathcal{I})$ and $V(\mathcal{J})$ are in bijection does not guarantee that the ideal degrees of $\mathcal{I}$ and $\mathcal{J}$ (in their respective polynomial rings) are equal. As an example, consider the ideals $\langle x \rangle$ and $\langle x^2 \rangle$ in the polynomial ring $K[x]$. Both ideals have the same affine variety ($V(\langle x \rangle) = V(\langle x^2 \rangle) = \{0\}$), but their ideal degrees differ, since $K[x]/\langle x \rangle = \operatorname{span}_K \{1\}$ and $K[x]/\langle x^2 \rangle = \operatorname{span}_K \{1, x\}$. The following lemma shows that $\mathcal{R}/\mathcal{I}$ and $\mathcal{S}/\mathcal{J}$ are isomorphic as K-algebras, from which it follows that $\dim_K \mathcal{S}/\mathcal{J} = \dim_K \mathcal{R}/\mathcal{I}$. The lemma actually proves a slightly stronger result, which will be needed to prove that $h_1(\mathbf{f}),\ldots,h_n(\mathbf{f})$ forms a regular sequence.

Lemma 94. *Let $S = K[x_1,\ldots,x_n,y_1,\ldots,y_n]$. Let $\mathbf{f} = (f_1,\ldots,f_n) \in \mathcal{R}^n$ and $\mathbf{h} = (h_1,\ldots,h_n) \in \mathcal{R}^n$ be polynomial sequences. For all $0 \leq i \leq n$, let*

$$\mathcal{J}_i = \langle h_1(\mathbf{y}),\ldots,h_i(\mathbf{y}), f_1 - y_1,\ldots,f_n - y_n \rangle \subseteq \mathcal{S}$$

and

$$\mathcal{I}_i = \langle h_1(\mathbf{f}), \ldots, h_i(\mathbf{f}) \rangle \subseteq \mathcal{R}.$$

Then $\mathcal{S}/\mathcal{J}_i \cong \mathcal{R}/\mathcal{I}_i$ *as* K*-algebras.*

Proof. Fix $0 \leq i \leq n$ and let $\phi_i \colon \mathcal{S} \to \mathcal{R}/\mathcal{I}_i$ be the unique K-algebra homomorphism defined by $x_j \mapsto \left[x_j\right]_{\mathcal{R}/\mathcal{I}_i}$ and $y_j \mapsto \left[f_j\right]_{\mathcal{R}/\mathcal{I}_i}$ for all $1 \leq j \leq n$. This homomorphism is surjective, since for any $[f\,] \in \mathcal{R}/\mathcal{I}_i$ we have $f \in \mathcal{R} \subseteq \mathcal{S}$ and $\phi_i(f) = [f\,]_{\mathcal{R}/\mathcal{I}_i}$. Moreover, we claim that $\operatorname{Ker}\phi_i = \mathcal{J}_i$, from which Lemma 94 follows using the First Isomorphism Theorem for Algebras (see Figure 5.1).

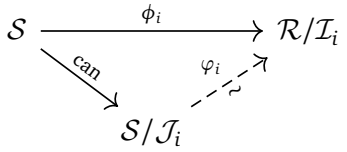

Figure 5.1: Depiction of the (surjective) K-algebra homomorphism ϕ_i and the induced isomorphism φ_i. Here, "can" denotes the canonical homomorphism which maps f to its equivalence class $[f\,]$.

We show that $\mathcal{J}_i \subseteq \operatorname{Ker}\phi_i$ (for all $1 \leq i \leq n$) by showing that ϕ_i maps the generators $h_1(\mathbf{y}), \ldots, h_i(\mathbf{y}), f_1 - y_1, \ldots, f_n - y_n$ of $\mathcal{J}_i$ to zero. First, note that for all $1 \leq j \leq i$, we have $\phi_i\big(h_j(\mathbf{y})\big) = h_j(\phi_i(y_1), \ldots, \phi_i(y_n))$, since ϕ_i is a K-algebra homomorphism. By definition, $h_j(\phi_i(y_1), \ldots, \phi_i(y_n))$ equals $h_j\big([f_1]_{\mathcal{R}/\mathcal{I}_i}, \ldots, [f_n]_{\mathcal{R}/\mathcal{I}_i}\big) = \left[h_j(f_1, \ldots, f_n)\right]_{\mathcal{R}/\mathcal{I}_i}$. (The last equality is again a consequence of the canonical homomorphism $\text{can}\colon \mathcal{R} \to \mathcal{R}/\mathcal{I}_i$ being a K-algebra homomorphism.) It follows that $\phi_i\big(h_j(\mathbf{y})\big) = \left[h_j(f_1, \ldots, f_n)\right]_{\mathcal{R}/\mathcal{I}_i} = \left[h_j(\mathbf{f})\right]_{\mathcal{R}/\mathcal{I}_i} = 0$. We also have

$$\phi_i\big(f_j - y_j\big) = \phi_i\big(f_j\big) - \phi_i\big(y_j\big) = \left[f_j\right]_{\mathcal{R}/\mathcal{I}_i} - \left[f_j\right]_{\mathcal{R}/\mathcal{I}_i} = 0$$

for all $1 \leq j \leq i$. Since ϕ_i maps a basis of $\mathcal{J}_i$ to zero, we must have $\mathcal{J}_i \subseteq \operatorname{Ker}\phi_i$.

For the converse inclusion, let $g = \sum_{\alpha,\beta} c_{\alpha,\beta}\mathbf{x}^\alpha\mathbf{y}^\beta \in \operatorname{Ker}\phi_i$. Then $\phi_i(g) = \left[\sum_{\alpha,\beta} c_{\alpha,\beta}\mathbf{x}^\alpha\mathbf{f}^\beta\right]_{\mathcal{R}/\mathcal{I}_i} = 0$ implies that $\sum_{\alpha,\beta} c_{\alpha,\beta}\mathbf{x}^\alpha\mathbf{f}^\beta \in \mathcal{I}_i$, so there exist polynomials $u_1, \ldots, u_i \in \mathcal{R}$ such that

$$\sum_{\alpha,\beta} c_{\alpha,\beta}\mathbf{x}^\alpha\mathbf{f}^\beta = \sum_{j=1}^{i} u_j h_j(\mathbf{f}). \tag{5.1}$$

It follows from $\left[y_j\right]_{\mathcal{S}/\mathcal{J}_i} = \left[f_j\right]_{\mathcal{S}/\mathcal{J}_i}$ (for all $1 \leq j \leq n$) that

$$[g]_{\mathcal{S}/\mathcal{J}_i} = \left[\sum_{\alpha,\beta} c_{\alpha,\beta}\mathbf{x}^\alpha\mathbf{y}^\beta\right]_{\mathcal{S}/\mathcal{J}_i} = \left[\sum_{\alpha,\beta} c_{\alpha,\beta}\mathbf{x}^\alpha\mathbf{f}^\beta\right]_{\mathcal{S}/\mathcal{J}_i}.$$

Applying Equation (5.1), followed by the equality $\left[y_j\right]_{\mathcal{S}/\mathcal{J}_i} = \left[f_j\right]_{\mathcal{S}/\mathcal{J}_i}$ again, then yields

$$[g]_{\mathcal{S}/\mathcal{J}_i} = \left[\sum_{j=1}^{i} u_j h_j(\mathbf{f})\right]_{\mathcal{S}/\mathcal{J}_i} = \left[\sum_{j=1}^{i} u_j h_j(\mathbf{y})\right]_{\mathcal{S}/\mathcal{J}_i} = 0.$$

It follows that $g \in \mathcal{J}_i$ and $\operatorname{Ker}\phi_i \subseteq \mathcal{J}_i$.

We now know that $\operatorname{Ker}\phi_i = \mathcal{J}_i$. The First Isomorphism Theorem for Algebras (Proposition 27) implies that the K-algebra homomorphism $\varphi_i \colon \mathcal{S}/\mathcal{J}_i \to \mathcal{R}/\mathcal{I}_i$ defined by $\left[x_j\right]_{\mathcal{S}/\mathcal{J}_i} \mapsto \left[x_j\right]_{\mathcal{R}/\mathcal{I}_i}$ and $\left[y_j\right]_{\mathcal{S}/\mathcal{J}_i} \mapsto \left[f_j\right]_{\mathcal{R}/\mathcal{I}_i}$ for all $1 \leq j \leq n$ is a well-defined K-algebra isomorphism. □

Lemma 94 allows us to reason about $\mathcal{R}/\langle \mathbf{h}(\mathbf{f})\rangle$ using the quotient space $\mathcal{S}/\langle \mathbf{h}(\mathbf{y}), \mathbf{f}-\mathbf{y}\rangle$. By assumption, $h_1, \ldots, h_n$ is a regular sequence of polynomials in $\mathcal{R}$. It is immediate that $h_1(\mathbf{y}), \ldots, h_n(\mathbf{y})$ forms a regular sequence in $K[y_1, \ldots, y_n]$ and Lemma 60 implies that the latter sequence also forms a regular sequence in $\mathcal{S} = K[x_1, \ldots, x_n, y_1, \ldots, y_n]$. This insight, together with the isomorphisms $\varphi_0, \ldots, \varphi_n \colon \mathcal{S}/\mathcal{J}_i \to \mathcal{R}/\mathcal{I}_i$ from Lemma 94 allows us to prove an important result on the composition of regular sequences.

Theorem 95. *Let $\mathbf{f} = (f_1, \ldots, f_n) \in \mathcal{R}^n$ and $\mathbf{h} = (h_1, \ldots, h_n) \in \mathcal{R}^n$ be regular sequences of homogeneous polynomials. Then $h_1(\mathbf{f}), \ldots, h_n(\mathbf{f})$ also forms a regular sequence in $\mathcal{R}$.*

Proof. To show that $h_1(\mathbf{f}), \ldots, h_n(\mathbf{f})$ forms a regular sequence in $\mathcal{R}$, we need to show that $\mathcal{I}_n = \langle h_1(\mathbf{f}), \ldots, h_n(\mathbf{f})\rangle$ is a proper ideal (i.e. it does not equal $\mathcal{R}$) and that for all $1 \leq i \leq n$, $[h_i(\mathbf{f})]$ is a non-zero-divisor for $\mathcal{R}/\mathcal{I}_{i-1}$. Here, $\mathcal{I}_{i-1} = \langle h_1(\mathbf{f}), \ldots, h_{i-1}(\mathbf{f})\rangle \subseteq \mathcal{R}$ is as in the proof of Lemma 94.

Recall that $\mathcal{I}_n$ is a proper ideal of $\mathcal{R}$ if and only if $1 \notin \mathcal{I}_n$. Every element of $\mathcal{I}_n$ can be written as $\sum_{i=1}^{n} u_i h_i(\mathbf{f})$, where $u_1, \ldots, u_n \in \mathcal{R}$. Since $\mathbf{f}$ and $\mathbf{h}$ form regular sequences, their degrees must be strictly positive (Remark 59). Moreover, the polynomials in $\mathbf{f}$ and $\mathbf{h}$ are homogeneous, so the 0-th degree homogeneous parts of these polynomials are zero. Consequently, the 0-th degree homogeneous part of every $h_i(\mathbf{f})$ is zero. Taking the 0-th degree homogeneous part of both sides of $\sum_{i=1}^{n} u_i h_i(\mathbf{f}) = 1$ yields the contradiction $0 = 1$. We conclude that $1 \notin \mathcal{I}_n$.

Next, suppose that $[u \cdot h_i(\mathbf{f})]_{\mathcal{R}/\mathcal{I}_{i-1}} = 0$ for some $1 \leq i \leq n$ and $u \in \mathcal{R}$. Let the ideal

$$\mathcal{J}_{i-1} = \langle h_1(\mathbf{y}), \ldots, h_{i-1}(\mathbf{y}), f_1 - y_1, \ldots, f_n - y_n\rangle \subseteq \mathcal{S}$$

and the K-algebra isomorphism $\varphi_{i-1} \colon \mathcal{S}/\mathcal{J}_{i-1} \to \mathcal{R}/\mathcal{I}_{i-1}$ be as in the proof of Lemma 94. Note that $\varphi_{i-1}\big([u \cdot h_i(\mathbf{y})]_{\mathcal{S}/\mathcal{J}_{i-1}}\big) = [u \cdot h_i(\mathbf{f})]_{\mathcal{R}/\mathcal{I}_{i-1}}$. Therefore,

$$[u \cdot h_i(\mathbf{y})]_{\mathcal{S}/\mathcal{J}_{i-1}} = \varphi_{i-1}^{-1}\big([u \cdot h_i(\mathbf{f})]_{\mathcal{R}/\mathcal{I}_{i-1}}\big) = \varphi_{i-1}^{-1}(0) = 0.$$

Since $[h_i(\mathbf{y})]_{\mathcal{S}/\mathcal{J}_{i-1}}$ is a non-zero-divisor, we must have $[u]_{\mathcal{S}/\mathcal{J}_{i-1}} = 0$. Applying φ_{i-1} to both sides then yields $[u]_{\mathcal{R}/\mathcal{I}_{i-1}} = \varphi_{i-1}\big([u]_{\mathcal{S}/\mathcal{J}_{i-1}}\big) = 0$, proving that every $[h_i(\mathbf{f})]_{\mathcal{R}/\mathcal{I}_{i-1}}$ is a non-zero-divisor. □

Corollary 96. *Let $f_1, \ldots, f_n \in \mathcal{R}$ be a regular sequence of d_f-th degree polynomials and let $h_1, \ldots, h_n \in \mathcal{R}$ be a regular sequence of d_h-th degree polynomials. Then $h_1(\mathbf{f}), \ldots, h_n(\mathbf{f})$ forms a regular sequence of homogeneous $d_h d_f$-th degree polynomials.*

Proof. It follows from Proposition 93 and Theorem 95 that $h_1(\mathbf{f}), \ldots, h_n(\mathbf{f})$ forms a regular sequence of $d_h d_f$-th degree polynomials. Homogeneity of the polynomials in the composition follows from the homogeneity of the polynomials $f_1, \ldots, f_n$ and $h_1, \ldots, h_n$. □

Corollary 97. *Let $\mathbf{f} = (f_1, \ldots, f_n) \in \mathcal{R}^n$ be a regular sequence of d-th degree homogeneous polynomials. For all $i \geq 0$, the i-th iteration $\mathbf{f}^{(i)} = \left(f_1^{(i)}, \ldots, f_n^{(i)}\right) \in \mathcal{R}^n$ of $\mathbf{f}$ forms a regular sequence of homogeneous d^i-th degree polynomials.*

Proof. The proof is by induction on $i \geq 0$. We know from Corollary 58 that $\mathbf{f}^{(0)} = (x_1, \ldots, x_n)$ forms a regular sequence of linear homogeneous polynomials. Next, suppose that $\mathbf{f}^{(i)}$ forms a regular sequence of d^i-th degree homogeneous polynomials for some $i \geq 0$. By Corollary 96, $\mathbf{f}^{(i+1)} = \mathbf{f} \circ \mathbf{f}^{(i)}$ forms a regular sequence of homogeneous polynomials of degree $d \cdot d^i = d^{i+1}$. □

The results obtained so far allow us to compute the ideal degree of the iterations of a polynomial system consisting of homogeneous d-th degree polynomials.

Theorem 98. *Let $f_1, \ldots, f_n \in \mathcal{R}$ be a regular sequence of homogeneous d-th degree polynomials. For all $i \geq 0$, let $\mathbf{f}^{(i)} = \left(f_1^{(i)}, \ldots, f_n^{(i)}\right) \in \mathcal{R}^n$ be the i-th iteration of $\mathbf{f}$. The ideal $\left\langle \mathbf{f}^{(i)} \right\rangle$ has ideal degree d^{in}.*

Proof. Fix $i \geq 0$ By Corollary 97, $f_1^{(i)}, \ldots, f_n^{(i)}$ forms a regular sequence of d^i-th degree homogeneous polynomials. It follows from Corollary 57 that $\left\langle \mathbf{f}^{(i)} \right\rangle$ has ideal degree d^{in}. □

5.3 Small Perturbations of Regular Systems

Let $\mathbf{f} = (f_1, \ldots, f_s) \in \mathcal{R}^s$ be a regular sequence of d-th degree homogeneous polynomials. We know from Corollary 57 that the ideal $\langle \mathbf{f} \rangle = \langle f_1, \ldots, f_s \rangle$ has ideal degree d^s. Recall that we call $\mathbf{g} = (f_1 + \epsilon_1, \ldots, f_n + \epsilon_s) \in \mathcal{R}^s$ a small perturbation of $\mathbf{f}$ if $\deg \epsilon_i < d$ for all $1 \leq i \leq s$. In this section, we will show that the ideal $\langle \mathbf{g} \rangle = \langle f_1 + \epsilon_1, \ldots, f_s + \epsilon_s \rangle$ generated by the small perturbation $\mathbf{g}$ of $\mathbf{f}$ has the same ideal degree as $\langle \mathbf{f} \rangle$.

We know from Lemma 38 that $\mathcal{R}/\langle \mathbf{f} \rangle$ is isomorphic as a K-vector space to $\operatorname{span}_K\{\mathbf{x}^\alpha \in \mathcal{R} \mid \mathbf{x}^\alpha \notin \langle \mathrm{LT}(\langle \mathbf{f} \rangle) \rangle\}$ for an arbitrary monomial ordering $\geq$. Similarly, $\mathcal{R}/\langle \mathbf{g} \rangle \cong \operatorname{span}_K\{\mathbf{x}^\alpha \in \mathcal{R} \mid \mathbf{x}^\alpha \notin \langle \mathrm{LT}(\langle \mathbf{g} \rangle) \rangle\}$ as K-vector spaces. We can therefore show that $\dim_K \mathcal{R}/\langle \mathbf{f} \rangle = \dim_K \mathcal{R}/\langle \mathbf{g} \rangle$ by showing that $\langle \mathrm{LT}(\langle \mathbf{f} \rangle) \rangle = \langle \mathrm{LT}(\langle \mathbf{g} \rangle) \rangle$ for

some monomial ordering $\geq$. This equality turns out to hold true when $\geq$ is an arbitrary graded monomial ordering. We prove it by separately showing the inclusions $\langle \mathrm{LT}(\langle \mathbf{f} \rangle) \rangle \subseteq \langle \mathrm{LT}(\langle \mathbf{g} \rangle) \rangle$ and $\langle \mathrm{LT}(\langle \mathbf{g} \rangle) \rangle \subseteq \langle \mathrm{LT}(\langle \mathbf{f} \rangle) \rangle$. We start by proving the former inclusion, which is the easiest of the two.

Lemma 99. *Let $\mathbf{f} = (f_1, \ldots, f_s) \in \mathcal{R}^s$ be a regular sequence of d-th degree homogeneous polynomials and let $\mathbf{g} = (f_1 + \epsilon_1, \ldots, f_s + \epsilon_s) \in \mathcal{R}^s$ be a small perturbation of $\mathbf{f}$. Then $\langle \mathrm{LT}(\langle \mathbf{f} \rangle) \rangle \subseteq \langle \mathrm{LT}(\langle \mathbf{g} \rangle) \rangle$ with respect to any graded monomial ordering $\geq$ on $\mathcal{R}$.*

Proof. Observe that it suffices to show the inclusion $\mathrm{LT}(\langle \mathbf{f} \rangle) \setminus \{0\} \subseteq \mathrm{LT}(\langle \mathbf{g} \rangle)$, since $\langle \mathrm{LT}(\langle \mathbf{f} \rangle) \rangle$ and $\langle \mathrm{LT}(\langle \mathbf{g} \rangle) \rangle$ are both ideals. Every non-zero element in $\mathrm{LT}(\langle \mathbf{f} \rangle)$ is of the form $h = \mathrm{LT}\left(\sum_{i=1}^{s} u_i f_i\right)$ for some polynomials $u_1, \ldots, u_s \in \mathcal{R}$. For all $1 \leq i \leq s$ and $d' \geq 0$, we define $(u_i)_{d'}$ to be the homogeneous part of degree d' of u_i. Every u_i can then be written as the finite sum $u_i = \sum_{d' \geq 0} (u_i)_{d'}$, so

$$h = \mathrm{LT}\left(\sum_{i=1}^{s} \sum_{d' \geq 0} (u_i)_{d'} f_i \right) = \mathrm{LT}\left(\sum_{d' \geq 0} \sum_{i=1}^{s} (u_i)_{d'} f_i \right).$$

Since $f_1, \ldots, f_s$ are homogeneous polynomials of degree d, every $(u_i)_{d'} f_i$ either has degree $d' + d$ or is zero. It follows that $h = \mathrm{LT}\left(\sum_{i=1}^{s} (u_i)_{d_0} f_i\right)$, where d_0 is the largest integer d' for which $\sum_{i=1}^{s} (u_i)_{d'} f_i$ is non-zero. (Such a d_0 exists, since $f \neq 0$.) This sum must have degree $d_0 + d$. The degree of $\sum_{i=1}^{s} (u_i)_{d_0} \epsilon_i$ is strictly smaller than $d_0 + d$, since the ϵ_i have degree $< d$. Therefore,

$$h = \mathrm{LT}\left(\sum_{i=1}^{s} (u_i)_{d_0} f_i + \sum_{i=1}^{s} (u_i)_{d_0} \epsilon_i \right) \in \mathrm{LT}(\langle \mathbf{g} \rangle).$$

We conclude that $\langle \mathrm{LT}(\langle \mathbf{f} \rangle) \rangle \subseteq \langle \mathrm{LT}(\langle \mathbf{g} \rangle) \rangle$. □

Lemma 100. *Let $\mathbf{f} = (f_1, \ldots, f_s) \in \mathcal{R}^s$ be a regular sequence of d-th degree homogeneous polynomials and let $\mathbf{g} = (f_1 + \epsilon_1, \ldots, f_s + \epsilon_s) \in \mathcal{R}^s$ be a small perturbation of $\mathbf{f}$. For all $u_1, \ldots, u_s \in \mathcal{R}$, we have*

$$\left(\sum_{i=1}^{s} u_i (f_i + \epsilon_i) \right)_{\mathrm{top}} \in \langle \mathbf{f} \rangle \tag{5.2}$$

with respect to any graded monomial ordering $\geq$ on $\mathcal{R}$. Consequently, $\langle \mathrm{LT}(\langle \mathbf{g} \rangle) \rangle \subseteq \langle \mathrm{LT}(\langle \mathbf{f} \rangle) \rangle$.

Proof. We first show how the inclusion $\langle \mathrm{LT}(\langle \mathbf{g} \rangle) \rangle \subseteq \langle \mathrm{LT}(\langle \mathbf{f} \rangle) \rangle$ follows from the first part of the lemma. As in the proof of Lemma 99, it suffices to show the inclusion $\mathrm{LT}(\langle \mathbf{g} \rangle) \setminus \{0\} \subseteq \mathrm{LT}(\langle \mathbf{f} \rangle)$. To this end, let $h = \sum_{i=1}^{s} u_i (f_i + \epsilon_i)$ be an arbitrary non-zero polynomial in $\langle \mathbf{g} \rangle$. By Equation (5.2), $h_{\mathrm{top}} \in \langle \mathbf{f} \rangle$. We conclude that $\mathrm{LT}(h) = \mathrm{LT}\left(h_{\mathrm{top}}\right) \in \mathrm{LT}(\langle \mathbf{f} \rangle)$.

It remains to show that Equation (5.2) holds for all $u_1, \ldots, u_s \in \mathcal{R}$. The proof is by induction on $d^* := \max_{1 \le i \le s} \deg u_i \ge -\infty$. In the base case $d^* = -\infty$, all u_i are zero, so Equation (5.2) holds trivially. For the induction step, let $d^* \ge 0$ and suppose that Equation (5.2) holds for all $u'_1, \ldots, u'_s \in \mathcal{R}$ with $\max_{1 \le i \le s} \deg u'_i < d^*$. We write

$$\begin{aligned}\sum_{i=1}^{s} u_i (f_i + \epsilon_i) &= \sum_{d'=0}^{d^*} \sum_{i=1}^{s} (u_i)_{d'} (f_i + \epsilon_i) \\ &= \underbrace{\sum_{i=1}^{s} (u_i)_{d^*} f_i}_{=:A} + \underbrace{\sum_{i=1}^{s} (u_i)_{d^*} \epsilon_i}_{=:B} + \underbrace{\sum_{d'=0}^{d^*-1} \sum_{i=1}^{s} (u_i)_{d'} f_i}_{=:C} + \underbrace{\sum_{d'=0}^{d^*-1} \sum_{i=1}^{s} (u_i)_{d'} \epsilon_i}_{=:D},\end{aligned}$$

where $(u_i)_{d'}$ denotes the d'-th degree homogeneous part of u_i. Since both $(u_1)_{d^*}, \ldots, (u_s)_{d^*}$ and $f_1, \ldots, f_s$ are homogeneous sequences of d^*-th degree and d-th degree polynomials, respectively, the sum $\sum_{i=1}^{s} (u_i)_{d^*} f_i$ either has degree $d^* + d$ or is zero. If the sum is non-zero, then $\left(\sum_{i=1}^{s} u_i (f_i + \epsilon_i)\right)_{\text{top}} = \left(\sum_{i=1}^{s} (u_i)_{d^*} f_i\right)_{\text{top}}$, because $\sum_{i=1}^{s} (u_i)_{d^*} \epsilon_i + \sum_{d'=0}^{d^*-1} \sum_{i=1}^{s} (u_i)_{d'} (f_i + \epsilon_i)$ has degree strictly less than $d^* + d$. Otherwise, if $\sum_{i=1}^{s} (u_i)_{d^*} f_i = 0$, then $(u_1)_{d^*}, \ldots, (u_s)_{d^*}$ forms a syzygy of the regular sequence. $f_1, \ldots, f_s$. By Lemma 62 there exist polynomials $w_{11}, \ldots, w_{ss} \in \mathcal{R}$ such that $(u_i)_{d^*} = \sum_{j=1}^{s} w_{ij} f_j$. The w_{ij} also satisfy $w_{ij} = -w_{ji}$ and $w_{ii} = 0$ for all $1 \le i, j \le s$. Without loss of generality, we assume that every w_{ij} is either zero or has degree $d^* - d$. We now have

$$B = \sum_{i=1}^{s} (u_i)_{d^*} \epsilon_i = \sum_{i=1}^{s} \sum_{j=1}^{s} w_{ij} f_j \epsilon_i = \sum_{i=1}^{s} v_i f_i,$$

where $v_i = \sum_{j=1}^{s} w_{ji} \epsilon_j$ for all $1 \le i \le s$. Define $u'_i = v_i + \sum_{d'=0}^{d^*-1} (u_i)_{d'}$ for all $1 \le i \le s$. We will show that $\sum_{i=1}^{s} u_i (f_i + \epsilon_i) = \sum_{i=1}^{s} u'_i (f_i + \epsilon_i)$ and that $\deg u'_i < d^*$ for all $1 \le i \le s$. Note that $\sum_{i=1}^{s} u'_i f_i = B + C$. We also have $\sum_{i=1}^{s} u'_i \epsilon_i = S + D$, where $S = \sum_{i=1}^{s} \sum_{j=1}^{s} w_{ji} \epsilon_j \epsilon_i$. However,

$$S = \sum_{1 \le i < j \le s} \left(w_{ij} + w_{ji}\right) \epsilon_i \epsilon_j + \sum_{i=1}^{s} w_{ii} \epsilon_i^2 = 0,$$

since both sums on the right-hand side consist of zero terms. It now follows that $\sum_{i=1}^{s} u'_i (f_i + \epsilon_i) = B + C + D = \sum_{i=1}^{s} u_i (f_i + \epsilon_i)$. (Recall that $A = \sum_{i=1}^{s} (u_i)_{d^*} f_i = 0$.) For all $1 \le i \le s$, we have $\deg v_i \le \max_j \deg w_{ji} \epsilon_j \le (d^* - d) + (d - 1) = d^* - 1$ and $\deg\left(\sum_{d'=0}^{d^*-1} (u_i)_{d'}\right) \le d^* - 1$, so $\deg u'_i \le d^* - 1$. We may therefore use the induction hypothesis to conclude that

$$\left(\sum_{i=1}^{s} u_i (f_i + \epsilon_i)\right)_{\text{top}} = \left(\sum_{i=1}^{s} u'_i (f_i + \epsilon_i)\right)_{\text{top}} \in \langle \mathbf{f} \rangle .$$

□

The following theorem is an immediate consequence of Lemma 99 and Lemma 100.

Theorem 101. *Let $\mathbf{f} = (f_1, \ldots, f_s) \in \mathcal{R}^s$ be a regular sequence of d-th degree homogeneous polynomials and let $\mathbf{g} = (f_1 + \epsilon_1, \ldots, f_s + \epsilon_s) \in \mathcal{R}^s$ be a small perturbation of $\mathbf{f}$. Then $\langle \mathrm{LT}(\langle \mathbf{f} \rangle) \rangle = \langle \mathrm{LT}(\langle \mathbf{g} \rangle) \rangle$ with respect to any graded monomial ordering $\geq$ on $\mathcal{R}$.*

Theorem 102. *Let $\mathbf{f} = (f_1, \ldots, f_s) \in \mathcal{R}^s$ be a regular sequence of d-th degree homogeneous polynomials and let $\mathbf{g} = (f_1 + \epsilon_1, \ldots, f_s + \epsilon_s) \in \mathcal{R}^s$ be a small perturbation of $\mathbf{f}$. Then $\langle \mathbf{f} \rangle$ and $\langle \mathbf{g} \rangle$ have the same ideal degree, which is equal to d^s.*

Proof. Fix an arbitrary graded monomial ordering $\geq$ on $\mathcal{R}$. By Lemma 38, $\mathcal{R}/\langle \mathbf{g} \rangle$ is isomorphic as a K-vector space to $S_1 = \mathrm{span}_K\{\mathbf{x}^\alpha \in \mathcal{R} \mid \mathbf{x}^\alpha \notin \langle \mathrm{LT}(\langle \mathbf{g} \rangle) \rangle\}$. By Theorem 101, this span equals $S_2 = \mathrm{span}_K\{\mathbf{x}^\alpha \in \mathcal{R} \mid \mathbf{x}^\alpha \notin \langle \mathrm{LT}(\langle \mathbf{f} \rangle) \rangle\}$. Applying Lemma 38 again, we see that S_2 is isomorphic as a K-vector space to $\mathcal{R}/\langle \mathbf{f} \rangle$. We now have $\mathcal{R}/\langle \mathbf{g} \rangle \cong S_1 = S_2 \cong \mathcal{R}/\langle \mathbf{f} \rangle$ (as K-vector spaces) and use Corollary 57 to conclude that $\dim_K \mathcal{R}/\langle \mathbf{g} \rangle = \dim_K \mathcal{R}/\langle \mathbf{f} \rangle = d^s$. □

5.4 Compositions of Perturbed Regular Systems

This section combines and summarizes results from Section 5.2 and Section 5.3 to derive results on compositions of small perturbations of regular sequences of d-th degree homogeneous polynomials.

We start with a proposition concerning compositions of small perturbations.

Proposition 103. *Let $\mathbf{f}, \mathbf{f}' \in \mathcal{R}^n$ be regular sequences of d-th and d'-th degree homogeneous polynomials, respectively, and let $\mathbf{g}, \mathbf{g}' \in \mathcal{R}^n$ be small perturbation of $\mathbf{f}$ and $\mathbf{f}'$, respectively. Then, the composition $\mathbf{f}' \circ \mathbf{f}$ is a regular sequence of $d'd$-th degree homogeneous polynomials and $\mathbf{g}' \circ \mathbf{g}$ is a small perturbation of $\mathbf{f}' \circ \mathbf{f}$.*

Proof. We know from Corollary 96 that $\mathbf{f}' \circ \mathbf{f}$ forms a regular sequence of $d'd$-th degree homogeneous polynomials. Applying Proposition 93, we find that the top homogeneous parts of $\mathbf{g}' \circ \mathbf{g}$ are equal to $\mathbf{f}' \circ \mathbf{f}$. This is equivalent to $\mathbf{g}' \circ \mathbf{g}$ being a small perturbation of $\mathbf{f}' \circ \mathbf{f}$. □

The following theorem summarizes the results from this chapter.

Theorem 104. *Let $\mathbf{f}_1, \ldots, \mathbf{f}_k \in \mathcal{R}^n$ be regular sequences such that $\mathbf{f}_i$ consists of d_i-th degree homogeneous polynomials, for all $1 \leq i \leq n$. Let $\mathbf{g}_1, \ldots, \mathbf{g}_k \in \mathcal{R}^n$ be small perturbations of $\mathbf{f}_1, \ldots, \mathbf{f}_k \in \mathcal{R}^n$, respectively. Then the following holds:*

1. *The composition $\mathbf{f}_k \circ \cdots \circ \mathbf{f}_1$ is a regular sequence of $d_1 \cdots d_k$-th degree homogeneous polynomials and $\mathbf{g}_k \circ \cdots \circ \mathbf{g}_1$ is a small perturbation of $\mathbf{f}_k \circ \cdots \circ \mathbf{f}_1$.*
2. *The ideal $\langle \mathbf{f}_k \circ \cdots \circ \mathbf{f}_1 \rangle$ has ideal degree $(d_1 \cdots d_k)^n$.*

3. *The ideal* $\langle \mathbf{g}_k \circ \cdots \circ \mathbf{g}_1 \rangle$ *has ideal degree* $(d_1 \cdots d_k)^n$.

Proof. Property 1 is obtained by repeatedly applying Proposition 103. Property 2 then follows from Corollary 57. Lastly, 3 follows from Theorem 102. □

Theorem 104 can be used to find the ideal degree corresponding to ATRAPOS$[R]$.

Corollary 105. *Let* $\ell > 3$ *be an odd number not divisible by* 3*, let* $p \geq 3$ *be a prime number, and let* $\mathcal{R} = \mathbb{F}_p[x_1, \ldots, x_\ell]$*. As defined in Section 5.1, let* $\mathcal{F}_{j,\text{inh}} \in \mathcal{R}^\ell$ *be the polynomials corresponding to the* j*-th round of* ATRAPOS*. Then, the ideal* $\left\langle \mathcal{F}_{R,\text{inh}} \circ \cdots \circ \mathcal{F}_{1,\text{inh}} \right\rangle$ *corresponding to the first* $R \geq 1$ *rounds of* ATRAPOS *has ideal degree* $2^{\ell R}$*.*

Proof. From our discussion in Section 5.1, we know that every $\mathcal{F}_{j,\text{inh}} \in \mathcal{R}^\ell$ is a small perturbation of the homogeneous system $\mathcal{F}_{\text{hom}} = (f_1, \ldots, f_\ell)$ defined in Section 4.1. By definition, $f_1, \ldots, f_\ell$ are homogeneous polynomials of degree $d = 2$. Additionally, these polynomials form a regular sequence (Lemma 87). By Theorem 104, the ideal $\left\langle \mathcal{F}_{R,\text{inh}} \circ \cdots \circ \mathcal{F}_{1,\text{inh}} \right\rangle$ has ideal degree $\left(\prod_{i=1}^{R} d\right)^\ell = 2^{\ell R}$. □

5.5 Minimal Number of Rounds for ATRAPOS

We conclude this chapter by discussing the minimal number of rounds R required for ATRAPOS$[R]$ to achieve a security of 128 bits against algebraic attacks.

From the discussion in Subsection 3.2.1 we know that the complexity of solving $\mathcal{F}_{\text{inh}}^{(R)}$ for a single solution is dominated by the FGLM step, which requires $\mathcal{C}_{\text{FGLM}} = d_{\mathcal{I}}^{\omega}$ field operations in $\mathbb{F}_p$. The conservative choice $\omega = 2$ and $d_{\mathcal{I}} = 2^{\ell R}$ (Corollary 105) together yield a conservative estimation of $\mathcal{C}_{\text{FGLM}} = 2^{2\ell R}$.

Given an odd $\ell \geq 3$ not divisible by 3, ATRAPOS-SPONGE has a security level against algebraic attacks of

$$\lambda = \log_2\left(\frac{\mathcal{C}_{\text{FGLM}}}{\text{P}}\right) = \log_2\left(2^{2\ell R}\right) = 2\ell R \text{ bits}$$

with respect to field operations in $\mathbb{F}_p$ (addition and multiplication). The success probability P of the attack equals 1. To obtain (at least) 128 bits of security against algebraic attacks, we need $2\ell R \geq 128$. The minimum R satisfying this equation is $R = \left\lceil \frac{128}{2\ell} \right\rceil$. For KYBER ($\ell = 17$), this yields $R = 4$ (resulting in 136 bits of security). For DILITHIUM ($\ell = 7$), we obtain $R = 10$ (resulting in 140 bits of security).

Note that these results do not yet include a security margin to account for possible improvements to the three-step algorithm discussed in Subsection 3.2.1. Also, there may be other attacks for which a higher number of rounds would be required. This is, however, outside the scope of this thesis.

Remark 106. The 128 bits of security against algebraic attacks discussed here is with respect to field operations in $\mathbb{F}_p$, while the 128 bits of security against brute-force attacks (as stated in Section 3.1) is with respect to evaluations of the ATRAPOS permutation. These security claims are therefore not directly comparable. One way to compare both claims is to estimate and compare the complexities (using e.g. "gate equivalents") of optimal circuits that realize the brute-force and algebraic attacks, respectively. We will not pursue this endeavor.

Chapter 6

Experimental Results

In Section 5.5, we estimated that the FGLM step when solving the CICO problem in Definition 70 for R-round ATRAPOS requires $\mathcal{C}_{\text{FGLM}} = 2^{2\ell R}$ field operations in $\mathbf{F}_p$, assuming that $\ell > 3$ is odd and not divisible by 3. Ultimately, this estimated time complexity stems from the *upper bound* $\mathcal{O}\left(nd_{\mathcal{I}}^{\omega}\right)$ for the full FGLM algorithm (Proposition 46), which is assumed to be tight. (We use $\mathcal{I}$ to denote the ideal corresponding to R-round ATRAPOS.) In this chapter we run small-scale experiments to verify the tightness of this bound for a specific instance of the FGLM algorithm (in Magma). We also measure the time spent on the DRL Gröbner basis computation (using F4) and find that it is negligible compared to the time spent on the lexicographic Gröbner basis computation (using FGLM). Since Magma does not implement the modified FGLM algorithm discussed in Section 2.8 (where we only solve for a single solution using an FGLM variant that exploits sparsity), we perform our experiments using the traditional FGLM algorithm.

The experiments were conducted using Magma (V2.28-8) [BCP97] on a computer with an Intel i9-9900K CPU (3.60 GHz) and 64 GiB of RAM. The code is included in Sections A.2 and A.3.

The ATRAPOS permutation defined in Chapter 3 allows either $\ell = 17$ (for KYBER) or $\ell = 7$ (for DILITHIUM), but is also a permutation for any odd $\ell \geq 3$. This permits us to perform small-scale experiments for small values of ℓ and R. The experiments were conducted using an implementation of the ATRAPOS permutation in Magma. The implementation uses the column shifts $r_0 = 0$, $r_1 = 1$, and $r_2 = r_1 + 1 = 2$. The round constants are set to $c_1 = 1$, $c_2 = 2$, etc., and p was set to the prime $p = 3329$ used in KYBER.

Remark 107. We state without proof that the ideals corresponding to a single round of ATRAPOS also have ideal degree 2^{ℓ} if $\ell \geq 3$ is a multiple of three. Moreover, the top homogeneous parts of their "natural" generators form a regular sequence of ℓ quadratic polynomials. For small ℓ, this can be directly verified by computing their Hilbert series.

A consequence is that the complexity estimate $\mathcal{C}_{\text{FGLM}}$ holds for all odd $\ell \geq 3$.

We will therefore also consider $\ell \geq 3$ divisible by three in our experiments.

See also Remark 89.

For several combinations of (odd) $\ell \geq 3$ and $R \geq 1$, the following steps were performed:

1. Calculate the state $\mathbf{a}'$ obtained by running ATRAPOS$[R]$ on the input state $\mathbf{a} = (\mathbf{0}, \mathbf{0}, \mathbf{x})^{\mathsf{T}}$. (That is, the first two rows, corresponding to the inner part of the state, are set to zero. The bottom row, which corresponds to the outer part, consists of variables $x_1, \ldots, x_\ell$.) We denote the outer part (bottom row) of $\mathbf{a}'$ by $\mathbf{g}$.

2. Compute a DRL Gröbner basis for $\langle \mathbf{g} \rangle \subseteq \mathbb{F}_p[x_1, \ldots, x_\ell]$ using the F4 algorithm.

3. Convert the DRL Gröbner basis to a lexicographic Gröbner basis using the FGLM algorithm.

The steps above were repeated five times for each combination of ℓ and R to account for random fluctuations in running times.

Table 6.1 lists the mean running times for the F4 (step 2) and FGLM (step 3) steps. In all cases, the total running time is dominated by the running time of FGLM. This observation justifies our choice of analyzing only the time complexity of the FGLM step.

	$R = 1$	$R = 2$	$R = 3$	$R = 4$	$R = 5$
$\ell = 3$	0.00s (0.00s)	0.01s (0.00s)	0.13s (0.00s)	50s (0.00s)	7h31m (0.00s)
$\ell = 5$	0.00s (0.01s)	2.40s (0.26s)			
$\ell = 7$	0.01s (0.01s)				
$\ell = 9$	0.15s (0.00s)				
$\ell = 11$	6.35s (0.09s)				
$\ell = 13$	6m36s (0.77s)				

Table 6.1: Running times for F4 (gray, between parenthesis) and FGLM for small (odd) $\ell \geq 3$ and $R \geq 1$. The numbers listed represent the mean over five identical experiments. Empty cells correspond to experiments that resulted in an out-of-memory error. Experiments with $R > 6$ or $\ell \geq 15$ resulted in out-of-memory errors as well.

The experimental results also help us reason about the actual time complexity of the FGLM step (in Magma). Assuming that the FGLM implementation in Magma requires roughly $\ell d_{\mathcal{I}}^{\omega} = \ell \cdot 2^{\omega \ell R}$ field operations in $\mathbb{F}_p$, it is reasonable to estimate the time spent on FGLM by $T_{\text{FGLM}} = c\ell \cdot 2^{\omega \ell R}$ for some constant $c > 0$. (The actual value of c is very implementation specific, as it depends on the combination of hardware and software used for the FGLM step.) This estimate depends on both ℓ and R. Dividing T_{FGLM} by ℓ, we obtain $T_{\text{FGLM}}/\ell = c \cdot 2^{\omega \ell R}$, which only depends on the product ℓR. Therefore, if we denote the actual running times for the FGLM step by T, we expect the points $(\ell R, T/\ell)$ to be close

to the curve $c \cdot 2^{\omega \ell R}$ for some $c > 0$ and $2 \leq \omega \leq 3$. Equivalently, for every point $(\ell R, T/\ell)$, we expect $\log_2(T/\ell) \approx \omega \ell R + b$, where $b = \log_2 c$.

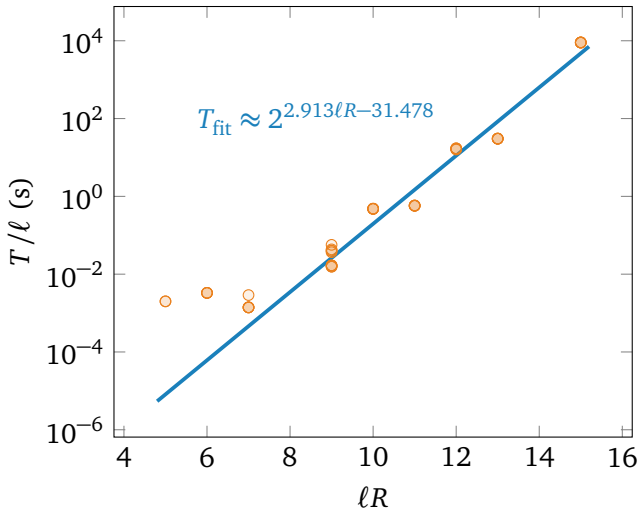

Figure 6.1: Scatter plot of measured FGLM running times T, divided by ℓ, as a function of ℓR. The blue line denotes the fitted line.

Figure 6.1 shows T/ℓ as a function of ℓR, together with a fitted curve $T_{\text{fit}} = 2^{\hat{\omega}\ell R + \hat{b}}$. The fitted curve was obtained by performing a polynomial fit of the linear polynomial $\omega \ell R + b$ on the points $\big(\ell R, \log_2(T/\ell)\big)$. For small ℓR, the results may be skewed, since the memory used by the FGLM step fits completely within the L3 cache. For this reason, only points with $\ell R \geq 9$ were considered for this fit. These points correspond to experiments where the memory used by Magma exceeded the L3 cache (16 MB).

The results in Figure 6.1 are surprisingly close to our expectations. For example, $\hat{\omega} \approx 2.913$, which is not too far off from the value $\omega = \log_2 7 \approx 2.807$ we would expect for Strassen multiplication. Note that for small ℓR, a larger part of the data fits in the CPU caches, which would speed up the algorithm. If we look at the observed running times, however, it seems that they are *higher* than the predicted running times when ℓR is small. A possible explanation is that the FGLM implementation in Magma involves a constant time component (e.g. a setup phase): when ℓR is small, the total running time is low, so the constant time component is relatively high, compared to the total running time.

As a concluding remark, we want to stress that these experimental results by themselves do not constitute a proof that solving the CICO problem is hard. They only show that the complexity of solving (small instances of) the CICO problem using the FGLM implementation in this version of Magma agrees with our expectations. Nevertheless, these experimental results give some assurance that the theoretical running time $T_{\text{FGLM}} = c\ell \cdot 2^{\omega \ell R}$ for the FGLM step when solving the CICO problem is reasonable. By extension, the results give some

assurance that the time complexity $\mathcal{C}_{\text{FGLM}} = 2^{2\ell R}$ for the optimized FGLM variant discussed in Section 2.8 is reasonable.

Chapter 7

Related Work

Since the ATRAPOS-SPONGE specification has not been published yet, no prior (publicly available) research on the security of ATRAPOS-SPONGE against algebraic attacks exists. There is, however, a vast amount of published research on algebraic attacks against other cryptographic algorithms.

The application of algebraic techniques in cryptanalysis is not new. For instance, the authors of a 2002 paper [CP02] modeled the AES block cipher using an overdetermined system of (quadratic) polynomial equations. Subsequently, they described a method to recover the secret key using an attack they call the "eXtended Sparse Linearization" (XSL) attack. The authors conjecture that the security of AES does not grow exponentially with the number of rounds and that their proposed attack seems to break AES-256. However, this claim has been disputed by [CL05].

A few years later, in 2006, [BPW06] presented a key recovery attack for AES using Gröbner basis techniques. The authors give a polynomial modeling of AES such that the polynomials already form a Gröbner basis with respect to the DRL ordering. The corresponding ideal degree is $254^{200} \approx 2^{1598}$, which makes obtaining a Gröbner basis with respect to the lexicographic ordering using FGLM prohibitively expensive. This technique does therefore not break AES.

More recently, algebraic cryptanalysis using Gröbner bases has received new attention in the context of "arithmetization-oriented" cryptography. Arithmetization-oriented primitives are used in e.g. zero-knowledge protocols and operate on elements in finite fields $\mathbb{F}_p$ for large p. For these primitives, Gröbner basis techniques are often assumed to be the most efficient attack techniques [KLR24].

As an example, [KLR24] investigates the security of Anemoi, an arithmetization-oriented permutation-based hash function, against Gröbner basis attacks. The authors define two polynomial modelings for Anemoi. For both of these modelings, they estimate the complexity of computing a DRL Gröbner basis as well as the complexity of the basis conversion step using FGLM, by extrapolating from small-scale experiments. In [CR25], an alternative modeling for Anemoi

is presented. The authors derive complexities for both the DRL Gröbner basis computation and the basis conversion step (based on the ideal degree). The ideal degree is determined by counting the number of monomials in the quotient space $\mathcal{R}/\mathcal{I}$ corresponding to their polynomial modeling of Anemoi. It is unclear whether the techniques used in [CR25] can be adapted to determine the ideal degree of our modeling of ATRAPOS.

Chapter 8

Conclusions

In this thesis, we described ATRAPOS-SPONGE and analyzed a specific CICO problem related to single-block preimage resistance. As we have seen, solving this CICO problem amounts to solving a system of polynomial equations. We estimated the complexity of this problem using the ideal degree corresponding to the ATRAPOS permutation in the sponge construction.

We showed that the homogeneous ideal corresponding to a single round of ATRAPOS is generated by a regular sequence of polynomials and used this fact to show that the homogeneous ideal corresponding to R-round ATRAPOS is generated by a regular sequence as well. In both cases, properties of regular sequences were used to extend the results to the inhomogeneous case.

Ultimately, we proved that the complexity of the CICO problem for R-round ATRAPOS is $2^{\omega \ell R}$ and used this to derive the minimum number of rounds required to obtain 128 bits of security (with respect to our CICO problem) when ATRAPOS-SPONGE is used in KYBER and DILITHIUM.

In Chapter 6 we confirmed the theoretical results obtained in earlier chapters.

8.1 Future Research

We outline three areas of interest for future research.

First, the research in this thesis is confined to instantiations of ATRAPOS where only the bottom row of the two-dimensional state is used for the outer part of the state ("the $r = \ell$ case"). The current ATRAPOS specification also allows the bottom *two* rows to be used for the outer part ("the $r = 2\ell$ case"). In the latter case, the polynomials $g_1, \ldots, g_{2\ell}$ that describe the outer part of the state after the ATRAPOS permutation (cf. Section 3.2) no longer all have the same total degree. Many of our results depend on all g_i having the same total degree and extending the results to the $r = 2\ell$ case seems non-trivial. Future research is needed to determine the security properties of ATRAPOS in the latter case, possibly purely experimentally.

Second, the polynomial systems we analyzed are "critically determined" in the sense that there are as many equations as there are unknowns. (The systems are neither overdetermined nor underdetermined.) An adversary can force the system to be overdetermined by simply guessing one or more unknowns and then solving the system for the remaining variables. In the extreme case, the attacker guesses every unknown, entirely bypassing the need for the F4/F5 and FGLM algorithms. Anecdotally, guessing unknowns does not reduce the complexity of solving the CICO problem below 128 bits, but more research is needed to properly verify this.

Finally, the CICO problem in Definition 70 is strongly related to the preimage resistance of ATRAPOS-SPONGE. Besides preimage resistance, second-preimage resistance and collision resistance are also important properties for an extendable-output function (XOF). Investigating the second-preimage resistance and collision resistance of ATRAPOS-SPONGE, for example by formulating these properties in terms of polynomial systems, would be an interesting direction for future research.

Bibliography

[Alb+19] Martin R. Albrecht et al. “Algebraic Cryptanalysis of STARK-Friendly Designs: Application to MARVELlous and MiMC”. In: *Advances in Cryptology – ASIACRYPT 2019*. Ed. by Steven D. Galbraith and Shiho Moriai. Springer International Publishing, 2019, pp. 371–397. ISBN: 978-3-030-34618-8. DOI: 10.1007/978-3-030-34618-8_13.

[Ava+21] Roberto Avanzi et al. *CRYSTALS-Kyber (version 3.02) – Submission to round 3 of the NIST post-quantum project*. 2021. URL: https://pq-crystals.org/kyber/data/kyber-specification-round3-20210804.pdf.

[BCP97] Wieb Bosma, John Cannon, and Catherine Playoust. “The Magma Algebra System I: The User Language”. In: *Journal of Symbolic Computation* 24.3 (1997), pp. 235–265. ISSN: 0747-7171. DOI: 10.1006/jsco.1996.0125.

[BDPV11a] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. *Cryptographic sponge functions*. 2011. URL: https://keccak.team/files/CSF-0.1.pdf.

[BDPV11b] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. *The Keccak reference*. 2011. URL: https://keccak.team/files/Keccak-reference-3.0.pdf.

[BPW06] Johannes Buchmann, Andrei Pyshkin, and Ralf-Philipp Weinmann. “A Zero-Dimensional Gröbner Basis for AES-128”. In: *Fast Software Encryption*. Ed. by Matthew Robshaw. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 78–88. ISBN: 978-3-540-36598-3. DOI: 10.1007/11799313_6.

[CL05] Carlos Cid and Gaëtan Leurent. “An Analysis of the XSL Algorithm”. In: *Advances in Cryptology - ASIACRYPT 2005*. Ed. by Bimal Roy. Berlin, Heidelberg: Springer Berlin Heidelberg, 2005, pp. 333–352. ISBN: 978-3-540-32267-2. DOI: 10.1007/11593447_18.

[CLO15] David A. Cox, John Little, and Donal O'Shea. *Ideals, Varieties, and Algorithms: An Introduction to Computational Algebraic Geometry and Commutative Algebra*. 4th ed. Springer International Publishing, 2015. ISBN: 978-3-319-16721-3. DOI: 10.1007/978-3-319-16721-3.

[CP02] Nicolas T. Courtois and Josef Pieprzyk. "Cryptanalysis of Block Ciphers with Overdefined Systems of Equations". In: *Advances in Cryptology — ASIACRYPT 2002*. Ed. by Yuliang Zheng. Berlin, Heidelberg: Springer Berlin Heidelberg, 2002, pp. 267–287. ISBN: 978-3-540-36178-7. DOI: 10.1007/3-540-36178-2_17.

[CR25] Luca Campa and Arnab Roy. "Gröbner Basis Cryptanalysis of Anemoi". In: *Advances in Cryptology – EUROCRYPT 2025*. Ed. by Serge Fehr and Pierre-Alain Fouque. Cham: Springer Nature Switzerland, 2025, pp. 303–332. ISBN: 978-3-031-91107-1. DOI: 10.1007/978-3-031-91107-1_11.

[DF04] David S. Dummit and Richard M. Foote. *Abstract Algebra*. 3rd ed. Wiley, 2004. ISBN: 978-0-471-43334-7.

[DMMØ25] Joan Daemen, Silvia Mella, Konstantina Miteloudi, and Morten Øygarden. Private correspondence. 2025.

[Duc+21] Léo Ducas et al. *CRYSTALS-Dilithium – Algorithm Specifications and Supporting Documentation (Version 3.1)*. 2021. URL: https://pq-crystals.org/dilithium/data/dilithium-specification-round3-20210208.pdf.

[Eis95] David Eisenbud. *Commutative Algebra: with a View Toward Algebraic Geometry*. Springer New York, 1995. ISBN: 978-1-4612-5350-1. DOI: 10.1007/978-1-4612-5350-1.

[Fau02] Jean-Charles Faugère. "A new efficient algorithm for computing Gröbner bases without reduction to zero (F5)". In: *Proceedings of the 2002 International Symposium on Symbolic and Algebraic Computation*. ISSAC '02. Lille, France: Association for Computing Machinery, 2002, pp. 75–83. ISBN: 1581134843. DOI: 10.1145/780506.780516.

[Fau99] Jean-Charles Faugère. "A new efficient algorithm for computing Gröbner bases (F4)". In: *Journal of Pure and Applied Algebra* 139.1 (1999), pp. 61–88. ISSN: 0022-4049. DOI: 10.1016/S0022-4049(99)00005-5.

[FGHR14] Jean-Charles Faugère, Pierrick Gaudry, Louise Huot, and Guénaël Renault. "Sub-cubic change of ordering for Gröbner basis: a probabilistic approach". In: *Proceedings of the 39th International Symposium on Symbolic and Algebraic Computation*. ISSAC '14. Kobe,

Japan: Association for Computing Machinery, 2014, pp. 170–177. ISBN: 9781450325011. DOI: 10.1145/2608628.2608669.

[FGLM93] J.C. Faugère, P. Gianni, D. Lazard, and T. Mora. "Efficient Computation of Zero-dimensional Gröbner Bases by Change of Ordering". In: *Journal of Symbolic Computation* 16.4 (1993), pp. 329–344. ISSN: 0747-7171. DOI: 10.1006/jsco.1993.1051.

[FM11] Jean-Charles Faugère and Chenqi Mou. "Fast algorithm for change of ordering of zero-dimensional Gröbner bases with sparse multiplication matrices". In: *Proceedings of the 36th International Symposium on Symbolic and Algebraic Computation*. ISSAC '11. San Jose, California, USA: Association for Computing Machinery, 2011, pp. 115–122. ISBN: 9781450306751. DOI: 10.1145/1993886.1993908.

[Kan+] Matthias J. Kannwischer et al. *PQM4: Post-quantum crypto library for the ARM Cortex-M4*. https://github.com/mupq/pqm4.

[KL20] Jonathan Katz and Yehuda Lindell. *Introduction to Modern Cryptography*. 3rd ed. Chapman and Hall/CRC, 2020. ISBN: 9781351133036. DOI: 10.1201/9781351133036.

[KLR24] Katharina Koschatko, Reinhard Lüftenegger, and Christian Rechberger. "Exploring the Six Worlds of Gröbner Basis Cryptanalysis: Application to Anemoi". In: *IACR Transactions on Symmetric Cryptology* 2024.4 (Dec. 2024), pp. 138–190. DOI: 10.46586/tosc.v2024.i4.138-190.

[KR05] Martin Kreuzer and Lorenzo Robbiano. *Computational Commutative Algebra 2*. Springer Berlin Heidelberg, 2005. ISBN: 978-3-540-28296-9. DOI: 10.1007/3-540-28296-3.

[Nat15] National Institute of Standards and Technology. *SHA-3 Standard: Permutation-Based Hash and Extendable-Output Functions*. Tech. rep. Federal Information Processing Standards Publications (FIPS) 202. Washington, D.C.: U.S. Department of Commerce, 2015. DOI: 10.6028/NIST.FIPS.202.

[Nat16] National Institute of Standards and Technology. *Submission Requirements and Evaluation Criteria for the Post-Quantum Cryptography Standardization Process*. 2016. URL: https://csrc.nist.gov/CSRC/media/Projects/Post-Quantum-Cryptography/documents/call-for-proposals-final-dec-2016.pdf.

[Nat24a] National Institute of Standards and Technology. *Module-Lattice-Based Digital Signature Standard*. Tech. rep. Federal Information Processing Standards Publications (FIPS) 204. Washington, D.C.: U.S. Department of Commerce, 2024. DOI: 10.6028/NIST.FIPS.204.

[Nat24b] National Institute of Standards and Technology. *Module-Lattice-Based Key-Encapsulation Mechanism Standard*. Tech. rep. Federal Information Processing Standards Publications (FIPS) 203. Washington, D.C.: U.S. Department of Commerce, 2024. DOI: 10.6028/NIST.FIPS.203.

[Sho94] Peter W. Shor. "Algorithms for quantum computation: discrete logarithms and factoring". In: *Proceedings 35th Annual Symposium on Foundations of Computer Science*. 1994, pp. 124–134. DOI: 10.1109/SFCS.1994.365700.

[Sta78] Richard P. Stanley. "Hilbert functions of graded algebras". In: *Advances in Mathematics* 28.1 (1978), pp. 57–83. ISSN: 0001-8708. DOI: 10.1016/0001-8708(78)90045-2.

[Str69] Volker Strassen. "Gaussian Elimination is not Optimal". In: *Numerische Mathematik* 13.4 (Aug. 1969), pp. 354–356. DOI: 10.1007/BF02165411.

[WXXZ24] Virginia Vassilevska Williams, Yinzhan Xu, Zixuan Xu, and Renfei Zhou. "New Bounds for Matrix Multiplication: from Alpha to Omega". In: *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2024, pp. 3792–3835. DOI: 10.1137/1.9781611977912.134.

Appendix A

Code

A.1 Hilbert Series Computation (`hilbert_series.py`)

The following Python file computes the Hilbert series of $\mathcal{R}/\mathcal{I}_{\text{hom}}$ for small odd $\ell \geq 3$ (see Section 4.2).

```python
from sage.all import GF, PolynomialRing, ideal

K = GF(3329)

for ell in range(3, 30, 2):
    R = PolynomialRing(K, ell, 'x', order='degrevlex')
    I = ideal([
        R.gen(i) * (R.gen(i) + R.gen((i + 1) % ell))
        for i in range(0, ell)
    ])
    print(f'{ell:02d}: {I.hilbert_series()}')
```

A.2 ATRAPOS Implementation (`atrapos.mag`)

The following Magma file implements the ATRAPOS permutation as specified in Subsection 3.1.1 and Chapter 6.

```
// wrapping for 1-based indices
WrapIndex := func<i, n | ((i - 1) mod n) + 1>;
WrapX := func<x, ell | WrapIndex(x, ell)>;
WrapY := func<y | WrapIndex(y, 3)>;

// computes the state S having the same dimensions as `state` such that
// `s_{xy} = f(state, x, y)`
MapState := function(state, f);
  return Matrix(
    [[f(state, x, y) : x in [1..Ncols(state)]] : y in [1..Nrows(state)]]
  );
end function;
```

```
Theta := function(state);
  ell := NumberOfColumns(state);
  return MapState(
    state,
    func<state, x, y |
      state[y, x]
      + state[WrapY(y + 1), WrapX(x + 1, ell)]
      + state[WrapY(y + 4), WrapX(x + 4, ell)]
      + state[WrapY(y + 5), WrapX(x + 5, ell)]
    >);
end function;

Rho := function(state, ry);
  ell := NumberOfColumns(state);
  return MapState(
    state,
    func<state, x, y |
      state[y, WrapX(x + ry[y], ell)]
    >);
end function;

Iota := function(state, c);
  return MapState(
    state,
    func<state, x, y |
      (x eq 1 and y eq 1) select state[y, x] + c else state[y, x]
    >);
end function;

Gamma := function(state);
  return MapState(
    state,
    func<state, x, y |
      (y eq 1) select state[y, x] + state[2, x] * state[3, x] else state[y, x]
    >);
end function;

// a single round of Atrapos with column shifts ry and round constant c
Atrapos := func<state, ry, c | Gamma(Iota(Rho(Theta(state), ry), c))>;
```

A.3 Experiments (`experiments.mag`)

The following Magma file performs the experiments described in Chapter 6.

```
SetVerbose("Groebner", 0);
SetVerbose("FGLM", 0);
SetNthreads(1);

load "./atrapos.mag";

// column shifts
ry := AssociativeArray(Integers());
ry[1] := 0;
```

```
ry[2] := 1;
ry[3] := ry[2] + 1;

// round constants
cj := [1..10];

// Limit the number of rounds (depending on ell) to avoid out-of-memory
// exceptions. The default maximum is 1.
maxRounds := AssociativeArray(Integers());
maxRounds[3] := 5;
maxRounds[5] := 2;

// ground field
K := GF(3329);

// how often to repeat every experiment
repetitions := 5;

ells := [3..13 by 2];

// Run and benchmark a single round of Atrapos. Returns the modified state.
BenchmarkRound := function(state, ell, round, ry, c, R);
  ResetMaximumMemoryUsage();
  start := Cputime();
  state := Atrapos(state, ry, c);
  outerPart := RowSequence(state)[1];
  stop := Cputime();
  maxMem := GetMaximumMemoryUsage();

  printf "[ell=%o, R=%o] Computed state : time=%.4o, max_mem=%o\n",
    ell, round, (stop - start), maxMem;

  ResetMaximumMemoryUsage();
  start := Cputime();
  I_drl := ideal<R | outerPart>;
  GB_drl := GroebnerBasis(I_drl);
  stop := Cputime();
  maxMem := GetMaximumMemoryUsage();
  printf "[ell=%o, R=%o] Computed DRL GB: time=%.4o, max_mem=%o\n",
    ell, round, (stop - start), maxMem;

  ResetMaximumMemoryUsage();
  start := Cputime();
  I_lex := ChangeOrder(I_drl, "lex");
  GB_lex := GroebnerBasis(I_lex: Al := "FGLM");
  stop := Cputime();
  maxMem := GetMaximumMemoryUsage();
  printf "[ell=%o, R=%o] Computed LEX GB: time=%.4o, max_mem=%o\n",
    ell, round, (stop - start), maxMem;

  printf "\n";

  return state;
end function;
```

```
for ell in ells do
  R := PolynomialRing(K, ell, "grevlex");

  for repetition in [1..repetitions] do
    printf "ell=%o, repetition=%o\n", ell, repetition;

    state := Matrix([
      [R.i : i in [1..ell]],
      [0   : i in [1..ell]],
      [0   : i in [1..ell]]
    ]);

    max := IsDefined(maxRounds, ell) select maxRounds[ell] else 1;
    for round in [1..max] do
      c := cj[round];
      state := BenchmarkRound(state, ell, round, ry, c, R);
    end for;
  end for;
end for;
```