

\mathcal{P} -schemes and Deterministic Polynomial Factoring over Finite Fields

Thesis by
Zeyu Guo

In Partial Fulfillment of the Requirements for the
degree of
Doctor of Philosophy

The logo for the California Institute of Technology (Caltech), featuring the word "Caltech" in a bold, orange, sans-serif font.

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2017
Defended May 22

© 2017

Zeyu Guo

ORCID: 0000-0001-7893-4346

All rights reserved

ACKNOWLEDGEMENTS

I would like to thank my advisor, Chris Umans, for his patient guidance and continual encouragement throughout my graduate studies. His passion, enthusiasm, and dedication for research are truly inspiring. I am very fortunate to have such a great teacher as my advisor.

I am indebted to Michael Aschbacher, who taught me a graduate algebra course, and Matthias Flach, who taught me a course on algebraic number theory. The knowledge I learned from their courses is crucial for the work presented in this thesis.

I also want to thank Leonard Schulman, Anand Kumar Narayanan, Manuel Arora, and Jenish Mehta for many helpful conversations. In particular, I am grateful to Manuel Arora for explaining to me his work on m -schemes.

Finally, I want to thank my family and friends for their continual support and encouragement.

ABSTRACT

We introduce a family of mathematical objects called \mathcal{P} -schemes, where \mathcal{P} is a poset of subgroups of a finite group G . A \mathcal{P} -scheme is a collection of partitions of the right coset spaces $H \backslash G$, indexed by $H \in \mathcal{P}$, that satisfies a list of axioms. These objects generalize the classical notion of association schemes [BI84] as well as the notion of m -schemes [IKS09].

Based on \mathcal{P} -schemes, we develop a unifying framework for the problem of deterministic factoring of univariate polynomials over finite field under the generalized Riemann hypothesis (GRH). More specifically, our results include the following:

- We show an equivalence between m -scheme as introduced in [IKS09] and \mathcal{P} -schemes in the special setting that G is an multiply transitive permutation group and \mathcal{P} is a poset of pointwise stabilizers, and therefore realize the theory of m -schemes as part of the richer theory of \mathcal{P} -schemes.
- We give a generic deterministic algorithm that computes the factorization of the input polynomial $f(X) \in \mathbb{F}_q[X]$ given a “lifted polynomial” $\tilde{f}(X)$ of $f(X)$ and a collection \mathcal{F} of “effectively constructible” subfields of the splitting field of $\tilde{f}(X)$ over a certain base field. It is routine to compute $\tilde{f}(X)$ from $f(X)$ by lifting the coefficients of $f(X)$ to a number ring. The algorithm then successfully factorizes $f(X)$ under GRH in time polynomial in the size of $\tilde{f}(X)$ and \mathcal{F} , provided that a certain condition concerning \mathcal{P} -schemes is satisfied, for \mathcal{P} being the poset of subgroups of the Galois group G of $\tilde{f}(X)$ defined by \mathcal{F} via the Galois correspondence. By considering various choices of G , \mathcal{P} and verifying the condition, we are able to derive the main results of known (GRH-based) deterministic factoring algorithms [Hua91a; Hua91b; Rón88; Rón92; Evd92; Evd94; IKS09] from our generic algorithm in a uniform way.
- We investigate the *schemes conjecture* in [IKS09] and formulate analogous conjectures associated with various families of permutation groups, each of which has applications on deterministic polynomial factoring. Using a technique called induction of \mathcal{P} -schemes, we establish reductions among these conjectures and show that they form a hierarchy of relaxations of the original schemes conjecture.

- We connect the complexity of deterministic polynomial factoring with the complexity of the Galois group G of $\tilde{f}(X)$. Specifically, using techniques from permutation group theory, we obtain a (GRH-based) deterministic factoring algorithm whose running time is bounded in terms of the noncyclic composition factors of G . In particular, this algorithm runs in polynomial time if G is in Γ_k for some $k = 2^{O(\sqrt{\log n})}$, where Γ_k denotes the family of finite groups whose noncyclic composition factors are all isomorphic to subgroups of the symmetric group of degree k . Previously, polynomial-time algorithms for Γ_k were known only for bounded k .
- We discuss various aspects of the theory of \mathcal{P} -schemes, including techniques of constructing new \mathcal{P} -schemes from old ones, \mathcal{P} -schemes for symmetric groups and linear groups, orbit \mathcal{P} -schemes, etc. For the closely related theory of m -schemes, we provide explicit constructions of strongly antisymmetric homogeneous m -schemes for $m \leq 3$. We also show that all antisymmetric homogeneous orbit 3-schemes have a matching for $m \geq 3$, improving a result in [IKS09] that confirms the same statement for $m \geq 4$.

In summary, our framework reduces the algorithmic problem of deterministic polynomial factoring over finite fields to a combinatorial problem concerning \mathcal{P} -schemes, allowing us to not only recover most of the known results but also discover new ones. We believe progress in understanding \mathcal{P} -schemes associated with various families of permutation groups will shed some light on the ultimate goal of solving deterministic polynomial factoring over finite fields in polynomial time.

TABLE OF CONTENTS

Acknowledgements	iii
Abstract	iv
Table of Contents	vi
Chapter I: Introduction	1
1.1 Previous work	1
1.2 Main results	3
1.3 Outline of the thesis.	8
1.4 Notations and preliminaries	10
Chapter II: \mathcal{P} -schemes	14
2.1 Preliminaries	15
2.2 \mathcal{P} -schemes	18
2.3 m -schemes	22
2.4 Orbit \mathcal{P} -schemes and m -schemes	32
2.5 Strongly antisymmetric homogeneous m -schemes for $m \leq 3$	38
Chapter III: The \mathcal{P} -scheme algorithm	44
3.1 Preliminaries	49
3.2 Algorithmic preliminaries	55
3.3 Reduction to computing an idempotent decomposition of $\bar{\mathcal{O}}_F$	58
3.4 Main algorithm	61
3.5 Compatibility and invariance test	65
3.6 Regularity test	67
3.7 Strong antisymmetry test	70
3.8 Constructing a collection of number fields	74
3.9 Putting it together	75
Chapter IV: Constructing number fields	79
4.1 Preliminaries	80
4.2 Adjoining roots of polynomials	86
4.3 Reduction to primitive group actions	89
4.4 Other techniques of constructing number fields	94
Chapter V: The generalized \mathcal{P} -scheme algorithm	100
5.1 Preliminaries	106
5.2 Algorithmic preliminaries	114
5.3 Reduction to computing an idempotent decomposition of R_F	117
5.4 Producing a \mathcal{P} -scheme of double cosets \mathcal{C}	123
5.5 Testing local constantness of ramification indices	126
5.6 Testing local constantness of inertia degrees	128
5.7 A \mathcal{P} -collection $\tilde{\mathcal{C}}$ induced from \mathcal{C} and auxiliary elements	129
5.8 $(\mathcal{C}, \mathcal{D})$ -separated \mathcal{P} -collections	131
5.9 Producing an ordinary \mathcal{P} -scheme	136

5.10 Putting it together	139
Chapter VI: Constructing new \mathcal{P} -schemes from old ones	146
6.1 Basic operations on \mathcal{P} -schemes	148
6.2 Induction of \mathcal{P} -schemes	153
6.3 Schemes conjectures	159
6.4 Extension to the closure of a subgroup system	162
6.5 Restricting to a subset	168
6.6 Primitivity of homogeneous m -schemes	171
6.7 Direct products and wreath products	175
Chapter VII: Symmetric groups and linear groups	181
7.1 The natural action of a symmetric group	183
7.2 Self-reduction of discreteness	186
7.3 The actions of symmetric groups on k -subsets or partitions	189
7.4 The natural actions of linear groups	194
Chapter VIII: Groups with restricted noncyclic composition factors	201
8.1 Proof of the main theorem	204
8.2 The O’Nan-Scott theorem for finite primitive permutation groups	208
8.3 Almost simple type	212
8.4 Affine type	215
8.5 Diagonal type	223
8.6 Product type and twisted wreath type	226
8.7 Future research	230
Appendix A: A unifying definition of \mathcal{P} -schemes	232
Appendix B: Proofs omitted from Chapter III	236
Appendix C: Proofs omitted from Chapter V	244
Appendix D: List of algorithms	267
Bibliography	269
List of Notations	276

Chapter 1

INTRODUCTION

We are interested in the problem of *deterministic univariate polynomial factoring* over finite fields: given a univariate polynomial f of degree $n \in \mathbb{N}^+$ over a finite field \mathbb{F}_q , our goal is to *deterministically* compute a factorization of f over \mathbb{F}_q

$$f(X) = c \cdot \prod_{i=1}^k f_i(X),$$

where $c \in \mathbb{F}_q$ is the leading coefficient of f and each factor f_i is *irreducible* over \mathbb{F}_q . This is called the *complete factorization* of f over \mathbb{F}_q . It is unique up to the order of the factors f_i , since $\mathbb{F}_q[X]$ is a *unique factorization domain*. In addition, we are also interested in the more moderate goal of deterministically computing a *proper factorization* of f , i.e., factoring f into more than one factors where each factor is allowed to be reducible.

1.1 Previous work

Univariate polynomial factoring over finite fields has been extensively studied over the years as one of the most fundamental problems in computer algebra and a common subroutine of many algorithms in coding theory, cryptography, computational number theory, etc. We review the previous work on this problem, with emphasis on *deterministic* factoring algorithms. For a detailed survey, see [GP01].

A truly polynomial-time factoring algorithm is required to factorize a degree- n polynomial $f(X) \in \mathbb{F}_q[X]$ in time $(n \log q)^{O(1)}$, since it takes $O(n \log q)$ bits to describe f . If randomness is allowed, such algorithms are well known: Berlekamp [Ber70] described a randomized algorithm that (completely) factorizes a univariate polynomial over \mathbb{F}_q in polynomial time. The same paper also gave a deterministic reduction from the problem of factoring f to the problem of finding the roots of certain other polynomials that split into n linear factors over \mathbb{F}_p , where $p = \text{char}(\mathbb{F}_q)$. More efficient randomized algorithms were discovered since then [CZ81; GS92; KS98; Uma08; KU11]. The current best known running time has the exponent $3/2$ in n , as achieved by [KU11] based on the technique of *fast modular composition*.

On the other hand, despite much effort, factoring polynomials over finite fields in *deterministic* polynomial time remains a long-standing open problem. Berlekamp

[Ber67] gave the first deterministic algorithm for the general problem, whose running time is polynomial in n and q (instead of n and $\log q$). His aforementioned paper [Ber70] gave a deterministic algorithm that runs in time polynomial in n , $\log q$ and $p = \text{char}(\mathbb{F}_q)$. Deterministic algorithms with running time $(n \log q)^{O(1)} p^{1/2}$ were given in [Sho90; BKS15]. Unfortunately, the $p^{1/2}$ -dependence on the characteristic p of the field remains the best known for *unconditional* deterministic factoring algorithms, even if we only consider quadratic polynomials. Faster algorithms are known when $p - 1$ is assumed to be a *smooth* number [Gat87; Rón89; Sho91]. In addition, there are deterministic algorithms for special polynomials based on the theory of elliptic curves or abelian varieties [Sch85; Pil90]. Finally, the paper [Iva+12] also unconditionally obtained some positive results on deterministic polynomial factoring in certain special cases.

A lot more is known if one accepts the generalized Riemann hypothesis (GRH): a deterministic polynomial-time algorithm that factorizes polynomials of the form $X^n - a \in \mathbb{F}_p[X]$ under GRH was given in [AMM77]. Several GRH-based deterministic algorithms were proposed since then. These algorithms factorize a polynomial $f(X) \in \mathbb{F}_p[X]$ using the auxiliary information of a *lifted polynomial*, i.e., a polynomial $\tilde{f}(X) \in \mathbb{Z}[X]$ satisfying $\tilde{f}(X) \bmod p = f(X)$. Huang [Hua91a; Hua91b] proved that a polynomial $f(X) \in \mathbb{F}_p[X]$ can be deterministically factorized in polynomial time under GRH provided that the Galois group of the lifted polynomial is abelian.¹ This was generalized in [Evd92] to the case of solvable Galois groups. For a general Galois group G , the work [Rón92] provided a deterministic algorithm that runs in time polynomial in $|G|$ and the size of the input under GRH. In general, however, the cardinality of G may be as large as $n!$, as attained by the symmetric group of degree n . Thus the algorithm in [Rón92] may take exponential time.

In a different approach, Rónyai [Rón88] showed that a polynomial $f(X) \in \mathbb{F}_q[X]$ of degree n can be factorized deterministically in time $(n^n \log q)^{O(1)}$ under GRH. The algorithm proceeds by manipulating tensor powers of the ring $\mathbb{F}_q[X]/(f(X))$, and does not need a lifted polynomial of f . Building on Rónyai's work, Evdokimov [Evd94] showed that the problem can be solved in quasipolynomial time by presenting a deterministic $(n^{\log n} \log q)^{O(1)}$ -time algorithm under GRH. Evdokimov's algorithm remains the best known result on GRH-based deterministic polynomial factoring, although the $O(\log n)$ exponent of the running time was later improved

¹In addition, p is assumed to be a “regular” prime in [Hua84; Hua91a; Hua91b] and also in [Rón92]. This condition can be removed. See Section 5.3 for a discussion.

by a certain constant factor [CH00; IKS09; Gua09; Aro13].

Efforts were made to understand the combinatorics behind Rónyai’s and Evdokimov’s algorithms [CH00; Gao01], culminating in the work [IKS09] that proposed the notion of m -schemes together with an algorithm that subsumes those in [Rón88; Evd94] (see also the follow-up work [Aro13; Aro+14]). An m -scheme, parametrized by $m \in \mathbb{N}^+$, is a collection of partitions of sets that satisfies a list of axioms. It was shown in [IKS09] that whenever the algorithm fails to produce a proper factorization, there always exists an m -scheme satisfying strict combinatorial properties. Evdokimov’s result can then be interpreted as the fact that such an m -scheme does not exist for sufficiently large $m = O(\log n)$. Finally, a conjecture on m -schemes, known as the *schemes conjecture*, was proposed in [IKS09], whose affirmative resolution would imply a polynomial-time factoring algorithm under GRH.

Role of GRH. GRH asserts that all nontrivial zeros of Dirichlet L-functions are on the line $\operatorname{Re}(z) = 1/2$. As noted in [Rón92], the known GRH-based algorithms (including our work) only need a consequence of GRH that finite fields can be efficiently constructed, and their k th power non-residues² can be efficiently found. Formally, for all the statements made under GRH throughout this thesis, we may use the following hypothesis instead.

Hypothesis ()*. There exists a deterministic algorithm that given a prime number p and an integer $d \in \mathbb{N}^+$, constructs³ the finite field \mathbb{F}_{p^d} in time polynomial in $d \log p$. In addition, given any prime factor k dividing $p^d - 1$, a k th power non-residue of \mathbb{F}_{p^d} can be found deterministically in time polynomial in k and $d \log p$.

See [Hua91b; LMO79] for the proof that Hypothesis (*) holds under GRH. By [Bha+17], it holds even under a weaker version of GRH, which asserts that all nontrivial zeros of Dirichlet L-functions are in the strip $\operatorname{Re}(z) \in [\frac{1}{2} - \epsilon, \frac{1}{2} + \epsilon]$ for some constant $\epsilon < 1/2$.

1.2 Main results

In this thesis, we introduce a family of mathematical objects called \mathcal{P} -schemes, generalizing the classical notion of association schemes [BI84] as well as the notion of m -schemes [IKS09]. Based on \mathcal{P} -schemes, we develop a unifying framework for deterministic univariate polynomial factoring over finite fields under GRH.

²For a prime factor k of $q - 1$, an element $x \in \mathbb{F}_q^\times$ is a k th power residue of \mathbb{F}_q if $x \in (\mathbb{F}_q^\times)^k$. Otherwise it is a k th power non-residue.

³By constructing \mathbb{F}_{p^d} , we mean finding its *structure constants* in some \mathbb{F}_p -basis. See [Len90].

\mathcal{P} -schemes. Roughly speaking, given a finite group G and a poset \mathcal{P} of subgroups of G , a \mathcal{P} -scheme is collection of partitions,

$$\mathcal{C} = \{C_H : H \in \mathcal{P}\},$$

satisfying certain constraints, where each C_H is a partition of the right coset space $H \backslash G = \{Hg : g \in G\}$. The formal definition is given in Definition 2.4. We also define various properties of \mathcal{P} -schemes, including *antisymmetry*, *strong antisymmetry*, *discreteness*, and *homogeneity*. These properties play important roles in our polynomial factoring algorithms.

When G is chosen to be a symmetric group and \mathcal{P} is a poset of *stabilizer subgroups* (with respect to the natural action of G), we recover the notion of m -schemes [IKS09]:

Theorem 1.1 (informal). *Suppose $G = \text{Sym}(S)$ acts naturally on a finite set S and \mathcal{P} consists of the (pointwise) stabilizers G_T for all subsets $T \subseteq S$ satisfying $1 \leq |T| \leq m$. Then a \mathcal{P} -scheme \mathcal{C} is equivalent to an m -scheme Π on S . Moreover, \mathcal{C} is antisymmetric (resp. strongly antisymmetric, discrete on G_x for $x \in S$, homogeneous on G_x for $x \in S$) iff Π has the corresponding property.*

This result in fact holds as long as G is k -transitive for sufficiently large k . See Theorem 2.1 for the formal statement.

In this way, we regard the theory of m -schemes [IKS09; Aro13; Aro+14] as part of the richer theory of \mathcal{P} -schemes. The advantage of adopting the notion of \mathcal{P} -schemes is that these objects capture not only the combinatorial structure of m -schemes but also the information provided by the group G and the poset \mathcal{P} , which allows us to carry out both the Galois-theoretic/group-theoretic approach [Hua91a; Hua91b; Evd92; Rón92] and the combinatorial approach [Evd94; IKS09] of deterministic polynomial factoring in a uniform way.

A unifying framework for deterministic polynomial factoring. The theory of \mathcal{P} -schemes is applied to deterministic polynomial factoring as follows. For simplicity, assume f is a degree- n polynomial that is defined over a prime field \mathbb{F}_p and factorizes into n distinct linear factors over \mathbb{F}_p . Let $\tilde{f}(X) \in \mathbb{Z}[X]$ be an *irreducible lifted polynomial* of f , defined as follows:

Definition 1.1 (lifted polynomial). *A lifted polynomial of a degree- n polynomial $f(X) \in \mathbb{F}_p[X]$ is a polynomial $\tilde{f}(X) \in \mathbb{Z}[X]$ of degree n satisfying $\tilde{f} \bmod p = f$.*

An irreducible lifted polynomial of f is a lifted polynomial of f that is irreducible over \mathbb{Q} .

Let L be the splitting field of $\tilde{f}(X)$ over \mathbb{Q} and let $G = \text{Gal}(L/\mathbb{Q})$. By Galois theory, we have a one-to-one correspondence between the subgroups of G and the subfields of L

$$H = \text{Gal}(L/K) \longleftrightarrow K = L^H,$$

where L^H denotes the fixed field of H .

In Chapter 3, we design a generic algorithm, which we refer to as the \mathcal{P} -scheme algorithm, that deterministically factorizes f under GRH given f and \tilde{f} . The generic part of the algorithm is a subroutine that uses \tilde{f} to construct a poset of subfields of L , which in turn corresponds to a poset \mathcal{P} of subgroups of G by Galois theory. We then prove that the algorithm always produces the complete factorization (resp. a proper factorization) of f under GRH, unless a combinatorial condition regarding \mathcal{P} -schemes fails to hold.⁴ Therefore the problem of deterministic polynomial factoring reduces to the problem of verifying this combinatorial condition about \mathcal{P} -schemes.

By choosing various posets \mathcal{P} and verifying the condition, we recover the main results of the previous work [Hua91a; Hua91b; Rón88; Rón92; Evd92; Evd94; IKS09] using the \mathcal{P} -scheme algorithm. Our algorithm thus provides a unifying framework for deterministic polynomial factoring over finite fields.

The generalized \mathcal{P} -scheme algorithm. The \mathcal{P} -scheme algorithm above is subject to the condition that the input polynomial is defined over a prime field \mathbb{F}_p and factorizes into distinct linear factors over \mathbb{F}_p . In Chapter 5, we extend it to the *generalized \mathcal{P} -scheme algorithm* that works for arbitrary polynomials $f(X) \in \mathbb{F}_q[X]$. The results obtained from the \mathcal{P} -scheme algorithm are then proved in full generality.

Several new ideas and a significant amount of work are required in the development of the generalized \mathcal{P} -scheme algorithm. See Chapter 5 for the details.

Constructing new \mathcal{P} -schemes from old ones. We develop various techniques of constructing new \mathcal{P} -schemes from old ones, including *restriction*, *induction*, *extension*, etc. These techniques are useful for investigating the existence of certain \mathcal{P} -schemes, allowing us to reduce one case to another.

⁴The condition requires all strongly antisymmetric \mathcal{P} -schemes to be discrete (resp. inhomogeneous) on G_x , where x is a root of \tilde{f} in L . See Theorem 3.9 for the formal statement.

In particular, using induction of \mathcal{P} -schemes, we show that for finite groups $H \subseteq G$ and a poset \mathcal{P} of subgroups of H , a \mathcal{P} -scheme with various properties (antisymmetry, strong antisymmetry, etc.) can be used to construct a \mathcal{P}' -scheme with the same properties, where \mathcal{P}' is a certain poset of G . Intuitively, this means polynomial factoring “becomes easier” if the Galois group G is replaced by a subgroup H . We make this intuition rigorous regarding the *schemes conjecture* proposed in [IKS09]. See below for a more detailed discussion.

In addition, we define the *direct product* and the *wreath product* of \mathcal{P} -schemes, generalizing the corresponding operations of permutation groups and association schemes [SS98; Bai04]. We also define the direct product and the wreath product of m -schemes. A consequence of these operations is that either the schemes conjecture in [IKS09] holds, or it has infinitely many counterexamples.

Schemes conjectures for families of permutation groups. The work [IKS09] proposed a combinatorial conjecture on m -schemes, called the *schemes conjecture*, whose positive resolution would imply a deterministic polynomial-time factoring algorithm under GRH. Proving this conjecture appears to be difficult. However, as noted in Theorem 1.1 above, an m -scheme is essentially a \mathcal{P} -scheme in the (worst) case of symmetric groups, with respect to a poset \mathcal{P} of pointwise stabilizers. This observation suggests that one should first formulate and attack the analogous conjectures for “less complex” Galois groups.

For each family \mathcal{G} of finite permutation groups, we formulate an analogous conjecture, called the *schemes conjecture for \mathcal{G}* . Like the original scheme conjecture, the schemes conjecture for \mathcal{G} also implies a deterministic polynomial-time factoring algorithm under GRH, provided that that Galois group of the lifted polynomial \tilde{f} , as a permutation group on the set of roots of \tilde{f} , is a member of \mathcal{G} . Moreover, we show that these conjectures form a hierarchy of relaxations of the original schemes conjecture in [IKS09]. More specifically, for two families of finite permutation groups \mathcal{G} and \mathcal{G}' such that every member of \mathcal{G} is (permutation isomorphic to) a subgroup of member in \mathcal{G}' , the schemes conjecture for \mathcal{G} is implied by that for \mathcal{G}' . The worst case occurs when \mathcal{G} is the family of symmetric groups, which yields (a slight relaxation of) the original scheme conjecture. We hope progress on this hierarchy of conjectures will shed some light on the original schemes conjecture and pave the way for solving deterministic polynomial factoring over finite fields in polynomial time under GRH.

Galois groups with restricted noncyclic composition factors. Using our framework of \mathcal{P} -schemes, we design a GRH-based deterministic factoring algorithm that completely factorizes a polynomial f using a lifted polynomial \tilde{f} , such that the running time of the algorithm is controlled by the *noncyclic composition factors*⁵ of the Galois group of \tilde{f} . More specifically, we have

Theorem 1.2 (informal). *Under GRH, there exists a deterministic algorithm that given $f(X) \in \mathbb{F}_q[X]$ and a lifted polynomial⁶ \tilde{f} of f with the Galois group G , completely factorizes f in time polynomial in $k(G)^{\log k(G)}$, $r(G)$ and the size of the input, where $k(G)$ (resp. $r(G)$) is the maximum degree (resp. maximum order) of the alternating groups (resp. classical groups) among the composition factors of G .*

See Theorem 8.2 for the formal statement. Now fix $k \in \mathbb{N}^+$ and consider the family of finite groups whose noncyclic composition factors are all isomorphic to subgroups of $\text{Sym}(k)$. This family is commonly denoted by Γ_k in the literature, and plays a significant role in graph isomorphism testing [Luk82; Mil83], asymptotic group theory [BCP82; Pyb93; PS97] and computational group theory [Luk93; Ser03]. It is known that a classical group of order r lies in Γ_k only if $r = k^{O(\log k)}$ [Coo78]. So Theorem 1.2 implies

Theorem 1.3 (informal). *Under GRH, there exists a deterministic algorithm that given $f(X) \in \mathbb{F}_q[X]$ of degree n and a lifted polynomial \tilde{f} of f , completely factorizes f in time polynomial in n , $\log q$ and $k^{\log k}$, where k is the smallest positive integer such that the Galois group of \tilde{f} is in Γ_k .*

See Theorem 8.3 for the formal statement. It refines and generalizes the main results of [Hua91a; Hua91b; Evd92; Rón92; Evd94]. Note that the algorithm runs in polynomial time under GRH provided that $k = 2^{O(\sqrt{\log n})}$. Previously, polynomial-time factoring algorithms for Γ_k were known only for bounded k under GRH [Evd92; BCP82].

Other results. Finally, we list some other results obtained in this thesis.

⁵Recall that a composition factor of a finite group is a *finite simple group*, and by the *classification of finite simple groups* (CFSG) it is isomorphic to one of the following groups: a cyclic group of prime order, an alternating group, a classical group, an exceptional group of Lie type, or one of the 26 sporadic simple groups.

⁶For a general (not necessarily prime) finite field \mathbb{F}_q , we use a more general definition of lifted polynomials (Definition 5.1) instead of Definition 1.1.

1. The schemes conjecture in [IKS09] asserts that if a homogeneous antisymmetric orbit m -schemes on a set S has no matching, then $m = O(1)$ (see Chapter 2 for the definition of matchings). Currently, the best known upper bound for m is $m \leq c \log |S| + O(1)$, where $c = \frac{2}{\log 12} = 0.5578 \dots$. We consider the analogous problem for a general linear group $\text{GL}(V)$ over a finite field \mathbb{F}_q acting naturally on $S = V - \{0\}$, and show that for this new problem, we have a slightly improved bound $m \leq c' \log |S| + O(1)$ where $c' = \frac{4}{4 \log q + \log 12} \leq 0.5273 \dots$ (Theorem 7.5). In addition, we consider the analogous problems for the groups $\text{GL}(V)$, $\Gamma\text{L}(V)$, $\text{PGL}(V)$, and $\text{PTL}(V)$, and show that these problems are equivalent, in the sense that the optimal values of m for them differ from each other by at most a constant (Theorem 7.4).
2. We generalize the notion of orbit schemes in [IKS09], or what we call *orbit m -schemes*, to the notion of *orbit \mathcal{P} -schemes*. We also prove that an orbit m -scheme associated with a group K is antisymmetric iff the order of K is coprime to $1, 2, \dots, m$ (Lemma 2.16), which in turn shows that a result of [Rón88; IKS09] on antisymmetric m -schemes is tight (cf. Lemma 2.17 and Example 2.2).
3. The paper [IKS09] showed that the schemes conjecture is true when restricted to orbit schemes, by proving that all antisymmetric homogeneous orbit m -schemes on a set of cardinality greater than one have a matching for $m \geq 4$. We prove that the later statement in fact holds for $m \geq 3$ (Theorem 6.6).

1.3 Outline of the thesis.

Basic notations and preliminaries are given in the next section, and additional preliminaries are given at the beginning of subsequent chapters.

Chapter 2 introduces definitions and develops basic results about \mathcal{P} -schemes: we first define \mathcal{P} -schemes and their various properties. After reviewing the notion of m -schemes in [IKS09] and their connection with association schemes, we prove the formal version of Theorem 1.1 above. Then we investigate the notion of *orbit schemes* in [IKS09], and extend it to our framework of \mathcal{P} -schemes. Finally, some concrete examples of *strongly antisymmetric homogeneous m -schemes* are given for small m .

The rest of the thesis is divided into two parts: Chapters 3–5 constitute the algorithmic part of the thesis, whereas Chapters 6–8 focus on further development of the theory of \mathcal{P} -schemes. The latter is mostly algorithm-free, except that Section 8.1

contains an algorithm that depends on Section 4.2, Section 4.3, and Theorem 5.9. The dependencies among chapters are roughly illustrated in Figure 1.1.




Figure 1.1: Dependencies among chapters

In Chapter 3, we develop the \mathcal{P} -scheme algorithm, and use it to reprove the main results of [Hua91a; Hua91b; Rón88; Rón92; Evd94; IKS09]. As mentioned above, the results in Chapter 3 are subject to the condition that the input polynomial is defined over a prime field \mathbb{F}_p and factorizes into distinct linear factors over \mathbb{F}_p .

The \mathcal{P} -scheme algorithm requires a subroutine that constructs a collection of number fields. In Chapter 4, we discuss various ways of implementing this subroutine and survey techniques of constructing number fields in the literature [Len83; Lan84; Lan85; LM85; Evd92].

In Chapter 5, we develop the generalized \mathcal{P} -scheme algorithm where the condition about the input polynomial is no longer needed. The results in Chapter 3 are then proved in full generality.

Chapter 6 develops various techniques of constructing new \mathcal{P} -schemes from old ones. In Section 6.3, we formulate the schemes conjectures for families of finite permutation groups and show that these conjectures form a hierarchy of relaxations of the scheme conjecture proposed in [IKS09]. Our result that an antisymmetric homogeneous orbit m -scheme on a set of cardinality $n > 1$ has a matching for $m \geq 3$ is proved in Section 6.6, where we also discuss *primitivity* of m -schemes.

Chapter 7 discusses the (non-)existence of certain \mathcal{P} -schemes for symmetric groups and linear groups. In particular, we review the result in [Aro13] on m -schemes (based on the work of [Evd94; IKS09], and independently discovered in [Gua09]), and interpret it as a result about \mathcal{P} -schemes with respect to the natural action of symmetric groups. We also extend it to a more general result about \mathcal{P} -schemes with respect to *standard actions* of symmetric groups. The analysis employs a technical “self-reduction lemma” proven in Section 7.2, which is also heavily used in Chapter 8. Some results about \mathcal{P} -schemes for linear groups are also given.

Finally, in Chapter 8, we describe our deterministic factoring algorithm for Galois groups with restricted noncyclic composition factors. More specifically, we give the algorithm and its analysis in Section 8.1, assuming a statement about \mathcal{P} -schemes for primitive permutation groups (Theorem 8.4). The rest of Chapter 8 then focuses on verifying this statement.

1.4 Notations and preliminaries

Denote by \mathbb{N}^+ the set of positive integers. For $k \in \mathbb{N}^+$, we denote by $[k]$ the set $\{1, 2, \dots, k\}$. For two sets A and B , write $A - B$ for the set difference $\{x : x \in A \text{ and } x \notin B\}$.⁷ The cardinality of a finite set S is denoted by $|S|$. Denote by \log the logarithmic function with base 2.

A partition of a finite set S is a set P of nonempty subsets of S satisfying $S = \bigsqcup_{B \in P} B$, where \bigsqcup denotes the disjoint union. Each $B \in P$ is called a *block* of P . For two partitions P and P' of S , we say P *refines* P' , or P is a *refinement* of P' , if every block in P' is a disjoint union of blocks in P . We say the refinement is *proper* if $P \neq P'$. Denote by 0_S the coarsest partition of S , i.e. the one consisting of a single block S . Denote by ∞_S the finest partition of S , i.e., $\infty_S = \{\{x\} : x \in S\}$. For $T \subseteq S$ and a partition P of S , define $P|_T := \{B \cap T : B \in P\} - \{\emptyset\}$ which is a partition of T , called the *restriction* of P to T . For a set S and $k \in \mathbb{N}^+$, define the set $S^{(k)} := \{(x_1, \dots, x_k) \in S^k : x_i \neq x_j \text{ for } i \neq j\}$ consisting of k -tuples of distinct elements.

Write $f \circ g$ for the composition of two functions f and g , from right to left. We note that this is the common convention, although group theorists often use the opposite convention gf . For a function f and a subset T of the domain of f , denote by $f|_T$ the restriction of f to T . For a field K , denote the characteristic of K by $\text{char}(K)$.

⁷This is often denoted by $A \setminus B$. We use $A - B$ to avoid confusion with a right coset space $H \setminus G$.

A polynomial is *monic* if its leading coefficient is one. For two polynomials $f(X), g(X) \in \mathbb{F}_q[X]$ over a finite field \mathbb{F}_q that are not both zero, define their *greatest common divisor* $\gcd(f, g)$ to be the unique monic polynomial $h(X) \in \mathbb{F}_q[X]$ of the greatest degree that divides both f and g . It is well defined since $\mathbb{F}_q[X]$ is a unique factorization domain, and can be computed efficiently from f and g using the *Euclidean algorithm* [GG13].

Basic notations about groups. All groups in this thesis are finite. Write e for the identity element of a group. For a group G , a subgroup H of G , and $g \in G$, write gH for the *left coset* $\{gh : h \in H\}$ and Hg for the *right coset* $\{hg : h \in H\}$. Write G/H for the *left coset space* $\{gH : g \in G\}$ and $H \backslash G$ for the *right coset space* $\{Hg : g \in G\}$. For two subgroups H, K of G and $g \in G$, write HgK for the *double coset* $\{hgh' : h \in H, h' \in K\}$, and write $H \backslash G / K$ for the *double coset space* $\{HgK : g \in G\}$. Define $[G : H] := |G|/|H|$, called the *index* of H in G . Write $\langle H_1, \dots, H_k \rangle$ for the *join* of subgroups H_1, \dots, H_k , i.e., the subgroup generated by H_1, \dots, H_k . Write $\langle g_1, \dots, g_k \rangle$ for the subgroup generated by the group elements g_1, \dots, g_k .

A *subquotient* of a group G is a quotient group of a subgroup of G . Two subgroups H and H' are said to be *conjugate* in G if $H' = gHg^{-1}$ for some $g \in G$. A subgroup H is said to be *normal* in G or a *normal subgroup* of G if $gHg^{-1} = H$ for all $g \in G$. Write $H \trianglelefteq G$ for H being normal in G . Define the *normalizer* of H in G to be $N_G(H) := \{g \in G : gHg^{-1} = H\}$. We have $H \trianglelefteq N_G(H)$, and indeed $N_G(H)$ is the unique maximal subgroup of G with this property. The *center* of G , denoted by $Z(G)$, is the subgroup $\{g \in G : gh = hg \text{ for all } h \in G\}$. A subgroup H of G is *maximal* if $H \neq G$ and there exists no subgroup H' of G satisfying $H \subsetneq H' \subsetneq G$.

For a finite set S , denote by $\text{Sym}(S)$ and $\text{Alt}(S)$ the symmetric group and the alternating group on S respectively. We also write $\text{Sym}(n)$ and $\text{Alt}(n)$ when $S = [n]$. Permutations are often written in the cycle notation, where $(a_1 a_2 \cdots a_n)$ denotes the cyclic permutation sending a_i to a_{i+1} for $1 \leq i < n$ and a_n to a_1 .

For a group G , denote by $\text{Aut}(G)$ the automorphism group of G , i.e., the group of invertible homomorphisms $\rho : G \rightarrow G$ where the group operation is defined by composition. For $g \in G$, the map $\tau_g : G \rightarrow G$ sending $h \in G$ to ghg^{-1} is an automorphism of G , called an *inner automorphism* of G . Define $\text{Inn}(G) := \{\tau_g : g \in G\}$, called the *inner automorphism group* of G , which is a normal subgroup of $\text{Aut}(G)$. Define $\text{Out}(G) := \text{Aut}(G)/\text{Inn}(G)$, called the *outer automorphism*

group of G .

Group actions. Let G be a group and S be a finite set. A (*left*) *group action* or an *action* of G on S is a function $\varphi : G \times S \rightarrow S$ satisfying (1) $\varphi(e, x) = x$ for all $x \in S$ and (2) $\varphi(g, \varphi(h, x)) = \varphi(gh, x)$ for all $x \in S$ and $g, h \in G$. We also say G *acts on* S and S is a G -*set*. We usually denote $\varphi(g, x)$ as ${}^g x$ when φ is clear from the context. For $T \subseteq S$, write ${}^g T$ for the set $\{{}^g x : x \in T\}$. Again, we note that group theorists commonly adopt the right action convention $xgh = (xg)h$ instead of our left action convention. One can switch between the two conventions by taking the inverse map $g \mapsto g^{-1}$.

Given a G -set S , the elements of G act as permutations of S . This gives a group homomorphism $\rho : G \rightarrow \text{Sym}(S)$, called a *permutation representation* of G on S . The action of G on S is *faithful* if ρ is injective. The image $\rho(G)$ is called a *permutation group* on S . When the action is faithful and clear from the context, we usually just say G is a permutation group on S .

Orbits and stabilizers. For a G -set S , the *orbit* or G -*orbit* of an element $x \in S$ is $Gx := \{{}^g x : g \in G\}$. The set S is a disjoint union of its G -orbits. The *stabilizer* of $x \in S$ is $G_x := \{g \in G : {}^g x = x\}$. For $T \subseteq S$, define the *pointwise stabilizer*

$$G_T := \{g \in G : {}^g x = x \text{ for all } x \in T\}$$

and the *setwise stabilizer*

$$G_{\{T\}} := \{g \in G : {}^g T = T\}.$$

For $T = \{x_1, \dots, x_k\} \subseteq S$ we also write G_{x_1, \dots, x_k} for G_T . Let $S^G := \{x \in S : {}^g x = x \text{ for all } g \in G\}$ be the *set of fixed points* of G .

An action of G on a set S is *transitive* if it has only one orbit. It is *semiregular* if G_x is trivial for all $x \in S$. A group action is *regular* if it is both transitive and semiregular. For $k \in \mathbb{N}^+$, an action of G on S induces an action on $S^{(k)}$ via

$${}^g(x_1, \dots, x_k) = ({}^g x_1, \dots, {}^g x_k),$$

called the *diagonal action* of G on $S^{(k)}$. For $1 \leq k \leq |S|$, we say the action of G on S is k -*transitive* if the corresponding diagonal action of G on $S^{(k)}$ is transitive. We say it is $(k + 1/2)$ -*transitive* if it is k -transitive, and in addition for all $T \subseteq S$ of cardinality k , either the G_T -orbit of every $x \in S - T$ contains more than one element, or $|S - T| = 1$. A $(k + 1)$ -transitive action is also $(k + 1/2)$ -transitive. For more discussion about half transitivity, see [Wie64].

G -modules and G -invariant elements. Given a group G , an abelian group A is called a G -module if it has an action of G compatible with its abelian group structure, i.e., ${}^g x + {}^g y = {}^g(x + y)$ for $x, y \in A$ and $g \in G$. The set of fixed points A^G is a subgroup of A , known as the subgroup of G -invariant elements of A . Suppose in addition that A is a ring (resp. field) and the action of G respects the multiplication of A as well, then A^G is a subring (resp. subfield) of A , called the *fixed subring* (resp. *fixed subfield*) of A corresponding to G .

Chapter 2

\mathcal{P} -SCHEMES

We introduce the notion of \mathcal{P} -schemes in this chapter, which plays a central role throughout the thesis. A \mathcal{P} -scheme is a combinatorial structure associated with a group G and a conjugation-closed poset \mathcal{P} of subgroups of G . Roughly speaking, it contains a collection of partitions of right coset spaces $H \backslash G$ for $H \in \mathcal{P}$, and these partitions satisfy various consistency properties.

For every permutation group G , we define the integers $d(G), d'(G) \in \mathbb{N}^+$ in terms of \mathcal{P} -schemes associated with G , and show that they are bounded by the *minimum base size* of G . We will see in Chapter 3 that $d(G)$ and $d'(G)$ are closely related to deterministic polynomial factoring.

The work [IKS09] proposed the notion of m -schemes as a “higher-order” generalization of *association schemes* that are central in the field of *algebraic combinatorics* [BI84]. We show \mathcal{P} -schemes are further generalization of m -schemes: an m -scheme arises as a \mathcal{P} -scheme associated with a symmetric group, or more generally with a multiply transitive group action.

Other results in this chapter include:

- We define *orbit \mathcal{P} -schemes*, generalizing the notion of *orbit m -schemes* in [IKS09]. We also provide a simple and exact criterion for antisymmetry of orbit m -schemes. Using this criterion, we give examples of antisymmetric homogeneous orbit m -schemes on finite sets S for m up to $\ell - 1$, where ℓ is the least prime factor of $|S|$. This result matches the upper bound $m < \ell$ established by Rónyai [Rón88] for *arbitrary* antisymmetric homogeneous m -schemes. We reproduce Rónyai’s argument and extend it to \mathcal{P} -schemes.
- We also provide examples of m -schemes for small values of m . In particular, for $m \leq 3$, we give explicit constructions of m -schemes satisfying the properties of *strong antisymmetry* and *homogeneity* that are closely related to deterministic polynomial factoring.

Outline of the chapter. Preliminaries are given in Section 2.1. In Section 2.2, we define the notion of \mathcal{P} -schemes and its various properties. We also define $d(G)$ and

$d'(G)$ in terms of \mathcal{P} -schemes. In Section 2.3, we review the notion of m -schemes and prove the equivalence between m -schemes and a certain kind of \mathcal{P} -schemes. We also discuss the connection between m -schemes and association schemes. In Section 2.4, we define orbit m -schemes as well as orbit \mathcal{P} -schemes. An exact criterion of antisymmetry is given for orbit m -schemes. Then we discuss Rónyai's upper bound for m for antisymmetric homogeneous m -schemes and extend it to \mathcal{P} -schemes. Finally, in Section 2.5, we describe explicit constructions of strongly antisymmetric homogeneous m -schemes for $m \leq 3$.

2.1 Preliminaries

Let G be a group. A *partially ordered set* or *poset* of subgroups of G is simply a set of subgroups of G , partially ordered by inclusion. All posets of subgroups in this thesis are assumed to be conjugation-closed, and we give the following definition for such posets.

Definition 2.1 (subgroup system). *A poset \mathcal{P} of subgroups of G is called a subgroup system over G if it is closed under conjugation in G , i.e., $gHg^{-1} \in \mathcal{P}$ for all $H \in \mathcal{P}$ and $g \in G$.*

We introduce \mathcal{P} -schemes in next section, each associated with a subgroup system \mathcal{P} . While the definitions are formulated for general subgroup systems, those arising from the factoring algorithms have special forms. In particular, the following kind of subgroup systems are frequently used in the algorithms.

Definition 2.2 (system of stabilizers). *Suppose G is a finite group acting on a finite set S . For $m \in \mathbb{N}$, let \mathcal{P}_m be the set of pointwise stabilizers for nonempty subsets $T \subseteq S$ of cardinality up to m :*

$$\mathcal{P}_m := \{G_T : T \subseteq S, 1 \leq |T| \leq m\}.$$

Then \mathcal{P}_m is a subgroup system over G , called the system of stabilizers of depth m (with respect to the action of G on S).

Left and inverse right translation. Let H be a subgroup of G . There is an action of G on the right coset space $H \backslash G$ defined by

$${}^g Hh = Hhg^{-1} \quad \text{for } Hh \in H \backslash G \text{ and } g \in G,$$

called the action of G on $H \backslash G$ by *inverse right translation*. More generally, for a subgroup $G' \subseteq G$, we have the action of G' on $H \backslash G$ by inverse right translation, defined by restricting the previous action of G to G' .

We also have an action of the normalizer $N_G(H)$ on $H \backslash G$ defined by

$${}^g Hh = Hgh \quad \text{for } Hh \in H \backslash G \text{ and } g \in N_G(H),$$

called the action of $N_G(H)$ on $H \backslash G$ by *left translation*.

It is easy to see that they are indeed well defined group actions. For example, we check that for left translation, the coset ${}^g Hh = Hgh$ is independent of the representative h of Hh : Suppose a different representative h' is chosen such that $Hh = Hh'$, then we have $gh'(gh)^{-1} = gh'h^{-1}g^{-1} \in gHg^{-1} = H$ for $g \in N_G(H)$ and hence $Hgh = Hgh'$.

For any $h \in G$, it holds that $Hgh = Hh$ iff $g \in H$. So the action of $N_G(H)$ on $H \backslash G$ induces a semiregular action of $N_G(H)/H$ on $H \backslash G$, defined by ${}^{gH} Hh = Hgh$, called the action of $N_G(H)/H$ on $H \backslash G$ by left translation.

Equivalent actions and permutation isomorphic actions. Let G be a group and let S, T be G -sets. We say the actions of G on S and T are *equivalent* if there exists a bijective map $\lambda : S \rightarrow T$ satisfying $\lambda({}^g x) = {}^g(\lambda(x))$ for all $x \in S$ and $g \in G$. And λ is said to be an *equivalence* between the two actions.

More generally, suppose $\phi : G \rightarrow H$ is a group isomorphism, S is a G -set, and T is an H -set. We say the action of G on S is *permutation isomorphic* to the action of H on T (with respect to ϕ) if there exists a bijective map $\lambda : S \rightarrow T$ satisfying $\lambda({}^g x) = \phi({}^g(\lambda(x)))$ for all $x \in S$ and $g \in G$.

The following lemma states that any transitive group action is equivalent to the action on a right coset space by inverse right translation.

Lemma 2.1. *Let G be a group acting transitively on a set S . For any $x \in S$, the map $\lambda_x : S \rightarrow G_x \backslash G$ sending ${}^g x$ to $G_x g^{-1}$ for $g \in G$ is well defined and is an equivalence between the action of G on S and that on $G_x \backslash G$ by inverse right translation.*

Proof. As the action of G on S is transitive, for any $y \in S$ we can choose $g \in G$ such that $y = {}^g x$. Suppose g, g' are two such choices. We have ${}^{g^{-1}g'} x = {}^{g^{-1}} y = x$

and hence $g^{-1}g' \in G_x$. So $G_x g^{-1} = G_x g'^{-1}$. Therefore λ_x is well defined. It is surjective since any coset $G_x g \in G_x \backslash G$ is the image of $g^{-1}x$ for a representative g of $G_x g$. And it is injective since $G_x g^{-1} = G_x g'^{-1}$ implies $g^{-1}g' \in G_x$ and hence $g'x = g(g^{-1}g')x = gx$. Finally we check that for any $y = gx$ and $h \in G$, it holds that

$$\lambda_x(hy) = \lambda_x(hgx) = G_x(hg)^{-1} = (G_x g^{-1})h^{-1} = {}^h(\lambda(y))$$

as desired. \square

Corollary 2.1 (orbit-stabilizer theorem). *Let S be a G -set for a finite group G . Then $|Gx| = |G|/|G_x|$ for any $x \in S$.*

Projections and conjugations. We define the following two kinds of maps between right coset spaces $H \backslash G$ for various subgroups $H \subseteq G$:

- (projection) for $H \subseteq H' \subseteq G$, define the *projection* $\pi_{H,H'} : H \backslash G \rightarrow H' \backslash G$ to be the map sending $Hg \in H \backslash G$ to $H'g \in H' \backslash G$, and
- (conjugation) for $H \subseteq G$ and $g \in G$, define the *conjugation* $c_{H,g} : H \backslash G \rightarrow gHg^{-1} \backslash G$ to be the map sending $Hh \in H \backslash G$ to $(gHg^{-1})gh \in gHg^{-1} \backslash G$.

Lemma 2.2. *The maps $\pi_{H,H'}$ and $c_{H,g}$ are well defined and satisfy the following properties:*

- *The maps $\pi_{H,H'}$ are surjective and $c_{H,g}$ are bijective.*
- $c_{H',g} \circ \pi_{H,H'} = \pi_{gHg^{-1},gH'g^{-1}} \circ c_{H,g}$.
- (transitivity) $\pi_{H',H''} \circ \pi_{H,H'} = \pi_{H,H''}$ and $c_{gHg^{-1},g'} \circ c_{H,g} = c_{H,g'}$.
- (G -equivariance) $\pi_{H,H'}({}^g Hh) = {}^g \pi_{H,H'}(Hh)$ and $c_{H,g'}({}^g Hh) = {}^g c_{H,g'}(Hh)$ with respect to the action of G on $H \backslash G$ by inverse right translation.

Proof. The proof is straightforward from the definitions. We check $c_{H',g} \circ \pi_{H,H'} = \pi_{gHg^{-1},gH'g^{-1}} \circ c_{H,g}$ and leave the rest to the reader: For $Hh \in H \backslash G$, we have

$$c_{H',g} \circ \pi_{H,H'}(Hh) = c_{H',g}(H'h) = (gH'g^{-1})gh$$

and

$$\pi_{gHg^{-1},gH'g^{-1}} \circ c_{H,g}(Hh) = \pi_{gHg^{-1},gH'g^{-1}}((gHg^{-1})gh) = (gH'g^{-1})gh$$

as desired. \square

Note that for $g \in N_G(H)$, the map $c_{H,g}$ is the permutation of $H \backslash G$ sending each Hh to gHh with respect to the action of $N_G(H)$ on $H \backslash G$ by left translation.

2.2 \mathcal{P} -schemes

We start with the definition of a \mathcal{P} -collection, which is a collection of partitions of right coset spaces.

Definition 2.3 (\mathcal{P} -collection). *Let \mathcal{P} be a subgroup system over a finite group G . A \mathcal{P} -collection \mathcal{C} is a family $\{C_H : H \in \mathcal{P}\}$ indexed by \mathcal{P} where each C_H is a partition of $H \backslash G$.*

We are now ready to define the central object of this thesis.

Definition 2.4 (\mathcal{P} -scheme). *A \mathcal{P} -collection $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ is a \mathcal{P} -scheme if it has the following properties:*

- (compatibility) *for $H, H' \in \mathcal{P}$ with $H \subseteq H'$ and $x, x' \in H \backslash G$ in the same block of C_H , the images $\pi_{H,H'}(x)$ and $\pi_{H,H'}(x')$ are in the same block of $C_{H'}$.*
- (invariance) *for $H \in \mathcal{P}$ and $g \in G$, the map $c_{H,g} : H \backslash G \rightarrow gHg^{-1} \backslash G$ maps any block of C_H to a block of $C_{gHg^{-1}}$.*
- (regularity) *for $H, H' \in \mathcal{P}$ with $H \subseteq H'$, any block $B \in C_H$, $B' \in C_{H'}$, the number of $x \in B$ satisfying $\pi_{H,H'}(x) = y$ is a constant when y ranges over the elements of B' .*

It is worth noting that in a \mathcal{P} -scheme, the partition of $H \backslash G$ for some $H \in \mathcal{P}$ determines the partitions of $H' \backslash G$ for all $H' \in \mathcal{P}$ containing H :

Lemma 2.3. *Let $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ be a \mathcal{P} -scheme. For $H, H' \in \mathcal{P}$ with $H \subseteq H'$, the blocks of $C_{H'}$ are exactly the images of the blocks of C_H under $\pi_{H,H'}$.*

Proof. Let B' be a block of $C_{H'}$. By compatibility, B' is a union of $\pi_{H,H'}(B)$ for one or more blocks $B \in C_H$. Assume $\pi_{H,H'}(B) \subsetneq B'$ for some $B \in C_H$ and choose $y \in \pi_{H,H'}(B)$, $y' \in B' - \pi_{H,H'}(B)$. Then we have $|\{x \in B : \pi_{H,H'}(x) = y\}| > 0$ but $|\{x \in B : \pi_{H,H'}(x) = y'\}| = 0$, which contradicts regularity. \square

In particular, if \mathcal{P} has the property that all minimal subgroups in \mathcal{P} are conjugate in G , then by invariance and Lemma 2.3, the partition for one of the minimal

subgroups determines the whole \mathcal{P} -scheme. For instance, this holds if \mathcal{P} is a system of stabilizers \mathcal{P}_m with respect to an m -transitive group action.

Remark. Besides the set-theoretic definition of \mathcal{P} -schemes given in Definition 2.4, there also exists an equivalent “algebraic” or ring-theoretic definition of \mathcal{P} -schemes. It formulates the three defining properties (compatibility, invariance, and regularity) in a unifying way as closedness of rings under three kinds of maps, respectively: inclusions, conjugations, and trace maps. The interested reader is referred to Appendix A for further discussion.

Next we define some optional properties of \mathcal{P} -schemes.

Homogeneity and discreteness. Recall that for a finite S , we denote by 0_S the coarsest partition of S and ∞_S the finest partition of S .

Definition 2.5. A \mathcal{P} -scheme $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ is homogeneous on a subgroup $H \in \mathcal{P}$ if $C_H = 0_{H \setminus G}$, and otherwise inhomogeneous on H . It is discrete on H if $C_H = \infty_{H \setminus G}$, and otherwise non-discrete on H .

We will see in Chapter 3 that homogeneity (resp. discreteness) of \mathcal{P} -schemes is closely related to whether or not the factoring algorithm always produces a proper factorization (resp. the complete factorization) of the input polynomial.

Symmetry and antisymmetry. Invariance of \mathcal{P} -schemes states that maps $c_{H,g} : Hh \mapsto (gHg^{-1})gh$ always send blocks to blocks. When $g \in N_G(H)$, the map $c_{H,g}$ is a permutation of $H \setminus G$, and we can impose on a \mathcal{P} -scheme the constraint that $c_{H,g}$ always sends a block to itself. Alternatively, we may require $c_{H,g}$ to always send a block to a different block when it is not the trivial permutation. These two constraints are captured by *symmetry* and *antisymmetry* of \mathcal{P} -schemes, respectively.

Definition 2.6. A \mathcal{P} -scheme $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ is symmetric if for $H \in \mathcal{P}$ and $g \in N_G(H)$, the permutation $c_{H,g}$ of $H \setminus G$ maps every block of C_H to itself. And \mathcal{C} is antisymmetric if for $H \in \mathcal{P}$ and g in $N_G(H)$ but not in H , the permutation $c_{H,g}$ maps every block of C_H to a different block.

Symmetry (resp. antisymmetry) is equivalent to the property that for all $H \in \mathcal{P}$, elements in each $(N_G(H)/H)$ -orbit of $H \setminus G$ belong to the same block (resp. distinct blocks) of C_H , where $N_G(H)/H$ acts on $H \setminus G$ by left translation.

As will be seen in Chapter 3, antisymmetry of \mathcal{P} -schemes is important for deterministic polynomial factoring [Rón88; Rón92; Evd94; IKS09]. For now we show that an antisymmetric \mathcal{P} -scheme is discrete on H for any $H \in \mathcal{P}$ provided that \mathcal{P} contains the trivial subgroup of G .

Lemma 2.4. *Suppose \mathcal{P} is a subgroup system over a finite group G that contains the trivial subgroup $\{e\}$. For $H \in \mathcal{P}$, all antisymmetric \mathcal{P} -schemes are discrete on H .*

Proof. Let $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ be an antisymmetric \mathcal{P} -scheme. As $N_G(\{e\}) = G$ acts transitively on $\{e\} \backslash G$ by left translation, we have $C_{\{e\}} = \infty_{\{e\} \backslash G}$ by antisymmetry. Now consider an arbitrary subgroup $H \in \mathcal{P}$. By Lemma 2.3, we have $C_H = \{\pi_{\{e\}, H}(B) : B \in C_{\{e\}}\} = \infty_{H \backslash G}$. So \mathcal{C} is discrete on H . \square

On the other hand, symmetry of \mathcal{P} -schemes plays no role in polynomial factoring as far as we know, and we only discuss it within this chapter.

Strong antisymmetry. We introduce another property called *strong antisymmetry*, which is a strengthening of antisymmetry define above. It is based on an idea introduced by Evdokimov [Evd94] which leads to his quasipolynomial-time factoring algorithm.

Antisymmetry states that no nontrivial permutation of blocks arises from a conjugation $c_{H,g}$ where $g \in N_G(H)$: For such a map $c_{H,g}$ and a block $B \in C_H$, either the image $c_{H,g}(B)$ is a different block, or $c_{H,g}$ is the identity map. We strengthen this property by considering permutations arising from compositions of not only conjugations, but also projections and their inverses. Of course, a projection $\pi_{H,H'}$ is not invertible whenever $H \subsetneq H'$. Nevertheless, it is possible that the restriction of $\pi_{H,H'}$ to some block $B \in C_H$ maps B bijectively to some block $B' \in C_{H'}$, in which case the inverse map $(\pi_{H,H'}|_B)^{-1}$ is well defined.

Definition 2.7. *A \mathcal{P} -scheme $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ is strongly antisymmetric if for any sequence of subgroups $H_0, \dots, H_k \in \mathcal{P}$, $B_0 \in C_{H_0}, \dots, B_k \in C_{H_k}$, and maps $\sigma_1, \dots, \sigma_k$ satisfying*

- σ_i is a bijective map from B_{i-1} to B_i ,
- σ_i is of the form $c_{H_{i-1},g}|_{B_{i-1}}$, $\pi_{H_{i-1},H_i}|_{B_{i-1}}$, or $(\pi_{H_i,H_{i-1}}|_{B_i})^{-1}$,

- $H_0 = H_k$ and $B_0 = B_k$,

the composition $\sigma_k \circ \cdots \circ \sigma_1$ is the identity map on $B_0 = B_k$.

In other words, no nontrivial permutation could be obtained by composing maps of the form $c_{H_{i-1},g}|_{B_{i-1}}$, $\pi_{H_{i-1},H_i}|_{B_{i-1}}$, or $(\pi_{H_i,H_{i-1}}|_{B_i})^{-1}$.

A strongly antisymmetric \mathcal{P} -scheme is indeed antisymmetric: Assume $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ is not antisymmetric, then there exist $H \in \mathcal{P}$, $g \in N_G(H) - H$ and $B \in C_H$ such that $c_{H,g}(B) = B$. Let σ_1 be the map $c_{H,g}|_B : B \rightarrow c_{H,g}(B) = B$. It sends $x \in B$ to ${}^g x$ with respect to the action of $N_G(H)/H$ on $H \setminus G$ by left translation. As this action is semiregular and $gH \in N_G(H)/H$ is not the identity element, the map σ_1 is a nontrivial permutation of B . So \mathcal{C} is not strongly antisymmetric.

$d(G)$ and $d'(G)$. For every finite permutation group G , we define $d(G), d'(G) \in \mathbb{N}^+$ which are closely related to deterministic polynomial factoring, as will be seen in Chapter 3.

Definition 2.8. Let G be a finite permutation group on a finite set S . For $m \in \mathbb{N}^+$, let \mathcal{P}_m be the system of stabilizers of depth m with respect to this action. Define $d(G), d'(G) \in \mathbb{N}^+$ as follows.

- Define $d(G)$ to be the smallest integer $m \in \mathbb{N}^+$ such that all strongly antisymmetric \mathcal{P}_m -schemes are discrete on G_x for all $x \in S$.
- If G acts transitively on S and $|S| > 1$, define $d'(G)$ to be the smallest integer $m \in \mathbb{N}^+$ such that all strongly antisymmetric \mathcal{P}_m -schemes are inhomogeneous on G_x for all $x \in S$. Otherwise let $d'(G) = 1$.

We have $1 \leq d'(G) \leq d(G) \leq \max\{|S| - 1, 1\}$ for any finite permutation group G on a finite set S . The first two inequalities are obvious and the last one follows from Lemma 2.4 and the fact that any $g \in G$ fixing $|S| - 1$ elements of S is the identity. A better upper bound for $d(G)$ is given by the *minimal base size* of G .

Definition 2.9 (base). Let G be a finite permutation group on a finite set S . A base of G is a set $B \subseteq S$ for which G_B equals the trivial subgroup $\{e\}$. The minimal base size of G , denoted by $b(G)$, is the minimum cardinality of a base of G .

By Lemma 2.4, we have

Lemma 2.5. $d(G) \leq \max\{b(G), 1\}$ for any finite permutation group G .

We also prove the following bound in latter chapters based on the work of [Evd94; IKS09; Gua09; Aro13].

Lemma 2.6. *There exists an absolute constant $c > 0$ such that $d(G) \leq c \log n + O(1)$ for any finite permutation group G on a set of cardinality $n \in \mathbb{N}^+$.*

The best known upper bound for c is $\frac{2}{\log 12} = 0.55788\dots$, proved by [Gua09; Aro13]. See Section 7.1 for more details.

2.3 m -schemes

The paper [IKS09] proposed the notion of m -schemes. In this section, we present their definition and show that it is generalized by the notion of \mathcal{P} -schemes: roughly speaking, an m -scheme could be regarded as a \mathcal{P} -scheme where \mathcal{P} is a system of stabilizers with respect to an m -transitive group action.

We use the following notations:

Let S be a finite set and let $m \in \mathbb{N}^+$. Define an m -collection on S to be a collection of partitions P_1, \dots, P_m of $S^{(1)}, \dots, S^{(m)}$ respectively.

For $k \in [m]$, the symmetric group $\text{Sym}(k)$ acts on the set $S^{(k)}$ by permuting the k coordinates, i.e., for $g \in \text{Sym}(k)$ and $x = (x_1, \dots, x_k) \in S^{(k)}$, we have ${}^g x = (y_1, \dots, y_k)$ where $y_{g_i} = x_i$, or equivalently $y_i = x_{g^{-1}_i}$.

For $1 < k \leq m$ and $i \in [k]$, let $\pi_i^k : S^{(k)} \rightarrow S^{(k-1)}$ be the projection omitting the k th coordinate. More generally, for a proper subset T of $[k]$, let $\pi_T^k : S^{(k)} \rightarrow S^{(k-r)}$ be the projection omitting the coordinates whose indices are in T .

For $k \in [m]$ and $g \in \text{Sym}(k)$, let c_g^k be the permutation of $S^{(k)}$ sending x to ${}^g x$, with respect to the above action of $\text{Sym}(k)$ on $S^{(k)}$.

Definition 2.10 (m -scheme [IKS09]). *An m -collection $\Pi = \{P_1, \dots, P_m\}$ on S is an m -scheme if it has the following properties:*

- (compatibility) for $1 < k \leq m, i \in [k]$ and elements $x, x' \in S^{(k)}$ in the same block of P_k , the elements $\pi_i^k(x), \pi_i^k(x')$ are in the same block of P_{k-1} .
- (invariance) for $k \in [m]$ and $g \in \text{Sym}(k)$, the permutation c_g^k of $S^{(k)}$ sends blocks of P_k to blocks.

- (regularity) for $1 < k \leq m$, $i \in [k]$ and blocks $B \in P_k$, $B' \in P_{k-1}$, the number of $x \in B$ satisfying $\pi_i^k(x) = y$ is a constant when y ranges over the elements of B' .

Furthermore, we say Π is symmetric (resp. antisymmetric) if for all $k \in [m]$ and $g \in \text{Sym}(k) - \{e\}$, the permutation c_g^k of $S^{(k)}$ sends every block of P_k to itself (resp. a different block). And Π is said to be homogeneous if P_1 equals the coarsest partition 0_S .

We also introduce the following definitions which did not appear in [IKS09].

Definition 2.11. An m -scheme $\Pi = \{P_1, \dots, P_m\}$ on S is said to be discrete if P_1 equals the finest partition ∞_S . It is said to be strongly antisymmetric if no nontrivial permutation of any block of P_k for any $k \in [m]$ can be obtained by composing maps of the form $c_g^i|_B$, $\pi_T^i|_B$, or $(\pi_T^i|_B)^{-1}$, where B is a block of P_i .

Remark. The parameter m is allowed to be arbitrarily large in our definition. Nevertheless, the sets $S^{(k)}$ for $k = |S| + 1, \dots, m$ are empty and hence the corresponding partitions P_k contain no information. By discarding these partitions and replacing m with $\min\{m, |S|\}$, we may assume $m \leq |S|$.

The connection of m -schemes with \mathcal{P} -schemes

Given a finite set S and $m \in \mathbb{N}^+$, let G be a group acting m' -transitively on S where $m' := \min\{m, |S|\}$.¹ Choose $\mathcal{P} = \mathcal{P}_m$ to be the system of stabilizers of depth m with respect to this action (see Definition 2.2). We prove that for such G and \mathcal{P} , every \mathcal{P} -scheme gives rise to an m -scheme on S , and (under an additional assumption), there is a one-to-one correspondence between m -schemes on S and \mathcal{P} -schemes, with various properties (symmetry, antisymmetry, etc.) preserved.

For $k \in [m']$ and $x = (x_1, \dots, x_k) \in S^{(k)}$, let $T_x = \{x_1, \dots, x_k\}$. The stabilizer G_x with respect to the diagonal action of G on $S^{(k)}$ equals the pointwise stabilizer G_{T_x} with respect to the action of G on S , and therefore $G_x = G_{T_x} \in \mathcal{P}$. As the action of G on $S^{(k)}$ is transitive (which follows from m' -transitivity of G on S), by Lemma 2.1, we have an equivalence of group actions

$$\lambda_x : S^{(k)} \rightarrow G_x \backslash G$$

¹In particular, we can take $G = \text{Sym}(S)$ acting naturally on S , which is $|S|$ -transitive.

between the diagonal action of G on $S^{(k)}$ and the action on $G_x \backslash G$ by inverse right translation. It sends ${}^g x$ to $G_x g^{-1}$ for $g \in G$. We use these maps λ_x to construct an m -scheme on S from a \mathcal{P} -scheme, and vice versa.

From a \mathcal{P} -scheme to an m -scheme. We construct an m -scheme on S from a \mathcal{P} -scheme as follows.

Definition 2.12. Given a \mathcal{P} -scheme $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$, define an m -collection $\Pi(\mathcal{C}) = \{P_1, \dots, P_m\}$ on S as follows: for each $k \in [m']$ where $m' = \min\{m, |S|\}$, pick $x = (x_1, \dots, x_k) \in S^{(k)}$, and define $P_k = \{\lambda_x^{-1}(B) : B \in C_{G_x}\}$. For $m' < k \leq m$, the partition P_k is a partition of the empty set $S^{(k)}$ and is unique.

Lemma 2.7. $\Pi(\mathcal{C})$ as defined above is independent of the choices of elements x and is an m -scheme. It is symmetric (resp. antisymmetric, strongly antisymmetric) if \mathcal{C} is symmetric (resp. antisymmetric, strongly antisymmetric). And it is homogeneous (resp. discrete) iff \mathcal{C} is homogeneous on G_x (resp. discrete on G_x) for $x \in S$.

Proof. We may assume $m \leq |S|$. Fix $k \in [m]$ and we show that P_k does not depend on the choice of $x \in S^{(k)}$. Consider two elements $x, x' \in S^{(k)}$. Choose $h \in G$ such that $x' = {}^h x$. Such h exists since G acts transitively on $S^{(k)}$. Then $G_{x'} = hG_x h^{-1}$ and we have the conjugation $c_{G_{x'}, h} : G_x \backslash G \rightarrow G_{x'} \backslash G$ sending $G_x g$ to $G_{x'} hg$. We check that $\lambda_{x'} = c_{G_{x'}, h} \circ \lambda_x$. This holds since for $y = {}^g x' \in S^{(k)}$, we have $\lambda_{x'}(y) = G_{x'} g^{-1}$ and

$$c_{G_{x'}, h} \circ \lambda_x(y) = c_{G_{x'}, h} \circ \lambda_x({}^g x') = c_{G_{x'}, h} \circ \lambda_x({}^{gh} x) = c_{G_{x'}, h}(G_x (gh)^{-1}) = G_{x'} g^{-1}.$$

So $\lambda_{x'}^{-1} = \lambda_x^{-1} \circ c_{G_{x'}, h}^{-1} = \lambda_x^{-1} \circ c_{G_x, h^{-1}}$. As \mathcal{C} is invariant, the conjugation $c_{G_{x'}, h^{-1}}$ sends blocks of $C_{G_{x'}}$ to blocks of C_{G_x} . So the two partitions $\{\lambda_x^{-1}(B) : B \in C_{G_x}\}$ and $\{\lambda_{x'}^{-1}(B) : B \in C_{G_{x'}}\}$ are identical, i.e., the elements x and x' define the same partition P_k .

Next we check that $\Pi(\mathcal{C})$ is an m -scheme. For $1 < k \leq m$, consider the elements $x \in S^{(k)}$ and $x' \in S^{(k-1)}$ as picked in Definition 2.12. Let $\bar{x} = \pi_i^k(x) \in S^{(k-1)}$ so that $G_x \subseteq G_{\bar{x}}$. Choose $h \in G$, satisfying $x' = {}^h \bar{x}$ so that $G_{x'} = hG_{\bar{x}}h^{-1}$. Then the following diagram commutes:

$$\begin{array}{ccc} S^{(k)} & \xrightarrow{\pi_i^k} & S^{(k-1)} \\ \lambda_x \downarrow & & \downarrow \lambda_{x'} \\ G_x \backslash G & \xrightarrow{c_{G_{\bar{x}}, h} \circ \pi_{G_x, G_{\bar{x}}}} & G_{x'} \backslash G. \end{array}$$

To see this, note that for any $y = {}^g x \in S^{(k)}$ where $g \in G$, we have

$$c_{G_{\bar{x},h}} \circ \pi_{G_x, G_{\bar{x}}} \circ \lambda_x(y) = c_{G_{\bar{x},h}} \circ \pi_{G_x, G_{\bar{x}}}(G_x g^{-1}) = c_{G_{\bar{x},h}}(G_{\bar{x}} g^{-1}) = G_{x'} h g^{-1},$$

and

$$\lambda_{x'} \circ \pi_i^k(y) = \lambda_{x'} \circ \pi_i^k({}^g x) = \lambda_{x'}({}^g(\pi_i^k(x))) = \lambda_{x'}({}^g \bar{x}) = \lambda_{x'}({}^{gh^{-1}} x') = G_{x'} h g^{-1},$$

as desired. Also note that the maps λ_x and $\lambda_{x'}$ are bijections, sending blocks to blocks. Compatibility and regularity of $\Pi(\mathcal{C})$ then follow from compatibility, regularity, and invariance of \mathcal{C} .

For $k \in [m]$, $\tau \in \text{Sym}(k)$ and $x = (x_1, \dots, x_k) \in S^{(k)}$, let $x' = c_\tau^k(x) \in S^{(k)}$. Choose $h \in G$ such that $x = {}^h x'$. Then $G_x = h G_{x'} h^{-1}$. We also have $G_x = G_{x'}$ since they are both the pointwise stabilizer G_T with respect to the action of G on S , where $T = \{x_1, \dots, x_k\}$. So $h \in N_G(G_x)$. We claim that the following diagram commutes:

$$\begin{array}{ccc} S^{(k)} & \xrightarrow{c_\tau^k} & S^{(k)} \\ \lambda_x \downarrow & & \downarrow \lambda_{x'} \\ G_x \backslash G & \xrightarrow{c_{G_x, h}} & G_x \backslash G. \end{array}$$

To see this, note that for any $y = {}^g x \in S^{(k)}$ where $g \in G$, we have $c_{G_x, h} \circ \lambda_x(y) = c_{G_x, h}(G_x g^{-1}) = G_x h g^{-1}$, and

$$\lambda_x \circ c_\tau^k(y) = \lambda_x \circ c_\tau^k({}^g x) = \lambda_x({}^g(c_\tau^k(x))) = \lambda_x({}^g x') = \lambda_x({}^{gh^{-1}} x) = G_x h g^{-1},$$

as desired. Invariance of $\Pi(\mathcal{C})$ then follows from invariance of \mathcal{C} . So $\Pi(\mathcal{C})$ is an m -scheme.

The previous diagram also shows that if \mathcal{C} is symmetric (resp. antisymmetric) then so is $\Pi(\mathcal{C})$. Suppose a nontrivial permutation of some block of P_k for some $k \in [m]$ can be obtained by composing maps of the form $c_g^i|_B$, $\pi_T^i|_B$, or $(\pi_T^i|_B)^{-1}$, then using the two diagrams above, we also obtain a nontrivial permutation of some block of $G_x \backslash G$ (where $x \in S^{(k)}$ is as chosen in Definition 2.12) by composing conjugations, projections, and their inverses (restricted to blocks). Therefore, if \mathcal{C} is strongly symmetric, so is $\Pi(\mathcal{C})$.

Finally, for any $x \in S$, the partition P_1 of S is constructed using the bijection $\lambda_x : S \rightarrow G_x \backslash G$ and the partition C_{G_x} of $G_x \backslash G$. Therefore $\Pi(\mathcal{C})$ is homogeneous (resp. discrete) iff \mathcal{C} is homogeneous on G_x (resp. discrete on G_x). \square

From an m -scheme to a \mathcal{P} -scheme. Conversely, we could also construct a \mathcal{P} -scheme from an m -scheme on S . Here we need an additional assumption that $m \neq |S| - 1$ and G acts $\min\{|S|, m + 1/2\}$ -transitively on S .²

Lemma 2.8. *Assume $m \neq |S| - 1$ and G acts $\min\{|S|, m + 1/2\}$ -transitively on S . For $T, T' \subseteq S$ of cardinality at most m , we have $G_T = G_{T'}$ iff $T = T'$. And the normalizer $N_G(G_T)$ of G_T is the setwise stabilizer $G_{\{T\}}$.*

Proof. The assumption implies that set of elements in S fixed by G_T (resp. $G_{T'}$) is precisely T (resp. T'). So $G_T = G_{T'}$ implies $T = T'$. The other direction is trivial.

For $g \in G$, we have $gG_Tg^{-1} = G_{gT}$. So $g \in N_G(G_T)$ iff $G_T = G_{gT}$, which holds iff $T = {}^gT$ by the first part. So $N_G(G_T) = G_{\{T\}}$. \square

Definition 2.13. *Assume $m \neq |S| - 1$ and G acts $\min\{|S|, m + 1/2\}$ -transitively on S . Given an m -scheme $\Pi = \{P_1, \dots, P_m\}$ on S , define a \mathcal{P} -collection $\mathcal{C}(\Pi) = \{C_H : H \in \mathcal{P}\}$ as follows: For $H \in \mathcal{P}$, pick $T \subseteq S$ of cardinality $k \in [m]$ such that $H = G_T$. By Lemma 2.8, such a set T is unique. Pick $x = (x_1, \dots, x_k) \in S^{(k)}$ such that $T = \{x_1, \dots, x_k\}$. Then $G_x = G_T = H$ and we have the map $\lambda_x : S^{(k)} \rightarrow H \backslash G$. Define $C_H = \{\lambda_x(B) : B \in P_k\}$.*

Lemma 2.9. *$\mathcal{C}(\Pi)$ as defined above is independent of the choices of elements x and is an \mathcal{P} -scheme. It is symmetric (resp. antisymmetric, strongly antisymmetric) if Π is symmetric (resp. antisymmetric, strongly antisymmetric). And it is homogeneous on G_x (resp. discrete on G_x) for $x \in S$ iff Π is homogeneous (resp. discrete).*

Proof. Fix $H = G_T \in \mathcal{P}$ and we show that C_H does not depend on the choices of x . Consider two elements $x = (x_1, \dots, x_k), x' = (x'_1, \dots, x'_k) \in S^{(k)}$ such that $T = \{x_1, \dots, x_k\} = \{x'_1, \dots, x'_k\}$. Then there exists $\rho \in \text{Sym}(k)$ such that $c_\rho^k(x') = x$. We check that $\lambda_{x'} = \lambda_x \circ c_\rho^k$: For any $y = {}^g x' \in S^{(k)}$ where $g \in G$, we have $\lambda_{x'}(y) = Hg^{-1}$ and

$$\lambda_x \circ c_\rho^k(y) = \lambda_x \circ c_\rho^k({}^g x') = \lambda_x \left({}^g (c_\rho^k(x')) \right) = \lambda_x({}^g x) = Hg^{-1}$$

as desired. As Π is invariant, the map c_ρ^k sends blocks to blocks. Therefore $\{\lambda_x(B) : B \in P_k\} = \{\lambda_{x'}(B) : B \in P_k\}$. So the two elements x and x' define the same partition C_H .

²Recall that a group action of G on S is $(k + 1/2)$ -transitive if it is k -transitive, and in addition for all $T \subseteq S$ of cardinality k , either the G_T -orbit of every $x \in S - T$ contains more than one element, or $|S - T| = 1$.

Next we check that $\mathcal{C}(\Pi)$ is a \mathcal{P} -scheme. Consider a projection $\pi_{H,H'} : H \backslash G \rightarrow H' \backslash G$ where $H, H' \in \mathcal{P}$ and $H \subseteq H'$. Then there exist $T' \subseteq T \subseteq S$ such that $H = G_T, H' = G_{T'}$. We may assume $|T| = |T'| + 1$ by decomposing $\pi_{H,H'}$ into the composition of more projections if necessary. Let $k = |T|$ and pick $x = (x_1, \dots, x_k)$ such that $T = \{x_1, \dots, x_k\}$. Choose the unique $i \in [k]$ such that $x_i \notin T'$. Let $x' = \pi_i^k(x)$. Then $H = G_T = G_x$ and $H' = G_{T'} = G_{x'}$. We claim that the following diagram commutes:

$$\begin{array}{ccc} S^{(k)} & \xrightarrow{\pi_i^k} & S^{(k-1)} \\ \lambda_x \downarrow & & \downarrow \lambda_{x'} \\ H \backslash G & \xrightarrow{\pi_{H,H'}} & H' \backslash G. \end{array}$$

To see this, note that for any $y = g x \in S^{(k)}$ where $g \in G$, we have $\pi_{H,H'} \circ \lambda_x(y) = \pi_{H,H'}(H g^{-1}) = H' g^{-1}$, and

$$\lambda_{x'} \circ \pi_i^k(y) = \lambda_{x'} \circ \pi_i^k(g x) = \lambda_{x'}(g(\pi_i^k(x))) = \lambda_{x'}(g x') = H' g^{-1},$$

as desired. And $\lambda_x, \lambda_{x'}$ are bijections that send blocks to blocks. Compatibility and regularity of $\mathcal{C}(\Pi)$ then follow from those of Π .

Now consider a conjugation $c_{H,h} : H \backslash G \rightarrow H' \backslash G$ for $H \in \mathcal{P}$ and $h \in G$, where $H' = h H h^{-1}$. Choose $x \in S^{(k)}$ for some $k \in [m]$ such that $H = G_x$. Let $x' = h x$ so that $H' = G_{x'}$. Then the following diagram commutes:

$$\begin{array}{ccc} & S^{(k)} & \\ \lambda_x \swarrow & & \searrow \lambda_{x'} \\ H \backslash G & \xrightarrow{c_{H,h}} & H' \backslash G. \end{array}$$

To see this, note that for any $y = g x' \in S^{(k)}$ where $g \in G$, we have $\lambda_{x'}(y) = H' g^{-1}$ and

$$c_{H,h} \circ \lambda_x(y) = c_{H,h} \circ \lambda_x(g^h x) = c_{H,h}(H(g h)^{-1}) = H' g^{-1},$$

as desired. So $\mathcal{C}(\Pi)$ is invariant. Therefore $\mathcal{C}(\Pi)$ is a \mathcal{P} -scheme.

Now we prove the claim that strongly antisymmetry is preserved. Assume that a map $\tau : B_1 \rightarrow B_2$ between blocks $B_1, B_2 \in C_H$ for some $H = G_T \in \mathcal{P}$ is obtained by composing conjugations, projections and their inverses (restricted to blocks). Let $k = |T|$. By the two diagrams above, we can obtain a map $\tau' : B'_1 \rightarrow B'_2$ between blocks $B'_1, B'_2 \in P_k$ by composing maps of the form $c_g^i|_B, \pi_T^i|_B$, or $(\pi_T^i|_B)^{-1}$, such

that the following diagram commutes

$$\begin{array}{ccc} B'_1 & \xrightarrow{\tau'} & B'_2 \\ \lambda_x \downarrow & & \downarrow \lambda_{x'} \\ B_1 & \xrightarrow{\tau} & B_2 \end{array}$$

for some $x, x' \in S^{(k)}$. We showed in the beginning that there exists $\rho \in \text{Sym}(k)$ satisfying $\lambda_{x'} = \lambda_x \circ c_\rho^k$. By replacing τ' with $c_\rho^k \circ \tau'$, B'_2 with $c_\rho^k(B'_2)$, and $\lambda_{x'}$ with λ_x , we may assume $x = x'$. Then if $B_1 = B_2$ and τ is a nontrivial permutation of B_1 , we also know that $B'_1 = B'_2$ and τ' is a nontrivial permutation of B'_1 . Therefore, if Π is strongly antisymmetric, so is $\mathcal{C}(\Pi)$. The claim for antisymmetry is proved in the same way, except that we only consider maps τ arising from conjugations but not projections. And if $B_1 \neq B_2$ for such τ , we also get a map $\tau' : B'_1 \rightarrow B'_2$ arising from $c_{\tau_0}^k$ for some $\tau_0 \in \text{Sym}(k)$ such that $B'_1 \neq B'_2$. So symmetry is also preserved.

Finally, for any $x \in S$, the partition C_{G_x} of $G_x \backslash G$ is constructed using the bijection $\lambda_x : S \rightarrow G_x \backslash G$ and the partition P_1 . Therefore $\mathcal{C}(\Pi)$ is homogeneous on G_x (resp. discrete on G_x) iff Π is homogeneous (resp. discrete). \square

The maps $\mathcal{C} \mapsto \Pi(\mathcal{C})$ and $\Pi \mapsto \mathcal{C}(\Pi)$ are inverse to each other by construction. So Lemma 2.7 and Lemma 2.9 together establish the one-to-one correspondence between \mathcal{P} -schemes and m -schemes on S .

Theorem 2.1. *Suppose $m \neq |S| - 1$ and G is a finite group acting $\min\{|S|, m + 1/2\}$ -transitively on S , and $\mathcal{P} = \mathcal{P}_m$ is the system of stabilizers of depth m with respect to this action. The map $\mathcal{C} \mapsto \Pi(\mathcal{C})$ in Definition 2.12 is a one-to-one correspondence between \mathcal{P} -schemes and m -schemes on S , with the inverse map $\Pi \mapsto \mathcal{C}(\Pi)$ as defined in Definition 2.13. And $\Pi(\mathcal{C})$ is symmetric (resp. antisymmetric, strongly antisymmetric, homogeneous, discrete) iff \mathcal{C} is symmetric (resp. antisymmetric, strongly antisymmetric, homogeneous on G_x for $x \in S$, discrete on G_x for $x \in S$).*

Remark. The unpleasant assumption $m \neq |S| - 1$ in Theorem 2.1 is due to the technical fact that when $T \subseteq S$ has cardinality $|S| - 1$, the pointwise stabilizer G_T fixes not only T but also the whole set S . This assumption is needed if we want the correspondence in Theorem 2.1 to preserve antisymmetry and homogeneity: Suppose G is a permutation group on S and $|S| = \ell$ is a prime number. Then for $m = \ell - 1$, there exists an antisymmetric homogeneous m -scheme on S (see

Example 2.2 in Section 2.4). On the other hand, note that $\mathcal{P} = \mathcal{P}_m$ contains the trivial subgroup $G_S = \{e\}$. So by Lemma 2.4, all antisymmetric \mathcal{P} -schemes are discrete on G_x for any $x \in S$.

Matchings. The papers [IKS09; Aro+14] formulated the idea of [Evd94] with a notion called a *matching*. We use the more general definition in [Aro+14] (where it is called a *generalized matching*).

Definition 2.14 (matching). *Let $\Pi = \{P_1, \dots, P_m\}$ be an m -scheme on S . A block $B \in P_k$ for some $k \in [m]$ is called a matching of Π if there exist two distinct proper subsets T, T' of $[k]$ of the same cardinality such that $\pi_T^k(B) = \pi_{T'}^k(B)$ and $|B| = |\pi_T^k(B)|$.*

The work [IKS09; Aro+14] designed algorithms leading to m -schemes with no matching. Now we explain the connection between this property and strong antisymmetry of m -schemes.

Given a matching $B \in P_k$ of Π , let $T, T' \subseteq [k]$ be as in Definition 2.14 and let $k' := k - |T|$. Then $B' := \pi_T^k(B) = \pi_{T'}^k(B)$ is a block of $P_{k'}$. We have two maps $\pi_T^k|_B$ and $\pi_{T'}^k|_B$ from B to B' , both of which are bijective by the condition $|B| = |\pi_T^k(B)|$. Moreover $\pi_T^k|_B \neq \pi_{T'}^k|_B$ as they omit different subsets of coordinates and the k coordinates of elements in $S^{(k)}$ are all distinct. So $\pi_{T'}^k|_B \circ (\pi_T^k|_B)^{-1}$ is a nontrivial permutation of B' . We conclude:

Lemma 2.10. *A strongly antisymmetric m -scheme has no matching.*

So our definition of strong antisymmetry of m -schemes (or that of \mathcal{P} -schemes by Lemma 2.7) subsumes the property that no matching exists.

We will use strong antisymmetry instead of (non-existence of) matchings throughout this thesis. The advantage of this comes from transitivity: Suppose $x \in H \setminus G$ is sent to a different element $y \in H \setminus G$ by a map τ that is a composition of conjugations, projections and their inverses (restricted to blocks), then x and y belong to different blocks by strong antisymmetry. Suppose we also separate y from another element $z \in H \setminus G$ in the same way. Then since the set of maps we consider are closed under composition, we get a map sending x to z and hence are able to separate them as well. The analyses in Chapter 7 and Chapter 8 crucially exploit this property.

The connection of m -schemes with association schemes

As shown in [IKS09; Aro+14], m -schemes are closely related to the notion of *association schemes* [BI84].

Definition 2.15. *An association scheme on a finite set S is a partition P of $S \times S$ such that*

- $1_S := \{(x, x) : x \in S\}$ is a block of P ,
- for a block $g \in P$, the set $g^* := \{(y, x) : (x, y) \in g\}$ is also a block, and
- for every triple of blocks $g, g', g'' \in P$, there exists an integer $c_{g', g''}^g \geq 0$ such that for any $(x, y) \in g$, the number of $z \in S$ satisfying $(x, z) \in g'$ and $(z, y) \in g''$ is $c_{g', g''}^g$.

An association scheme P is *symmetric* if $g = g^*$ for all $g \in P$, and *antisymmetric* if $g \neq g^*$ for all $g \in P - \{1_S\}$. The integer $c_{g, g^*}^{1_S}$ is called the *valency* of g .

We can obtain a homogeneous 3-scheme from an association scheme and vice versa using the following constructions.

Definition 2.16. *For a finite set S and a partition P of $S \times S$ such that $1_S \in P$, define the partition P' of $S^{(3)}$ such that two elements $(x_1, x_2, x_3), (x'_1, x'_2, x'_3) \in S^{(3)}$ are in the same block of P' iff (x_i, x_j) and (x'_i, x'_j) are in the same block of P for all $1 \leq i, j \leq 3$. And define a 3-collection $\Pi(P) = \{P_1, P_2, P_3\}$ on S by choosing $P_1 = S$, $P_2 = P - \{1_S\}$, $P_3 = P'$. Conversely, given a 3-collection $\Pi = \{P_1, P_2, P_3\}$ on S , define a partition $P(\Pi)$ of $S \times S$ by $P(\Pi) := P_2 \cup \{1_S\}$*

Lemma 2.11 ([IKS09; Aro+14]). *If P is an association scheme, then $\Pi(P)$ is a homogeneous 3-scheme. Conversely, if Π is a homogeneous 3-scheme, then $P(\Pi)$ is an association scheme.*

By construction, this gives a one-to-one correspondence between association schemes on S and equivalent classes of homogeneous 3-schemes on S , where two homogeneous 3-schemes $\{P_1, P_2, P_3\}$ and $\{P'_1, P'_2, P'_3\}$ on S are said to be equivalent if $P_1 = P'_1$ and $P_2 = P'_2$.

In addition, we obviously have

Lemma 2.12. *If Π is symmetric (resp. antisymmetric), so is $P(\Pi)$.*

Next we discuss the relation between symmetry and antisymmetry of an association scheme P and those of $\Pi(P)$. Obviously, for $\Pi(P)$ to be symmetric (resp. antisymmetric), it is necessary that P is also symmetric (resp. antisymmetric). The exact condition is given as follows.

Lemma 2.13. *The 3-scheme $\Pi(P)$ is symmetric iff P is the trivial association scheme $\{1_S, S \times S - 1_S\}$. It is antisymmetric iff P is antisymmetric and $c_{g,g}^{g^*} = 0$ for all $g \in P - \{1_S\}$.*

Proof. The trivial association scheme $P = \{1_S, S \times S - 1_S\}$ gives rise to the 3-scheme $\Pi(P) = \{0_S, 0_{S(2)}, 0_{S(3)}\}$ which is symmetric. Suppose $P \neq \{1_S, S \times S - 1_S\}$. Let g_1 and g_2 be two distinct blocks in $P - \{1_S\}$. If $g_1 = g_2^*$ then P is not symmetric and hence neither is $\Pi(P)$. So assume $g_1 \neq g_2^*$. Fix $x \in S$. Then $(x, y) \in g_1$ and $(x, z) \in g_2$ for some $y, z \in S - \{x\}$, and $y \neq z$. Consider the element $t = (x, y, z) \in S^{(3)}$. Let $h = (1\ 2\ 3) \in \text{Sym}(3)$ so that ${}^h t = (z, x, y)$. We have $\pi_3^3(t) = (x, y) \in g_1$ and $\pi_3^3({}^h t) = (z, x) \in g_2^* \neq g_1$. By compatibility of $\Pi(P)$, the elements t and ${}^h t$ are in different blocks. So $\Pi(P)$ is not symmetric.

Suppose $\Pi(P)$ is antisymmetric, then so is P . We check that $c_{g,g}^{g^*} = 0$ for all $g \in P - \{1_S\}$. Assume to the contrary that $c_{g,g}^{g^*} > 0$ for some $g \in P - \{1_S\}$. Fix $(x, y) \in g^*$. Then there exists $z \in S$ such that $(x, z), (z, y) \in g$. Then for $t = (x, y, z) \in S^{(3)}$ and $h = (1\ 2\ 3) \in \text{Sym}(3)$, we have $\pi_i^3(t), \pi_i^3({}^h t) \in g^*$ for all $1 \leq i \leq 3$. It follows by definition that t and ${}^h t$ are in the same block, contradicting antisymmetry of $\Pi(P)$.

Conversely, suppose P is antisymmetric and $c_{g,g}^{g^*} = 0$ for all $g \in P - \{1_S\}$. To prove $\Pi(P)$ is antisymmetric, it suffices to show that for any $t = (x, y, z) \in S^{(3)}$ and $h \in \text{Sym}(3)$, the elements t and ${}^h t$ are in different blocks. First assume h is a transposition, e.g., $(1\ 2)$ (the other cases are symmetric). Then $\pi_3^3(t) = (x, y)$ and $\pi_3^3({}^h t) = (y, x)$ are in different blocks by antisymmetry of P , and the claim follows by compatibility of $\Pi(P)$. Next assume h is a 3-cycle, e.g., $(1\ 2\ 3)$ (the other case is symmetric), so that ${}^h t = (z, x, y)$. Let g be the block in $P - \{1_S\}$ containing (y, x) , so that $(x, y) \in g^*$. As $c_{g,g}^{g^*} = 0$, either (x, z) or (z, y) is not in g . If $(x, z) \notin g$, we have $\pi_3^3(t) = (x, y) \in g^*$ and $\pi_3^3({}^h t) = (z, x) \notin g^*$. If $(x, z) \in g$ but $(z, y) \notin g$, we have $\pi_3^3(t) = (x, z) \in g$ and $\pi_3^3({}^h t) = (z, y) \notin g$. In either case t and ${}^h t$ are in different blocks by compatibility of $\Pi(P)$. \square

Example 2.1. Let S be a finite dimensional vector space over a finite field \mathbb{F}_q where $\text{char}(\mathbb{F}_q) \notin \{2, 3\}$. Let P be the partition of $S \times S$ such that (x, y) and (x', y') are in

the same block iff $x - y = x' - y'$, which is an association scheme [BI84]. We check that P satisfies the condition of Lemma 2.13, and hence $\Pi(P)$ is antisymmetric. For any $(x, y) \notin 1_S$, we have $x - y \neq y - x$ since $x \neq y$ and $\text{char}(\mathbb{F}_q) \neq 2$, and therefore (x, y) and (y, x) are in different blocks. So P is antisymmetric. Then we check that $c_{g,g}^{g^*} = 0$ for all $g \in P - \{1_S\}$. Assume to the contrary that $c_{g,g}^{g^*} > 0$ for some $g \in P - \{1_S\}$. Fix $(x, y) \in g^*$ and choose $z \in S$ such that $(x, z), (z, y) \in g$. Then $x - z = z - y = y - x$, implying $3(x - z) = 0$. This is impossible since $x \neq z$ and $\text{char}(\mathbb{F}_q) \neq 3$.

The antisymmetric 3-scheme $\Pi(P)$ in Example 2.1 is not strongly antisymmetric: For any distinct $x, y \in S$, let $B \in P - \{1_S\}$ be the block containing $t = (x, y)$. Then $\pi_1^2|_B$ and $\pi_2^2|_B$ are bijections from B to S sending t to y and x , respectively. So $\pi_1^2|_B \circ (\pi_2^2|_B)^{-1}$ is a permutation of the unique block $S \in 0_S$ sending x to y .

We do not know any example of an association scheme P for which $\Pi(P)$ is strongly antisymmetric. The following lemma gives a sufficient condition for the existence of such an association scheme.

Lemma 2.14. *Suppose P is an antisymmetric association scheme satisfying (1) $c_{g,g}^{g^*} = 0$ for all $g \in P - \{1_S\}$, and (2) for all blocks $g \in P$ and $g', g'' \in P - \{1_S\}$, either $c_{g',g''}^g = 0$ or $c_{g',g''}^g > 1$. Then $\Pi(P)$ is strongly antisymmetric.*

Proof. By Lemma 2.13, the 3-scheme $\Pi(P)$ is antisymmetric. And (2) implies that none of the projections π_i^2 and π_i^3 are invertible even restricted to blocks of $S^{(2)}$ and $S^{(3)}$ respectively. Strong antisymmetry of $\Pi(P)$ then follows from antisymmetry. \square

In general, strongly antisymmetric 3-schemes do exist. See Example 2.4 in Section 2.5.

2.4 Orbit \mathcal{P} -schemes and m -schemes

An important family of m -schemes called *orbit schemes*, or what we call *orbit m -schemes*, was proposed and studied in [IKS09]. The blocks of such m -schemes are orbits of group actions.

Definition 2.17 (orbit m -scheme [IKS09]). *Given a finite set S , $m \in \mathbb{N}^+$, and a group $K \subseteq \text{Sym}(S)$ acting naturally on K , for each $k \in [m]$, define the partition P_k of $S^{(k)}$ to be the partition into K -orbits with respect to the diagonal action of*

K on $S^{(k)}$. The m -collection $\Pi = \{P_1, \dots, P_m\}$ is called the orbit m -scheme on S associated with the group K .

This is indeed an m -scheme:

Theorem 2.2 ([IKS09]). *The m -collection Π in Definition 2.17 is an m -scheme on S .*

We define orbit \mathcal{P} -schemes in a similar way, except that the subgroup K of $\text{Sym}(S)$ is now replaced with a subgroup of G , and the diagonal actions on $S^{(k)}$, $k \in [m]$ are replaced with the actions on right coset spaces by inverse right translation.

Definition 2.18 (orbit \mathcal{P} -scheme). *Let \mathcal{P} be a subgroup system over a finite group G , and let K be a subgroup of G . For $H \in \mathcal{P}$, define the partition C_H of $H \backslash G$ to be the partition into K -orbits, with respect to the action of K on $H \backslash G$ by inverse right translation. The \mathcal{P} -collection $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ is called the orbit \mathcal{P} -scheme associated with the group K .*

This construction indeed yields a \mathcal{P} -scheme:

Theorem 2.3. *The \mathcal{P} -collection \mathcal{C} in Definition 2.18 is a \mathcal{P} -scheme.*

Proof. Let K act on each right coset space $H \backslash G$ by inverse right translation. For $H, H' \in \mathcal{P}$ with $H \subseteq H'$, $g \in K$ and $x \in H \backslash G$, we have $\pi_{H,H'}(g x) = {}^g(\pi_{H,H'}(x))$ by Lemma 2.2. Therefore if $x, x' \in H \backslash G$ are in the same block of C_H (i.e., the same K -orbit of $H \backslash G$), then $\pi_{H,H'}(x)$ and $\pi_{H,H'}(x')$ are in the same block of $C_{H'}$ (i.e., the same K -orbit of $H' \backslash G$). So \mathcal{C} is compatible.

Similarly, for $H \in \mathcal{P}$, $h \in G$, $g \in K$ and $x \in H \backslash G$, we have $c_{H,h}(g x) = {}^g(c_{H,h}(x))$ by Lemma 2.2. Therefore if $x, x' \in H \backslash G$ are in the same block of C_H , then $c_{H,h}(x)$ and $c_{H,h}(x')$ are in the same block of $C_{gHg^{-1}}$. So \mathcal{C} is invariant.

For $H' \in \mathcal{P}$ and $y, y' \in H' \backslash G$ in the same block B of $C_{H'}$, choose $g \in K$ such that $y' = {}^g y$. As $g \in K$, we have ${}^g B = B$. For $H \in \mathcal{P}$ with $H \subseteq H'$ and $x \in H \backslash G$, we have $x \in B$ and $\pi_{H,H'}(x) = y$ iff ${}^g x \in {}^g B = B$ and $\pi_{H,H'}(g x) = {}^g(\pi_{H,H'}(x)) = {}^g y = y'$. So the map $x \mapsto {}^g x$ is a one-to-one correspondence between $B \cap \pi_{H,H'}^{-1}(y)$ and $B \cap \pi_{H,H'}^{-1}(y')$, and hence the two sets have the same cardinality. So \mathcal{C} is regular. \square

The connection between Definition 2.17 and Definition 2.18 is given by the following lemma.

Lemma 2.15. *For a finite set S , $m \in \mathbb{N}^+$, and a subgroup $K \subseteq \text{Sym}(S)$, let $\mathcal{P} = \mathcal{P}_m$ be the system of stabilizers of depth m with respect to the natural action of $\text{Sym}(S)$ on S , and let $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ be the orbit \mathcal{P} -scheme associated with K . Then the orbit m -scheme associated with K is exactly $\Pi(\mathcal{C})$ as defined in Definition 2.12.*

Proof. We may assume $m \leq |S|$. Let G be the symmetric group $\text{Sym}(S)$ acting naturally on S . Suppose $\Pi(\mathcal{C}) = \{P_1, \dots, P_m\}$ where P_k is a partition of $S^{(k)}$ for $k \in [m]$. By Definition 2.12, each partition P_k is given by $P_k = \{\lambda_x^{-1}(B) : B \in C_{G_x}\}$ for some $x = (x_1, \dots, x_k) \in S^{(k)}$, where $\lambda_x : S^{(k)} \rightarrow G_x \backslash G$ is an equivalence between the diagonal action of G on $S^{(k)}$ and the action on $G_x \backslash G$ by inverse right translation. It follows that P_k is the partition into K -orbits with respect to the diagonal action, since C_{G_x} is the partition into K -orbits with respect to the action by inverse right translation. \square

Antisymmetry of orbit m -schemes. We prove a simple and exact criterion for antisymmetry of orbit m -schemes.

Lemma 2.16. *The orbit m -scheme on S associated with $K \subseteq \text{Sym}(S)$ is antisymmetric iff the order of K is coprime to $1, 2, \dots, m$.*

Proof. Let $\Pi = \{P_1, \dots, P_m\}$ be the orbit m -scheme on S associated with K . Suppose the order of K is divisible by an integer k satisfying $1 < k \leq m$. We may assume that k is a prime integer. By Cauchy's theorem (see, e.g., [Lan02]), the group K contains an element g of order k . The element g , as a permutation of S , has at least one k -cycle $(x_1 x_2 \cdots x_k)$. Consider the element $x = (x_1, \dots, x_k) \in S^{(k)}$, and let B be the block of P_k containing x . By definition, the element ${}^g x = ({}^g x_1, \dots, {}^g x_k) = (x_2, \dots, x_k, x_1)$ is also in B . On the other hand, let $h = (1 \ 2 \ \cdots \ k)^{-1} \in \text{Sym}(k)$. The permutation c_h^k of $S^{(k)}$ sends $x = (x_1, \dots, x_k)$ to $y = (y_1, \dots, y_k)$ defined by $y_i = x_{h^{-1}i}$ for $i \in [k]$. So $c_h^k(x) = (x_2, \dots, x_k, x_1) \in B$. Therefore Π is not antisymmetric.

Conversely, assume Π is not antisymmetric. Then for some integer k satisfying $1 < k \leq \min\{m, |S|\}$, $h \in \text{Sym}(k) - \{e\}$, and some element $x = (x_1, \dots, x_k) \in S^{(k)}$ lying in a block B of P_k , we have $c_h^k(x) \in B$, i.e., $c_h^k(x) = {}^g x$ for some $g \in K$ with

respect to the diagonal action of K on $S^{(k)}$. As the permutation c_h^k of $S^{(k)}$ sends x to $y = (y_1, \dots, y_k)$ defined by $y_i = x_{h^{-1}_i}$ for $i \in [k]$, we see ${}^g x_i = x_{h^{-1}_i}$ for $i \in [k]$. Then g preserves the set $T := \{x_1, \dots, x_k\}$ and restricts to a nontrivial permutation $g|_T \in \text{Sym}(T)$ of T . Let e be the order of $g|_T$. Then e is not coprime to some integer t where $t \leq |T| \leq m$. The order of K is a multiple of the order of g , which is a multiple of e . So the order of K is not coprime to t either. \square

Example 2.2. Let S be a finite set satisfying $|S| > 1$. Let K be a subgroup of $\text{Sym}(S)$ generated by a single $|S|$ -cycle so that it acts regularly on S . Denote by ℓ the least prime factor of $|S|$. Let Π be the orbit m -scheme on S associated with K where m is an integer satisfying $1 \leq m < \ell$. Then Π is homogeneous since K acts transitively on S . The order of K is $|S|$, which is coprime to $1, \dots, \ell - 1$. So Π is also antisymmetric by Lemma 2.16 and the fact $m \leq \ell - 1$.

Upper bound of m for antisymmetric homogeneous m -schemes. Let S be a finite set satisfying $|S| > 1$, and let ℓ be the least prime factor of $|S|$. For $m \geq \ell$, the orbit m -schemes on S in Example 2.2 are still homogeneous but no longer antisymmetric. Indeed, an argument of Rónyai [Rón88] shows that for $m \geq \ell$, even general m -schemes on S cannot be both homogeneous and antisymmetric. This was reproduced in [IKS09] and we present it here.

Lemma 2.17 ([Rón88; IKS09]). *Let S be a finite set satisfying $|S| > 1$, and let ℓ be the least prime factor of $|S|$. There exists no antisymmetric homogeneous m -scheme on S for $m \geq \ell$.*

Proof. Assume to the contrary that such an m -scheme $\Pi = \{P_1, \dots, P_m\}$ exists. The group $\text{Sym}(\ell)$ acts on $S^{(\ell)}$ by ${}^g x = c_g^k(x)$. By antisymmetry of Π , this action induces a semiregular action on the set of blocks in P_ℓ . Let $B_1, \dots, B_k \in P_\ell$ be a complete set of representatives for the $\text{Sym}(\ell)$ -orbits, i.e., each orbit contains exactly one B_i . Then we have

$$\sum_{i=1}^k |B_i| = \frac{|S^{(\ell)}|}{|\text{Sym}(\ell)|} = \frac{|S|(|S| - 1) \cdots (|S| - \ell + 1)}{\ell!}$$

Let π be the projection from $S^{(\ell)}$ to S sending (x_1, \dots, x_ℓ) to x_1 . By regularity and homogeneity of Π , for each $i \in [k]$, the cardinality of $B_i \cap \pi^{-1}(y)$ is a constant $d_i \in \mathbb{N}^+$ independent of $y \in S$. Then

$$\sum_{i=1}^k d_i = \sum_{i=1}^k \frac{|B_i|}{|S|} = \frac{(|S| - 1) \cdots (|S| - \ell + 1)}{\ell!}.$$

As $|S|$ is a multiple of ℓ , none of the factors $|S| - 1, \dots, |S| - \ell + 1$ of the numerator is divisible by the prime number ℓ appeared in the denominator. This contradicts the integrality of $\sum_{i=1}^k d_i$. \square

The condition $m \geq \ell$ in Lemma 2.17 is tight, since Example 2.2 shows that antisymmetric homogeneous m -schemes exist for $m = \ell - 1$.

Rónyai's result can be extended to \mathcal{P} -schemes in the case that \mathcal{P} is a system of stabilizers with respect to a transitive group action.

Lemma 2.18. *Let G be a finite group acting transitively on a set S of cardinality $n > 1$. Let $\mathcal{P} = \mathcal{P}_m$ be the corresponding system of stabilizers of depth m for some $m \geq \ell$, where ℓ is the least prime factor of n . Then for any $x \in S$, there exists no antisymmetric \mathcal{P} -scheme that is homogeneous on G_x . In particular, $d'(G) < \ell$.*

Lemma 2.18 can be easily proven using a technique called the *induction of \mathcal{P} -schemes*, to be discussed in Chapter 6. It allows us to reduce to the case $G = \text{Sym}(S)$. The claim then follows immediately, since by Lemma 2.7, for $G = \text{Sym}(S)$, the existence of an antisymmetric \mathcal{P} -scheme homogeneous on G_x implies the existence of an antisymmetric homogeneous m -scheme on S , which contradicts Lemma 2.17. For now, we just provide a direct proof.

Proof of Lemma 2.18. Assume to the contrary that $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ is an antisymmetric \mathcal{P} -scheme that is homogeneous on G_x for some $x \in S$. As \mathcal{C} is invariant and G acts transitively on S (and hence all one-point stabilizers G_x are conjugate in G), we know \mathcal{C} is homogeneous on G_x for all $x \in S$.

Consider the set $S^{(\ell)}$ equipped with two actions: the diagonal action of G and the action of $\text{Sym}(\ell)$ permuting the ℓ coordinates. The latter action is defined by $g(x_1, \dots, x_\ell) = (x_{g^{-1}1}, \dots, x_{g^{-1}\ell})$ for $g \in \text{Sym}(\ell)$ and $(x_1, \dots, x_\ell) \in S^{(\ell)}$. Note that these two actions commute with each other and combine to an action of $G \times \text{Sym}(\ell)$ on $S^{(\ell)}$. For $z \in S^{(\ell)}$, we have ${}^g Gz = G{}^g z$ for all $g \in \text{Sym}(\ell)$ and hence the action of $\text{Sym}(\ell)$ permutes the G -orbits within the $(G \times \text{Sym}(\ell))$ -orbit $(G \times \text{Sym}(\ell))z$.

Now fix $z \in S^{(\ell)}$. We have the bijection $\lambda_z : Gz \rightarrow G_z \backslash G$ which is an equivalence between the action of G on the G -orbit Gz and the action on $G_z \backslash G$ by inverse right translation. We also have a semiregular action of $N_G(G_z)/G_z$ on $G_z \backslash G$ by left translation. This gives an injective group homomorphism $\phi : N_G(G_z)/G_z \hookrightarrow \text{Sym}(G_z \backslash G)$, and we denote its image by N . Then $|N| = |N_G(G_z)/G_z|$.

Let H be the subgroup of $\text{Sym}(\ell)$ fixing Gz setwisely, i.e., $H = \{g \in \text{Sym}(\ell) : {}^gGz = Gz\}$. The action of $H \subseteq \text{Sym}(\ell)$ on $S^{(\ell)}$ restricts to an action on Gz and hence we have a group homomorphism $H \rightarrow \text{Sym}(Gz)$. It is injective since elements in $Gz \subseteq S^{(\ell)}$ have distinct coordinates. Now, identifying Gz with $G_z \backslash G$ via λ_z , we have an action of H on $G_z \backslash G$ as well, defined by ${}^g\lambda_z(x) = \lambda_z({}^gx)$ for $x \in Gz$. This gives an injective group homomorphism $\phi' : H \hookrightarrow \text{Sym}(G_z \backslash G)$.

We claim that $\phi'(H) \subseteq N$. To see this, pick any $g \in H$. We have ${}^gG_z e = G_z h_0$ for some $h_0 \in G$, or equivalently ${}^gz = h_0^{-1}z$. Then for any $h \in G$, we have

$${}^gG_z h = {}^g\left(\lambda_z(h^{-1}z)\right) = \lambda_z\left({}^g(h^{-1}z)\right) = \lambda_z\left(h^{-1}({}^gz)\right) = \lambda_z\left((h_0h)^{-1}z\right) = G_z h_0 h.$$

In particular, for any $h \in G_z$, we have ${}^gG_z h = G_z h_0 h$ and other other hand ${}^gG_z h = {}^gG_z e = G_z h_0$. So $h_0 h h_0^{-1} \in G_z$. Therefore $h_0 \in N_G(G_z)$. Furthermore, note that $h_0 G_z \in N_G(G_z)/G_z$ sends any $G_z h \in G_z \backslash G$ to $G_z h_0 h = {}^gG_z h$ by left translation. So $\phi'(g) = \phi(h_0 G_z) \in N$. Therefore $\phi'(H) \subseteq N$, as desired.

By antisymmetry, the action of N on $G_z \backslash G$ induces a semiregular action on the set of blocks of C_{G_z} , which induces a semiregular action of $\phi'(H)$ on the set of blocks of C_{G_z} . Let $B_1, \dots, B_k \in C_{G_z}$ be a complete set of representatives for the $\phi'(H)$ -orbits. Then we have

$$\sum_{i=1}^k |B_i| = \frac{|G_z \backslash G|}{|\phi'(H)|} = \frac{|Gz|}{|H|}.$$

Choose $x \in S$ such that $G_z \subseteq G_x$. By regularity and homogeneity on G_x , for each $i \in [k]$, the cardinality of $B_i \cap \pi_{G_z, G_x}^{-1}(y)$ is a constant $d_i \in \mathbb{N}^+$ independent of $y \in G_x \backslash G$, and hence $|B_i|$ is a multiple of $|G_x \backslash G| = n$. Therefore $|Gz|$ is a multiple of $n \cdot |H|$.

By the orbit-stabilizer theorem, the number of G -orbits contained in $(G \times \text{Sym}(\ell))z$ is $|\text{Sym}(\ell)|/|H|$, and these G -orbits all have the same cardinality $|Gz|$. So

$$|(G \times \text{Sym}(\ell))z| = \frac{|\text{Sym}(\ell)|}{|H|} \cdot |Gz|,$$

which is a multiple of $n \cdot |\text{Sym}(\ell)| = n\ell!$ since $|Gz|$ is a multiple of $n \cdot |H|$. As this holds for arbitrary $z \in S^{(\ell)}$, we know $|S^{(\ell)}| = n(n-1) \cdots (n-\ell+1)$ is also a multiple of $n\ell!$. But this is not possible since $n-1, \dots, n-\ell+1$ are not divisible by the prime number ℓ . \square

2.5 Strongly antisymmetric homogeneous m -schemes for $m \leq 3$

In this section, we give examples of strongly antisymmetric homogeneous m -schemes on a finite set S where $|S| > 1$ and $m \in \{1, 2, 3\}$.

The case $m = 1$. For all finite sets S , there exists a unique homogeneous 1-scheme $\Pi = \{P_1\}$ on S , given by $P_1 = 0_S$. It is obviously antisymmetric since $\text{Sym}(1)$ is the trivial group. And it is also strongly antisymmetric since there exists no projection π_i^k for $m = 1$.

The case $m = 2$. We discuss the following explicit construction of orbit 2-schemes.

Example 2.3. Let q be a prime power of the form $q = 4k + 3$ for some $k \in \mathbb{N}$.³ The multiplicative group \mathbb{F}_q^\times is a cyclic group of order $4k + 2$. Denote by $\chi_2 : \mathbb{F}_q^\times \rightarrow \mathbb{C}$ the unique nontrivial quadratic character of \mathbb{F}_q^\times , which sends quadratic residues to 1 and non-residues to -1 . Its kernel $\text{Ker}(\chi_2)$ is the unique subgroup of \mathbb{F}_q^\times of index two. For $u \in \mathbb{F}_q^\times$ and $v \in \mathbb{F}_q$, denote by $\phi_{u,v}$ the affine linear transformation of \mathbb{F}_q sending $x \in \mathbb{F}_q$ to $ux + v$. Define K by

$$K := \{\phi_{u,v} : u \in \text{Ker}(\chi_2), v \in \mathbb{F}_q\}.$$

Then K is a subgroup of $\text{Sym}(\mathbb{F}_q)$.⁴ Let $\Pi = \{P_1, P_2\}$ be the orbit 2-scheme on \mathbb{F}_q associated with the subgroup K .

The partitions P_1 and P_2 are given as follows: as K acts transitively on \mathbb{F}_q , we have $P_1 = 0_{\mathbb{F}_q}$ and Π is homogeneous. For $(a, b) \in \mathbb{F}_q^{(2)}$, we have $\phi_{1,b}(a - b, 0) = (a, b)$ and $\phi_{1,b} \in K$, and hence (a, b) and $(a - b, 0)$ are in the same block of P_2 . Two elements $(c, 0), (d, 0) \in \mathbb{F}_q^{(2)}$ are in the same block iff $^g c = d$ for some $g \in K_0$, where K_0 is the stabilizer of $0 \in \mathbb{F}_q$. As $K_0 = \{\phi_{u,0} : u \in \text{Ker}(\chi_2)\}$, we see that $(c, 0)$ and $(d, 0)$ are in the same block iff $\chi_2(c) = \chi_2(d)$. We conclude that P_2 contains two blocks B_{+1} and B_{-1} , where

$$B_s = \{(a, b) \in \mathbb{F}_q^{(2)} : \chi_2(a - b) = s\}$$

for $s = \pm 1$.

³In particular, we may choose q to be a prime number. By Dirichlet's theorem on arithmetic progressions [Neu99], there exist infinitely many prime numbers of the form $4k + 3$.

⁴The group K is also a subgroup of the *general affine group* $\text{AGL}_1(q)$ and is isomorphic to a semidirect product $\mathbb{F}_q \rtimes \text{Ker}(\chi_2)$.

The order of K is $q(q-1)/2 = (4k+3)(2k+1)$ which is odd. So Π is antisymmetric by Lemma 2.16. For every $y \in \mathbb{F}_q$, the number of elements in B_{+1} (or B_{-1}) mapped to y by the projection π_1^2 (or π_2^2) is $(q-1)/2$, which is greater than one when $q > 3$. Therefore when $q > 3$, the two projections π_1^2 and π_2^2 restricted to B_1 (or B_2) are not invertible, and hence Π is strongly antisymmetric. We conclude:

Lemma 2.19. *The orbit 2-scheme Π in Example 2.3 is homogeneous and antisymmetric. It is strongly antisymmetric when $q > 3$.*

We remark that the partition $P := P_2 \cup \{1_{\mathbb{F}_q}\}$ of $\mathbb{F}_q \times \mathbb{F}_q$ (where $1_{\mathbb{F}_q} = \{(a, a) : a \in \mathbb{F}_q\}$) is actually an antisymmetric association scheme on \mathbb{F}_q . It is known as an association scheme of *Paley tournaments* [ER63; BI84; BCN89], or more generally a *cyclotomic scheme* [BCN89].

Recall that for an association scheme P on a set S , blocks $g, g', g'' \in P$, and $(x, y) \in g$, we use $c_{g',g''}^g$ to denote the number of $z \in S$ satisfying $(x, z) \in g'$ and $(z, y) \in g''$. When P is antisymmetric and has only three blocks, the quantities $c_{g',g''}^g$ only depend on n .⁵ We state it formally for the cases $g, g', g'' \neq 1_S$.

Lemma 2.20. *Let P be an antisymmetric association scheme on a set S of cardinality n containing only three blocks $1_S, g$ and g^* . Then for $u, v, w \in \{g, g^*\}$, we have*

$$c_{v,w}^u = \begin{cases} (n+1)/4 & \text{if } u^* = v = w, \\ (n-3)/4 & \text{otherwise.} \end{cases}$$

Proof. From the basic properties of association schemes, we have $c_{g,g}^{g^*} = c_{g^*,g^*}^g$, $c_{g,g}^g = c_{g,g^*}^g = c_{g^*,g}^g = c_{g^*,g^*}^{g^*} = c_{g^*,g}^{g^*} = c_{g,g^*}^{g^*}$, and $\sum_{w \in P} c_{v,w}^u = (n-1)/2$ for $u, v \in \{g, g^*\}$.⁶ Also note that $c_{v,1_S}^u$ equals one when $u = v$ and zero otherwise.

The claim then follows by simple calculations. \square

In particular, Lemma 2.20 applies to the association scheme $P = P_2 \cup \{1_{\mathbb{F}_q}\}$ above. This is used in the next example for the proof of strong antisymmetry.

⁵This is a folklore result. Such an association scheme is equivalent to a *doubly regular tournament*. See, e.g., [RB72].

⁶See, e.g., [BI84, Section II.2, Proposition 2.2] and note that g, g^* have the same valency $(n-1)/2$.

The case $m = 3$. We have noted that for P_2 as defined in Example 2.3, the partition $P = P_2 \cup \{1_{\mathbb{F}_q}\}$ of $\mathbb{F}_q \times \mathbb{F}_q$ is an antisymmetric association scheme on \mathbb{F}_q . Thus by Lemma 2.11, we have a homogeneous 3-scheme $\Pi(P)$. Unfortunately, $\Pi(P)$ is not necessarily antisymmetric: there may exist distinct elements $a, b, c \in \mathbb{F}_q$ such that $\chi_2(a - b) = \chi_2(b - c) = \chi_2(c - a)$, and the block containing $(a, b, c) \in \mathbb{F}_q^{(3)}$ is preserved by the 3-cycles in $\text{Sym}(3)$.

However, it is possible to modify $\Pi(P)$ to get an explicit construction of strongly antisymmetric homogeneous 3-schemes. The idea is to use a nontrivial cubic character besides the quadratic character χ_2 .

Example 2.4. Let q be a prime power of the form $36k + 11$ or $36k + 23$ for some $k \in \mathbb{N}$.⁷ The congruence is chosen so that $q - 1$ is divisible by 2 but not by 3 or 4, and $q^2 - 1$ is divisible by 3 but not by 9. In particular, the condition $q \equiv 3 \pmod{4}$ in Example 2.3 still holds. Define a 3-collection $\Pi = \{P_1, P_2, P_3\}$ on \mathbb{F}_q as follows: P_1 and P_2 are constructed in the same way as in Example 2.3, i.e., $P_1 = 0_{\mathbb{F}_q}$ and P_2 contains two blocks, B_{+1} and B_{-1} , where

$$B_s = \{(a, b) \in \mathbb{F}_q^{(2)} : \chi_2(a - b) = s\}$$

for $s = \pm 1$, and $\chi_2 : \mathbb{F}_q^\times \rightarrow \mathbb{C}$ is the unique nontrivial quadratic character of \mathbb{F}_q^\times .

To construct P_3 , we consider the quadratic extension \mathbb{F}_{q^2} of \mathbb{F}_q . Its multiplicative group $\mathbb{F}_{q^2}^\times$ is a cyclic group of order $q^2 - 1$ which is divisible by 3. Choose a nontrivial cubic character $\chi_3 : \mathbb{F}_{q^2}^\times \rightarrow \mathbb{C}$. Let ω be a primitive third root of unity in \mathbb{F}_{q^2} so that $1 + \omega + \omega^2 = 0$. For $(a, b, c) \in \mathbb{F}_q^{(3)}$, we have $a + \omega b + \omega^2 c = (a - c) + \omega(b - c)$ which is nonzero since $\omega \notin \mathbb{F}_q$. So $a + \omega b + \omega^2 c \in \mathbb{F}_{q^2}^\times$. We define a function s on $\mathbb{F}_q^{(3)}$ by

$$s(a, b, c) := \begin{cases} \chi_3(a + \omega b + \omega^2 c) & \text{if } \chi_2(a - b) = \chi_2(b - c) = \chi_2(c - a), \\ 1 & \text{otherwise.} \end{cases}$$

For $(a, b, c) \in \mathbb{F}_q^{(3)}$, call the quadruple

$$(\chi_2(a - b), \chi_2(b - c), \chi_2(c - a), s(a, b, c))$$

the *signature* of (a, b, c) . Choose the partition P_3 of $\mathbb{F}_q^{(3)}$ such that two triples $(a, b, c), (a', b', c') \in \mathbb{F}_q^{(3)}$ are in the same block iff they have the same signature.

⁷Again, by Dirichlet's theorem on arithmetic progressions [Neu99], there exist infinitely many such q .

Lemma 2.21. *The 3-collection Π in Example 2.4 is an antisymmetric homogeneous 3-scheme on \mathbb{F}_q . It is strongly antisymmetric when $q > 11$.*

Proof. We first check that Π is an antisymmetric 3-scheme.

For compatibility, we need to verify that if $(a, b, c), (a', b', c') \in \mathbb{F}_q^{(3)}$ are in the same block of P_3 , then their images under π_i^3 are in the same block of P_2 , $i = 1, 2, 3$. This follows by construction.

For invariance and antisymmetry, we need to show that for any $g \in \text{Sym}(3) - \{e\}$, the signature of $(a, b, c) \in \mathbb{F}_q^{(3)}$ determines that of ${}^g(a, b, c)$, and they are different. We note that $\chi_2(-1) = -1$ since $|\mathbb{F}_q^\times| = q - 1$ is not divisible by 4, and $\chi_3(\omega)$ is a primitive third root of unity in \mathbb{C} since $|\mathbb{F}_{q^2}^\times| = q^2 - 1$ is not divisible by 9.

Suppose g is a transposition, e.g., the one sending (a, b, c) to (b, a, c) (the other cases are symmetric). Then the signature of ${}^g(a, b, c)$ is

$$\begin{aligned} & (\chi_2(b - a), \chi_2(a - c), \chi_2(c - b), s(b, a, c)) \\ &= (-\chi_2(a - b), -\chi_2(c - a), -\chi_2(b - c), s(b, a, c)). \end{aligned}$$

When $\chi_2(a - b), \chi_2(b - c), \chi_2(c - a)$ are not all equal, we have $s(b, a, c) = 1$. Otherwise

$$s(b, a, c) = \chi_3(b + \omega a + \omega^2 c) = \chi_3(\omega) \chi_3(a + \omega^{-1} b + \omega^{-2} c).$$

The automorphism $x \mapsto x^q$ of \mathbb{F}_{q^2} fixes $a, b, c \in \mathbb{F}_q$ and exchanges ω with ω^{-1} . So $\chi_3(a + \omega^{-1} b + \omega^{-2} c) = \chi_3((a + \omega b + \omega^2 c)^q) = \chi_3^q(a + \omega b + \omega^2 c)$. We see that in this case, the signature of $(a, b, c) \in \mathbb{F}_q^{(3)}$ determines that of ${}^g(a, b, c)$. And they are different since $\chi_2(b - a) = -\chi_2(a - b) \neq \chi_2(a - b)$.

Suppose g is a 3-cycle, e.g., the one sending (a, b, c) to (b, c, a) (the other case is symmetric). Then the signature of ${}^g(a, b, c)$ is $(\chi_2(b - c), \chi_2(c - a), \chi_2(a - b), s(b, c, a))$. When $\chi_2(a - b), \chi_2(b - c), \chi_2(c - a)$ are not all equal, we have $s(b, c, a) = 1$. Otherwise

$$s(b, c, a) = \chi_3(b + \omega c + \omega^2 a) = \chi_3(\omega^2) \chi_3(a + \omega b + \omega^2 c) \neq s(a, b, c).$$

So again the signature of ${}^g(a, b, c)$ is determined by and different from that of (a, b, c) .

To prove regularity, let K be the subgroup $\{\phi_{u,v} : u \in \text{Ker}(\chi_2), v \in \mathbb{F}_q\}$ of $\text{Sym}(\mathbb{F}_q)$ as in Example 2.3, and let $\Pi' = \{P'_1, P'_2, P'_3\}$ be the orbit 3-scheme on \mathbb{F}_q associated with K . Then $P_1 = P'_1$ and $P_2 = P'_2$. We claim that P_3 is a coarsening of P'_3 , i.e.,

each block B of P_3 is a disjoint union of a collection I of blocks in P'_3 . Assume the claim holds. Then for such a block B , an element $y \in \mathbb{F}_q^{(2)}$, and a projection π_i^3 , we have

$$|B \cap (\pi_i^3)^{-1}(y)| = \sum_{B' \in I} |B' \cap (\pi_i^3)^{-1}(y)|.$$

As Π' is regular, it follows that Π is also regular. So it remains to prove the claim.

The blocks of P'_3 are K -orbits. So it suffices to show that for $(a, b, c) \in \mathbb{F}_q^{(3)}$ and $\phi_{u,v} \in K$, the elements (a, b, c) and $\phi_{u,v}(a, b, c) = (ua + v, ub + v, uc + v)$ have the same signature. We have

$$\chi_2((ua + v) - (ub + v)) = \chi_2(u)\chi_2(a - b) = \chi_2(a - b)$$

since $u \in (\mathbb{F}_q^\times)^2$. Similarly $\chi_2((ub + v) - (uc + v)) = \chi_2(b - c)$ and $\chi_2((uc + v) - (ua + v)) = \chi_2(c - a)$. Also note that \mathbb{F}_q^\times is contained in the kernel of χ_3 ,⁸ and hence $\chi_3(u) = 1$. Therefore

$$\begin{aligned} & \chi_3((ua + v) + \omega(ub + v) + \omega^2(uc + v)) \\ &= \chi_3(u(a + \omega b + \omega^2 c) + v(1 + \omega + \omega^2)) \\ &= \chi_3(u(a + \omega b + \omega^2 c)) \\ &= \chi_3(a + \omega b + \omega^2 c) \end{aligned}$$

and hence $s(a, b, c) = s(ua + v, ub + v, uc + v)$, as desired.

Homogeneity holds since $P_1 = 0_{\mathbb{F}_q}$. Next we show that Π is strongly antisymmetric when $q > 11$. To prove this, it suffices to show that the projections π_i^2 and π_i^3 are not invertible even restricted to each block. For π_i^2 this holds when $q > 3$, as shown in the proof of Lemma 2.19. For π_i^3 we only need to check that the cardinalities of blocks of P_3 are greater than the cardinality $q(q-1)/2$ of blocks of P_2 . Let (a, b, c) be an element of $\mathbb{F}_q^{(3)}$ and let B be the block of P_3 containing it. Let (u, v, w, t) be the signature of B . By Lemma 2.20, if u, v, w are not all equal, the cardinality of B is $(q(q-1)/2)((q-3)/4) > q(q-1)/2$. If $u = v = w$, the block B and two other blocks, whose signatures are $(u, v, w, \chi_3(\omega)t)$ and $(u, v, w, \chi_3^2(\omega)t)$ respectively, are permuted by 3-cycles in $\text{Sym}(3)$, and their disjoint union has cardinality $(q(q-1)/2)((q+1)/4)$ by Lemma 2.20. So

$$|B| = \frac{1}{3} \cdot \frac{q(q-1)}{2} \cdot \frac{q+1}{4} > \frac{q(q-1)}{2}$$

as desired. □

⁸Otherwise the intersection of \mathbb{F}_q^\times with the kernel has order $(q-1)/3$, which is impossible as 3 does not divide $q-1$.

Unlike Example 2.19, the 3-schemes constructed in Example 2.4 are not orbit m -schemes. In fact, we prove in Theorem 6.6 later that no strongly antisymmetric homogeneous orbit m -schemes on S exist if $|S| > 1$ and $m \geq 3$. It strengthens the result in [IKS09] that no such m -schemes exist for $m \geq 4$.

For $m \geq 4$, there are no known examples of strongly antisymmetric homogeneous m -schemes on S (where $|S| > 1$), even for general m -schemes. It is conjectured in [IKS09] that such m -schemes do not exist for $m \geq C$ where C is an absolute constant. An affirmative solution to this conjecture would imply a polynomial-time deterministic factoring algorithm under GRH. See Theorem 6.2. Currently the best known upper bound for m is $O(\log |S| + 1)$ [Evd94; IKS09; Gua09; Aro13]. See Theorem 7.1.

Chapter 3

THE \mathcal{P} -SCHEME ALGORITHM

In this chapter, we present a generic deterministic factoring algorithm called the *\mathcal{P} -scheme algorithm*, based on the notion of \mathcal{P} -schemes introduced in Chapter 2.

A univariate polynomial over a finite field is said to be *square-free* if it has no repeated factors, and *completely reducible* over \mathbb{F}_q if it factorizes into linear factors over \mathbb{F}_q . For simplicity, the algorithm in this chapter assumes that the input polynomial satisfies the following condition:

Condition 3.1. *The input polynomial is defined over a prime field \mathbb{F}_p . In addition, it is square-free and completely reducible over \mathbb{F}_p .*

This assumption is commonly made in the literature (see, e.g., [Rón88; Evd94; CH00; IKS09; Aro13; Aro+14]) and is justified by standard reductions [Ber70; Yun76; Knu98]. Specifically, Berlekamp [Ber70] reduced the problem of completely factoring an arbitrary polynomial over a finite field to the problem of finding roots of certain other polynomials in \mathbb{F}_p . The latter problem further reduces to the problem of completely factoring polynomials satisfying Condition 3.1 by the technique of *square-free factorization* [Yun76; Knu98]. Alternatively, we develop an algorithm that works for arbitrary polynomials over finite fields in Chapter 5 without using these reductions.

Overview of the \mathcal{P} -scheme algorithm

The \mathcal{P} -scheme algorithm consists of three parts:

1. reducing to the problem of computing an “idempotent decomposition” of a certain ring,
2. computing idempotent decompositions of rings associated with a poset of number fields,
3. constructing the poset of number fields used in the previous part.

Now we elaborate on each part.

Reduction to computing an idempotent decomposition. It is well known that computing a factorization of f is equivalent to finding zero divisors of the ring $\mathbb{F}_p[X]/(f(X))$ [Rón88; Evd94; IKS09]. We focus on special zero divisors called *idempotent elements* or simply *idempotents*,¹ i.e., those elements x satisfying $x^2 = x$. Two idempotents x, y are said to be *orthogonal* if $xy = 0$. It can be shown that the problem of factoring f reduces to decomposing the unity of the ring $\mathbb{F}_p[X]/(f(X))$ into a sum of nonzero mutually orthogonal idempotent elements, called an *idempotent decomposition*.

Definition 3.1. An idempotent decomposition of a ring R is a set I of nonzero mutually orthogonal idempotent elements of R satisfying $\sum_{x \in I} x = 1$.

On the other hand, recall that our algorithm uses a lifted polynomial $\tilde{f}(X) \in \mathbb{Z}[X]$ of f , as mentioned in the introduction. Furthermore, we may assume \tilde{f} is an irreducible lifted polynomial (see Definition 1.1) by running the factoring algorithm for rational polynomials [LLL82] to factorize \tilde{f} into the irreducible factors over \mathbb{Q} . See Section 3.9 for more discussion. The polynomial \tilde{f} defines a number field $F := \mathbb{Q}[X]/(\tilde{f}(X))$. We show that, since f is square-free and completely reducible over \mathbb{F}_p , the ring $\mathbb{F}_p[X]/(f(X))$ is naturally isomorphic to $\bar{\mathcal{O}}_F := \mathcal{O}_F/p\mathcal{O}_F$, where \mathcal{O}_F is the *ring of integers* of the field F . Therefore the problem reduces to that of computing an idempotent decomposition of the ring $\bar{\mathcal{O}}_F$.

Computing idempotent decompositions for a poset of number fields. Denote by L the splitting field of \tilde{f} over \mathbb{Q} and G the Galois group of \tilde{f} over \mathbb{Q} , i.e., $G = \text{Gal}(L/\mathbb{Q})$. Conceptually, replacing $\mathbb{F}_p[X]/(f(X))$ with $\bar{\mathcal{O}}_F$ allows us to use the information provided by the Galois group G . By the work of Rónyai [Rón92], a zero divisor $a \neq 0$ (or, in our language, an idempotent decomposition) of $\bar{\mathcal{O}}_F$ can be found efficiently if an efficiently computable nontrivial automorphism of the ring $\bar{\mathcal{O}}_F$ is given. The Galois group G naturally provides automorphisms of $\bar{\mathcal{O}}_F$, at least when F is Galois over \mathbb{Q} . Moreover, these automorphisms can be efficiently computed thanks to the efficient polynomial factoring algorithms for number fields [Len83; Lan85]. Using this idea, Rónyai [Rón92] gave a polynomial-time factoring algorithm for the case that F is Galois over \mathbb{Q} .

When F is not Galois over \mathbb{Q} , not every automorphism in G restricts to an automorphism of F or $\bar{\mathcal{O}}_F$. One of our key observations is that F may still admit a

¹Strictly speaking, we need to exclude the unity of the ring which is the only idempotent element that is not a zero divisor.




Figure 3.1: The tower of fields and Galois groups. Denote by L the splitting field of \tilde{f} over \mathbb{Q} and regard F as a subfield of L .

nontrivial automorphism group, from which we can compute a partial factorization of f . Indeed, we regard F as a subfield of L and let H be the subgroup of G fixing F . Then the automorphism group of F is isomorphic to $N_G(H)/H$. The corresponding fixed subfield $F' = F^{N_G(H)/H}$ is the smallest subfield of F such that F/F' is Galois. See Figure 3.1 for an illustration.

In the worst case, we may have $N_G(H) = H$ and then the automorphism group of F is trivial. However, an extension K of F may still have a nontrivial automorphism group, and hence a nontrivial idempotent decomposition may be obtained for $\bar{\mathcal{O}}_K := \mathcal{O}_K/p\mathcal{O}_K$ instead of $\bar{\mathcal{O}}_F$, where \mathcal{O}_K is the ring of integers of K . For example, suppose G is the symmetric group $\text{Sym}(n)$ permuting the n roots of \tilde{f} . We identify F with $\mathbb{Q}(\alpha)$ for some root α of \tilde{f} , and then H is the stabilizer G_α . Let β be a root of \tilde{f} different from α . Then the automorphism group of $K = F(\beta) = \mathbb{Q}(\alpha, \beta)$ is $N_G(G_{\alpha,\beta})/G_{\alpha,\beta}$, which is nontrivial as $N_G(G_{\alpha,\beta})$ contains the permutations swapping α and β . Another example is the case that K equals the splitting field L of \tilde{f} . In this case, the automorphism group of K is just G .

Motivated by the above observation, we design the algorithm so that it computes idempotent decompositions not only for the number field F , but also simultaneously for a poset of subfields of L . Moreover, we compute homomorphisms between these fields, which induce homomorphisms between the rings $\bar{\mathcal{O}}_K$. Using these homomorphisms, we show that the idempotent decompositions can be properly refined, unless some consistency constraints between them are satisfied.

The connection with \mathcal{P} -schemes is as follows: by Galois theory, the poset of subfields used by the algorithm corresponds to a poset \mathcal{P} of subgroups of G . Suppose a field K in the former poset is associated with a subgroup H . It can be shown that an idempotent decomposition of $\bar{\mathcal{O}}_K$ corresponds to a partition of the coset space $H \backslash G$. These partitions for various $H \in \mathcal{P}$ altogether form a \mathcal{P} -collection. Then the consistency constraints between the idempotent decompositions are just the defining properties of \mathcal{P} -schemes in disguise, i.e. compatibility, regularity, and invariance. In addition, we incorporate in our algorithm Rónyai's technique [Rón92] as mentioned above as well as its extension by Evdokimov [Evd94]. They are characterized by antisymmetry and strongly antisymmetry of \mathcal{P} -schemes respectively.

The main part of the algorithm has the following structure: it constructs the rings $\bar{\mathcal{O}}_K$ and the homomorphisms between them, and then maintains the idempotent decompositions of these rings and iteratively refines them. Each time it calls a subroutine corresponding to some property of \mathcal{P} -schemes in attempt to obtain a refinement. Either the property is already satisfied, or strictly finer idempotent decompositions are obtained by the subroutine. The algorithm terminates when the decompositions cannot be properly refined any more, in which case we are guaranteed to have a strongly antisymmetric \mathcal{P} -scheme. This gives the following result.

Theorem 3.1 (informal). *Under GRH, there exists a deterministic algorithm that given a poset \mathcal{P}^\sharp of subfields of L corresponding to a poset \mathcal{P} of subgroups of G , outputs idempotent decompositions of $\bar{\mathcal{O}}_K$ for $K \in \mathcal{P}^\sharp$ corresponding to a strongly antisymmetric \mathcal{P} -scheme. The running time is polynomial in the size of the input.*

Suppose F is in the poset \mathcal{P}^\sharp , corresponding to a group $H \in \mathcal{P}$. Then in the strongly antisymmetric \mathcal{P} -scheme produced by the algorithm, the partition of $H \backslash G$ translates into an idempotent decomposition of the ring $\bar{\mathcal{O}}_F$. In particular, it follows from the reduction in the first part of the algorithm that if all strongly antisymmetric \mathcal{P} -schemes are discrete (resp. inhomogeneous) on H , then we always obtain the complete factorization (resp. a proper factorization) of f .

Constructing a collection of number fields. Theorem 3.1 is a generic result, as we may feed it any poset \mathcal{P}^\sharp of subfields of L and get a strongly antisymmetric \mathcal{P} -scheme, where \mathcal{P} is the corresponding poset of subgroups of G . To obtain an actual factoring algorithm, we need to construct such a poset. More precisely, we construct a collection \mathcal{F} of number fields that are representatives of isomorphism

classes of those in \mathcal{P}^\sharp , i.e., isomorphic fields in \mathcal{P}^\sharp are represented by the same element in \mathcal{F} . The posets \mathcal{P} and \mathcal{P}^\sharp are determined once \mathcal{F} is given.

Let H be the subgroup of G fixing F . The collection \mathcal{F} of number fields should satisfy the following two constraints: (1) \mathcal{F} contains the field F , so that we can convert the partition on $H \setminus G$ in the \mathcal{P} -scheme into a factorization of f , and (2) all strongly antisymmetric \mathcal{P} -schemes are discrete (resp. inhomogeneous) on H , so that the algorithm always produces the complete factorization (resp. a proper factorization) of f . In addition, we want to bound the running time spent in constructing the fields in \mathcal{P}^\sharp , which controls the running time of the whole algorithm.

We give various settings of \mathcal{F} in which the two constraints above are satisfied. One of them is to choose $\mathcal{F} = \{F, L\}$, where L is the splitting field of \tilde{f} . In another setting, we choose \mathcal{F} so that \mathcal{P} is a system of stabilizers of depth m for sufficiently large $m \in \mathbb{N}^+$. They lead to factoring algorithms with various running time.

For simplicity, we only state the results that \mathcal{F} can be constructed in certain amount of time (see Section 3.8). The proofs are deferred to Chapter 4, where we give a more comprehensive investigation on the problem of constructing number fields.

Summary. The actual factoring algorithm combines the three parts above in the opposite order: we first construct a collection \mathcal{F} of number fields which determines the posets \mathcal{P} and \mathcal{P}^\sharp . Then we run the algorithm in Theorem 3.1 to obtain a collection of idempotent decompositions corresponding to a strongly antisymmetric \mathcal{P} -scheme. Finally we extract a factorization of f from the idempotent decomposition of $\bar{\mathcal{O}}_F$. This yields the main result of this chapter:

Theorem 3.2 (informal). *Suppose there exists a deterministic algorithm that given a polynomial $g(X) \in \mathbb{Z}[X]$ irreducible over \mathbb{Q} , constructs in time $T(g)$ a collection \mathcal{F} of subfields of the splitting field L of g over \mathbb{Q} such that*

- $F = \mathbb{Q}[X]/(g(X))$ is in \mathcal{F} , and
- all strongly antisymmetric \mathcal{P} -schemes are discrete (resp. inhomogeneous) on $\text{Gal}(L/F) \in \mathcal{P}$, where \mathcal{P} is the subgroup system associated with \mathcal{F} .

Then under GRH, there exists a deterministic algorithm that given $f(X) \in \mathbb{F}_p[X]$ satisfying Condition 3.1 and an irreducible lifted polynomial $\tilde{f}(X) \in \mathbb{Z}[X]$ of f , outputs the complete factorization (resp. a proper factorization) of f over \mathbb{F}_p in time polynomial in $T(\tilde{f})$ and the size of the input.

We show that many results achieved by known factoring algorithms [Hua91a; Hua91b; Rón88; Rón92; Evd94; IKS09] can be derived from Theorem 3.2. Thus the \mathcal{P} -scheme algorithm provides a unifying approach to polynomial factoring over finite fields.

Outline of the chapter. Notations and mathematical preliminaries are given in Section 3.1, and algorithmic preliminaries are given in Section 3.2. We reduce the problem of factoring f to that of computing an idempotent decomposition of $\bar{\mathcal{O}}_F$ in Section 3.3. In Section 3.4, we give the main body of the algorithm that computes idempotent decompositions corresponding to a strongly antisymmetric \mathcal{P} -scheme, and use it to prove Theorem 3.1. The next three sections (Section 3.5, 3.6 and 3.7) describe three subroutines used by this algorithm. In Section 3.8 we state some results on constructing a collection \mathcal{F} of number fields using \tilde{f} . Finally, in Section 3.9, we combine the results developed in the previous sections to prove Theorem 3.2, and use it to derive the main results in [Hua91a; Hua91b; Rón88; Rón92; Evd94; IKS09].

3.1 Preliminaries

We first review basic notations and facts in algebra. They are standard and can be found in various textbooks, e.g., [Lan02; AM69; Mar77]. Then we discuss *splitting of prime ideals* in number field extensions. Finally, for the certain rings $\bar{\mathcal{O}}_K$, we establish a one-to-one correspondence between their idempotent decompositions and the partitions of certain right coset spaces.

All rings are assumed to be commutative rings with unity.

Ideals. Recall that a subset I of a ring R is an *ideal* of R if (1) I is a subgroup of the underlying additive abelian group of R , and (2) $R \cdot I = \{ra : r \in R, a \in I\} \subseteq I$. For $x \in R$, denote by (x) , xR or Rx the ideal $\{rx : r \in R\}$ of R generated by x .

An ideal of R is *proper* if it is a proper subset of R . Let I be a proper ideal of R . We say I is *prime* if $I \neq R$ and $ab \in I$ implies $a \in I$ or $b \in I$ for any $a, b \in R$. And I is *maximal* if $I \neq R$ and there exists no ideal I' of R satisfying $I \subsetneq I' \subsetneq R$. A proper ideal I is prime (resp. maximal) iff the quotient ring R/I is an integral domain (resp. a field). In particular, maximal ideals are prime. For an ideal I_0 of R , the map $I \mapsto I/I_0$ is a one-to-one correspondence between the ideals of R containing I_0 and the ideals of R/I_0 , and it preserves primality and maximality.

If $\mathfrak{m}_1, \dots, \mathfrak{m}_k$ and \mathfrak{m} are maximal ideals of R and $\bigcap_{i=1}^k \mathfrak{m}_i \subseteq \mathfrak{m}$, then $\mathfrak{m} = \mathfrak{m}_i$ for some $i \in [k]$.² In particular, if $\bigcap_{i=1}^k \mathfrak{m}_i = 0$, then $\mathfrak{m}_1, \dots, \mathfrak{m}_k$ are the only maximal ideals of R .

Two ideals I, I' of R are *coprime* if $I + I' = R$. In particular, distinct maximal ideals are always coprime. For pairwise coprime ideals I_1, \dots, I_k , it holds that $\bigcap_{i=1}^k I_i = \prod_{i=1}^k I_i$. We also have

Lemma 3.1 (Chinese remainder theorem). *Suppose I_1, \dots, I_k are pairwise coprime ideals of R . Then the ring homomorphism*

$$\phi : R / \bigcap_{i=1}^k I_i \rightarrow \prod_{i=1}^k R / I_i$$

sending $x + \bigcap_{i=1}^k I_i$ to $(x + I_1, \dots, x + I_k)$ is an isomorphism.

Semisimple rings. A (commutative) ring is *semisimple* if it is isomorphic to a finite product of fields. The following lemma provides a characterization of semisimple rings.

Lemma 3.2. *A ring R is semisimple iff it has finitely many maximal ideals $\mathfrak{m}_1, \dots, \mathfrak{m}_k$ and $\bigcap_{i=1}^k \mathfrak{m}_i = 0$, in which case R is isomorphic to $\prod_{i=1}^k R/\mathfrak{m}_i$ via the map $x \mapsto (x + \mathfrak{m}_1, \dots, x + \mathfrak{m}_k)$.*

Proof. Suppose $R \cong \prod_{i=1}^k F_i$ is semisimple where each F_i is a field. For $i \in [k]$, let $\pi_i : R \rightarrow F_i$ be the i th projection and \mathfrak{m}_i be its kernel. Then $R/\mathfrak{m}_i \cong F_i$ and hence each \mathfrak{m}_i is a maximal ideal of R . Moreover we have $\bigcap_{i=1}^k \mathfrak{m}_i = 0$ and hence $\mathfrak{m}_1, \dots, \mathfrak{m}_k$ are the only maximal ideals. Conversely, suppose R has finitely many maximal ideals $\mathfrak{m}_1, \dots, \mathfrak{m}_k$ and $\bigcap_{i=1}^k \mathfrak{m}_i = 0$. Then by the Chinese remainder theorem, the map $R \rightarrow \prod_{i=1}^k R/\mathfrak{m}_i$ sending $x \in R$ to $(x + \mathfrak{m}_1, \dots, x + \mathfrak{m}_k)$ is a ring isomorphism. Each direct factor R/\mathfrak{m}_i is a field, and hence R is semisimple. \square

The semisimple rings considered in this chapter are all *semisimple \mathbb{F}_p -algebras*, i.e. semisimple rings that are also \mathbb{F}_p -algebras.

Idempotent elements. An element x of a ring is an *idempotent element* (or just an *idempotent*) if $x^2 = x$. Two idempotents x, y are *orthogonal* if $xy = 0$. A nonzero idempotent x is *primitive* if it cannot be written as a sum of two nonzero orthogonal

²See [AM69, Proposition 1.11] for a more general statement for prime ideals.

idempotents. As already stated in Definition 3.1, an *idempotent decomposition* of a ring R is a set I of nonzero mutually orthogonal idempotents of R satisfying $\sum_{x \in I} x = 1$. We say such an idempotent decomposition is *proper* if $|I| > 1$ and *complete* if all idempotents in I are primitive.

Lemma 3.3. *Let R be a semisimple ring. For every maximal ideal \mathfrak{m} of R , there exists a unique primitive idempotent $\delta_{\mathfrak{m}} \in R$ satisfying $\delta_{\mathfrak{m}} \equiv 1 \pmod{\mathfrak{m}}$ and $\delta_{\mathfrak{m}} \equiv 0 \pmod{\mathfrak{m}'}$ for all maximal ideal $\mathfrak{m}' \neq \mathfrak{m}$. Two elements $\delta_{\mathfrak{m}}$ and $\delta_{\mathfrak{m}'}$ are orthogonal iff $\mathfrak{m} \neq \mathfrak{m}'$. Furthermore*

- *the map $\mathfrak{m} \mapsto \delta_{\mathfrak{m}}$ is a one-to-one correspondence between the maximal ideals of R and the primitive idempotents of R , and*
- *the map $B \mapsto \sum_{\mathfrak{m} \in B} \delta_{\mathfrak{m}}$ is a one-to-one correspondence between the sets of maximal ideals of R and the idempotents of R .*

Proof. This is clear from the isomorphism $R \cong \prod_{\mathfrak{m} \in S} R/\mathfrak{m}$, where S denotes the set of all the maximal ideals of R . □

We also need the following lemma.

Lemma 3.4. *Suppose $\phi : R' \rightarrow R$ is a ring homomorphism between two semisimple rings R, R' . Let δ, δ' be idempotents of R and R' respectively satisfying $\phi(\delta')\delta = \delta$. Then ϕ induces a ring homomorphism from $R'/(1 - \delta')$ to $R/(1 - \delta)$ sending $x + (1 - \delta')$ to $\phi(x) + (1 - \delta)$ for $x \in R'$.*

Proof. It suffices to show that $\phi(1 - \delta')$ is in the ideal $(1 - \delta)$ of R , which holds since $(1 - \phi(\delta'))(1 - \delta) = 1 - \phi(\delta') - \delta + \phi(\delta')\delta = 1 - \phi(\delta') = \phi(1 - \delta')$. □

Finitely generated modules and free modules. A subset S of an R -module M generates M if $\sum_{x \in S} Rx = M$. And M is *finitely generated* if it is generated by a finite subset S . A *basis* of M over R , or an *R -basis* of M is a subset $S \subseteq M$ generating M for which the sum $M = \sum_{x \in S} Rx$ is a direct sum. We say M is *free* (over R) if it admits an R -basis. The *rank* of a finitely generated free module over R is the cardinality of any R -basis of it, which is finite and depends only on R .

Number fields. Elements in the algebraic closure $\bar{\mathbb{Q}}$ of \mathbb{Q} are called *algebraic numbers*. An algebraic number is *integral* or an *algebraic integer* if it is a root of a monic polynomial in $\mathbb{Z}[X]$. The set of algebraic integers is a subring of $\bar{\mathbb{Q}}$, denoted by \mathbb{A} . A *number field* is a finite degree field extension of \mathbb{Q} in $\bar{\mathbb{Q}}$. For a number field K , the subring $\mathcal{O}_K := \mathbb{A} \cap K$ is called the *ring of integers* of K . It is embedded in the \mathbb{Q} -vector space K as a lattice of rank $[K : \mathbb{Q}]$.

Suppose K/K_0 is a number field extension. We say $\alpha \in K$ is a *primitive element* of K over K_0 if $K = K_0(\alpha)$. Primitive elements always exist for any number field extension by the *primitive element theorem*.

Galois theory. Let K/K_0 be a field extension. The set of automorphisms of K fixing K_0 is a group, called the *automorphism group* of K over K_0 , and is denoted by $\text{Aut}(K/K_0)$. We say K is *Galois* over K_0 if $|\text{Aut}(K/K_0)| = [K : K_0]$, in which case $\text{Aut}(K/K_0)$ is also called the *Galois group* of K over K_0 and denoted by $\text{Gal}(K/K_0)$.

Theorem 3.3 (fundamental theorem of Galois theory). *Let K/K_0 be a Galois extension. Then for any intermediate field $K_0 \subseteq E \subseteq K$, the extension K/E is also a Galois extension. Furthermore, the map $E \mapsto \text{Gal}(K/E)$ is an inclusion-reversing one-to-one correspondence between the poset of intermediate fields $K_0 \subseteq E \subseteq K$ and the poset of subgroups of $\text{Gal}(K/K_0)$, with the inverse map $H \mapsto K^H$.*

Given a Galois extension K/K_0 , two subfields E, E' between K and K_0 are *conjugate* over K_0 if there exists an isomorphism $\tau_0 : E \rightarrow E'$ fixing K_0 . Such an isomorphism always extends to an automorphism $\tau \in \text{Gal}(K/K_0)$ of K . The corresponding Galois groups $\text{Gal}(K/E)$ and $\text{Gal}(K/E')$ satisfy $\text{Gal}(K/E') = \tau \text{Gal}(K/E) \tau^{-1}$. So conjugate subfields of K over K_0 correspond to conjugate subgroups in $\text{Gal}(K/K_0)$.

Now we restrict to number field extensions. Let K/K_0 be a number field extension. There exists a unique minimal number field that contains K and is Galois over K_0 , called the *Galois closure* of K/K_0 . For a polynomial $f(X) \in K_0[X]$ with roots $\alpha_1, \dots, \alpha_k \in \bar{\mathbb{Q}}$, the number field $K' = K_0(\alpha_1, \dots, \alpha_k)$ is called the *splitting field* of f over K_0 and is Galois over K_0 . We also write $\text{Gal}(f/K_0)$ for the corresponding Galois group $\text{Gal}(K'/K_0)$, called the Galois group of f over K_0 . If f is the minimal polynomial of a primitive element of K over K_0 , the splitting field of f over K_0 is exactly the Galois closure of K/K_0 .

Suppose K/K_0 is a Galois extension with the Galois group G . If $x \in K$ is an algebraic integer, so is ${}^g x$ for any $g \in G$ since $\mathbb{Z} \subseteq K_0$ is fixed by G . So the action of G on K restricts to an action on \mathcal{O}_K .

Splitting of prime ideals. The ring of integers of a number field is an example of a *Dedekind domain* [AM69; Mar77]. An ideal of a Dedekind domain is a nonzero prime ideal iff it is a maximal ideal, and hence these two notions are interchangeable. By convention, we use the notion of (nonzero) prime ideals instead of maximal ideals.

Let K be a number field. It follows from the theory of Dedekind domains [Mar77] that the ideal $p\mathcal{O}_K$ of \mathcal{O}_K splits uniquely (up to the ordering) into a product of prime ideals of \mathcal{O}_K :

$$p\mathcal{O}_K = \prod_{i=1}^k \mathfrak{P}_i.$$

For $i \in [k]$, the quotient ring $\mathcal{O}_K/\mathfrak{P}_i$ is a finite field extension of degree $d_i \in \mathbb{N}^+$ over \mathbb{F}_p , and $\sum_{i=1}^k d_i = [K : \mathbb{Q}]$. We say $\mathfrak{P}_1, \dots, \mathfrak{P}_k$ are the prime ideals of \mathcal{O}_K lying over p . If $\mathfrak{P}_1, \dots, \mathfrak{P}_k$ are distinct and $\mathcal{O}_K/\mathfrak{P}_i \cong \mathbb{F}_p$ for all $i \in [k]$ (and hence $k = [K : \mathbb{Q}]$), we say p splits completely in K . It is known that if p splits completely in K , then it also splits completely in any subfield of the Galois closure of K/\mathbb{Q} . See, e.g., [Mar77, Chapter 4]. We also need the following result that identifies the set of prime ideals lying over p with a right coset space in the case that p splits completely in a Galois extension containing K .

Theorem 3.4. *Let L be a Galois extension of \mathbb{Q} such that p splits completely in L , and let $G = \text{Gal}(L/\mathbb{Q})$. Fix a prime ideal \mathfrak{Q}_0 of \mathcal{O}_L lying over p . For any subgroup $H \subseteq G$ and its fixed field $K = L^H$, the map $Hg \mapsto {}^g \mathfrak{Q}_0 \cap \mathcal{O}_K$ is a one-to-one correspondence between the right cosets in $H \backslash G$ and the prime ideals of \mathcal{O}_K lying over p .³*

See, e.g., [Mar77, Theorem 33]. As the prime ideals of \mathcal{O}_K lying over p are exactly those containing $p\mathcal{O}_K$, we get the following correspondence by passing to the quotient ring $\bar{\mathcal{O}}_K := \mathcal{O}_K/p\mathcal{O}_K$.

Corollary 3.1. *Let L, G, \mathfrak{Q}_0 be as in Theorem 3.4. For any subgroup $H \subseteq G$ and its fixed field $K = L^H$, the map $Hg \mapsto ({}^g \mathfrak{Q}_0 \cap \mathcal{O}_K)/p\mathcal{O}_K$ is a one-to-one*

³Note that this map is well defined: for another representative $hg \in G$ of Hg where $h \in H$, we have ${}^{hg} \mathfrak{Q}_0 \cap \mathcal{O}_K = {}^h ({}^g \mathfrak{Q}_0 \cap \mathcal{O}_K) = {}^g \mathfrak{Q}_0 \cap \mathcal{O}_K$ since \mathcal{O}_K is fixed by H .

correspondence between the right cosets in $H \backslash G$ and the prime (and maximal) ideals of $\bar{\mathcal{O}}_K$.

Idempotent decompositions vs. partitions of a right coset space. Suppose p splits completely into a product of prime ideals $\mathfrak{P}_1, \dots, \mathfrak{P}_k$ in a number field K . Then $\mathfrak{P}_1/p\mathcal{O}_K, \dots, \mathfrak{P}_k/p\mathcal{O}_K$ are the prime (and maximal) ideals of $\bar{\mathcal{O}}_K$. As the intersection of these ideals equals $p\mathcal{O}_K/p\mathcal{O}_K = 0$, the ring $\bar{\mathcal{O}}_K$ is semisimple by Lemma 3.2. The prime ideals $\mathfrak{P}_i/p\mathcal{O}_K$ correspond to the primitive idempotents of $\bar{\mathcal{O}}_K$ by Lemma 3.3 and also to the cosets in a right coset space by Corollary 3.1. We combine them and establish a correspondence between the idempotent decompositions of $\bar{\mathcal{O}}_K$ and the partitions of a certain right coset space.

For a number field extension L/K , the inclusion $\mathcal{O}_K \hookrightarrow \mathcal{O}_L$ induces a map

$$i_{K,L} : \bar{\mathcal{O}}_K \rightarrow \bar{\mathcal{O}}_L$$

with the kernel $(p\mathcal{O}_L \cap \mathcal{O}_K)/p\mathcal{O}_K$. As $p\mathcal{O}_L \cap \mathcal{O}_K = p\mathcal{O}_K$,⁴ this map is injective, which identifies $\bar{\mathcal{O}}_K$ with a subring of $\bar{\mathcal{O}}_L$. Also note that if L/\mathbb{Q} is a Galois extension with the Galois group G , the action of G on \mathcal{O}_L induces an action on $\bar{\mathcal{O}}_L$ and permutes the maximal ideals of $\bar{\mathcal{O}}_L$. These observations are used in Definition 3.2 below.

Fix the following notations: let L be a Galois extension of \mathbb{Q} with $\text{Gal}(L/\mathbb{Q}) = G$ and suppose p splits completely in L . For a nonzero prime ideal \mathfrak{Q} of \mathcal{O}_L lying over p , define $\bar{\mathfrak{Q}} := \mathfrak{Q}/p\mathcal{O}_L$ which is a prime (and hence maximal) ideal of $\bar{\mathcal{O}}_L$, and let $\delta_{\bar{\mathfrak{Q}}}$ be the primitive idempotent of $\bar{\mathcal{O}}_L$ satisfying $\delta_{\bar{\mathfrak{Q}}} \equiv 1 \pmod{\bar{\mathfrak{Q}}}$ and $\delta_{\bar{\mathfrak{Q}}} \equiv 0 \pmod{\bar{\mathfrak{Q}}'}$ for all maximal ideal $\bar{\mathfrak{Q}}' \neq \bar{\mathfrak{Q}}$ of $\bar{\mathcal{O}}_L$ (cf. Lemma 3.3).

Definition 3.2. *Suppose H is a subgroup of G and $K = L^H$. Fix a prime ideal \mathfrak{Q}_0 of \mathcal{O}_L lying over p . Then*

- *for an idempotent decomposition I of $\bar{\mathcal{O}}_K$, define $P(I)$ to be the partition of $H \backslash G$ such that $Hg, Hg' \in H \backslash G$ are in the same block iff $g^{-1}(i_{K,L}(\delta)) \equiv g'^{-1}(i_{K,L}(\delta)) \pmod{\bar{\mathfrak{Q}}_0}$ holds for all $\delta \in I$, and*
- *for a partition P of $H \backslash G$, define $I(P)$ to be the idempotent decomposition of $\bar{\mathcal{O}}_K$ consisting of the idempotents $\delta_B := i_{K,L}^{-1} \left(\sum_{g \in G: Hg \in B} g \delta_{\bar{\mathfrak{Q}}_0} \right)$, where B ranges over the blocks in P .⁵*

⁴To see this, note that if $x \in p\mathcal{O}_L \cap \mathcal{O}_K$, then $x/p \in \mathcal{O}_L \cap K = \mathcal{O}_K$.

⁵We show in the proof of Lemma 3.5 that $\sum_{g \in G: Hg \in B} g \delta_{\bar{\mathfrak{Q}}_0}$ does lie in the image of $i_{K,L}$, and hence δ_B is well defined.

We have the following lemmas.

Lemma 3.5. *The partitions $P(I)$ and the idempotent decompositions $I(P)$ are well defined. And for any idempotent decomposition I of $\bar{\mathcal{O}}_K$, the idempotents $\delta \in I$ correspond one-to-one to the blocks of $P(I)$ via the map $\delta \mapsto B_\delta := \{Hg \in H \setminus G : g^{-1}(i_{K,L}(\delta)) \equiv 1 \pmod{\bar{\mathcal{Q}}_0}\}$ with the inverse map $B \mapsto \delta_B$.*

Lemma 3.6. *The map $I \mapsto P(I)$ is a one-to-one correspondence between the idempotent decompositions of $\bar{\mathcal{O}}_K$ and the partitions of $H \setminus G$, with the inverse map $P \mapsto I(P)$.*

The proofs of Lemma 3.5 and Lemma 3.6 are routine and can be found in Appendix B. In particular, Lemma 3.6 establishes a one-to-one correspondence between the idempotent decompositions of $\bar{\mathcal{O}}_K$ and the partitions of $H \setminus G$.

3.2 Algorithmic preliminaries

In this section, we present some basic procedures used in the factoring algorithm, mostly related to number fields. Standard references include [Len92; Coh93].

Let A be an R -algebra that is a free R -module of finite rank. In the factoring algorithm, we represent such an algebra by maintaining an R -basis $B = \{b_1, \dots, b_d\}$ of it. The *structure constants* of A in the basis B are the constants $c_{ijk} \in R$ defined by $b_i b_j = \sum_{k=1}^d c_{ijk} b_k$. Given these structure constants, arithmetic operations of A can be performed in polynomial time, provided that the those of R can also be performed in polynomial time. In the discussion below, we use the phrase “computing A ” for the task of computing the structure constants of A in the R -basis B associated with A . And by “computing a ” for $a \in A$ we mean computing the constants $c_i \in R$ satisfying $a = \sum_{i=1}^d c_i b_i$. The interesting cases of R to us are \mathbb{Z} , \mathbb{Q} , and \mathbb{F}_p .

Now let R' be an R -algebra and let A' be an R' -algebra that is a free R' -module of finite rank. Let $\phi : A \rightarrow A'$ be an R -linear map. We use the phrase “computing ϕ ” for the task of computing $\phi(b_i) \in A'$ for all $b_i \in B$, in terms of the coefficients of $\phi(b_i)$ in the R' -basis B' associated with A' . The interesting cases to us are (1) $R = R' \in \{\mathbb{Z}, \mathbb{Q}, \mathbb{F}_p\}$, (2) $R = \mathbb{Z}$, $R' = \mathbb{Q}$ and ϕ is an inclusion that embeds a lattice in a vector space over \mathbb{Q} , and (3) $R = \mathbb{Z}$, $R' = \mathbb{Z}/p\mathbb{Z} \cong \mathbb{F}_p$, and ϕ is a quotient map from a lattice to a vector space over \mathbb{F}_p .

The *size* of an object used in the algorithm is the number of bits used to encode this object.

Encoding a number field. Let K be a number field of degree $d \in \mathbb{N}^+$ over \mathbb{Q} . We encode K using a primitive element $\alpha \in K$ over \mathbb{Q} , or more precisely, the minimal polynomial $g(X) \in \mathbb{Q}[X]$ of α over \mathbb{Q} . Given $g(X)$, we compute $\mathbb{Q}[X]/(g(X))$ in the standard \mathbb{Q} -basis $\{1 + (g(X)), X + (g(X)), \dots, X^{d-1} + (g(X))\}$ and use it to represent K . This is justified by the isomorphism $\mathbb{Q}[X]/(g(X)) \cong K$ sending $X + (g(X))$ to α .

Computing $\bar{\mathcal{O}}_K$. Given K and a prime number p , we want to compute the \mathbb{F}_p -algebra $\bar{\mathcal{O}}_K = \mathcal{O}_K/p\mathcal{O}_K$. It is natural to first compute the ring of integers \mathcal{O}_K and then pass to the quotient ring $\bar{\mathcal{O}}_K$. Unfortunately, computing a \mathbb{Z} -basis of \mathcal{O}_K in K is in general as hard as finding the largest square factor of a given integer [Chi89; Len92]. We overcome the difficulty by working with a subring $\mathcal{O}'_K \subseteq \mathcal{O}_K$ instead of \mathcal{O}_K such that $[\mathcal{O}_K : \mathcal{O}'_K]$ is finite and coprime to p . Such a subring is called a *p-maximal order* of K , which can be efficiently computed:

Theorem 3.5. *There exists a polynomial-time algorithm that given K and p , computes a p -maximal order \mathcal{O}'_K of K together with the inclusion $\mathcal{O}'_K \hookrightarrow K$.*

See, e.g., [Coh93, Chapter 6]. We may use \mathcal{O}'_K in place of \mathcal{O}_K thanks to the following lemma.

Lemma 3.7. *For a p -maximal order \mathcal{O}'_K of K , the ring homomorphism $\mathcal{O}'_K/p\mathcal{O}'_K \rightarrow \mathcal{O}_K/p\mathcal{O}_K = \bar{\mathcal{O}}_K$ induced from the inclusion $\mathcal{O}'_K \hookrightarrow \mathcal{O}_K$ is an isomorphism.*

Proof. To show surjectivity, it suffices to show that \mathcal{O}'_K and $p\mathcal{O}_K$ span \mathcal{O}_K over \mathbb{Z} . Note that $n_1 := [\mathcal{O}_K : \mathcal{O}'_K]$ is coprime to p and $n_2 := [\mathcal{O}_K : p\mathcal{O}_K]$ is a power of p . The index of the lattice spanned by \mathcal{O}'_K and $p\mathcal{O}_K$ in \mathcal{O}_K divides both n_1 and n_2 and hence equals one, as desired.

On the other hand, note that \mathcal{O}_K and \mathcal{O}'_K are both lattices of rank $[K : \mathbb{Q}]$. So $\bar{\mathcal{O}}_K$ and $\mathcal{O}'_K/p\mathcal{O}'_K$ are both vector spaces of dimension $[K : \mathbb{Q}]$ over \mathbb{F}_p . Therefore the map $\mathcal{O}'_K/p\mathcal{O}'_K \rightarrow \bar{\mathcal{O}}_K$ is an isomorphism. \square

This provides a method of computing the \mathbb{F}_p -algebra $\bar{\mathcal{O}}_K$:

Lemma 3.8. *There exists a polynomial-time algorithm `ComputeQuotientRing` that given K and p , computes the quotient ring $\bar{\mathcal{O}}_K$, a p -maximal order \mathcal{O}'_K , the inclusion $\mathcal{O}'_K \hookrightarrow K$, and the quotient map $\pi : \mathcal{O}'_K \rightarrow \bar{\mathcal{O}}_K$ sending $x \in \mathcal{O}'_K$ to $x + p\mathcal{O}_K$.*

Proof. Compute \mathcal{O}'_K and the inclusion $\mathcal{O}'_K \hookrightarrow K$ using Theorem 3.5. In particular the structure constants $c_{ijk} \in \mathbb{Z}$ of \mathcal{O}'_K in some \mathbb{Z} -basis $\{b_1, \dots, b_d\}$ are computed, where $d = [K : \mathbb{Q}]$. The structure constants of $\mathcal{O}'_K/p\mathcal{O}'_K$ in the \mathbb{F}_p -basis $\{b_1 + p\mathcal{O}'_K, \dots, b_k + p\mathcal{O}'_K\}$ are simply $c_{ijk} \bmod p$. By Lemma 3.7, they are also the structure constants of $\bar{\mathcal{O}}_K$ in the \mathbb{F}_p -basis $\{b_1 + p\mathcal{O}_K, \dots, b_k + p\mathcal{O}_K\}$. The map π is specified by the data $\phi(b_i) = b_i + p\mathcal{O}_K$. \square

Note that in addition to $\bar{\mathcal{O}}_K$, we also compute the auxiliary data of \mathcal{O}'_K and the maps from \mathcal{O}'_K to $\bar{\mathcal{O}}_K$ and K . They are used for the algorithms in Lemma 3.9 and Lemma 3.11 below.

Computing the residue of an algebraic integer modulo p . We need an algorithm computing the image of an algebraic integer $\alpha \in \mathcal{O}_K$ in $\bar{\mathcal{O}}_K$, where α is given as an element of K .

Lemma 3.9. *There exists a polynomial-time algorithm `ComputeResidue` that takes the following data as the input*

- a number fields K , a prime number p , and $\alpha \in \mathcal{O}_K$ given as an element of K ,
- the outputs of `ComputeQuotientRing` (see Lemma 3.8) on the inputs (K, p) , i.e., the quotient ring $\bar{\mathcal{O}}_K$, a maximal p -orders \mathcal{O}'_K , the inclusion $\mathcal{O}'_K \hookrightarrow K$, and the quotient map $\mathcal{O}'_K \rightarrow \bar{\mathcal{O}}_K$,

and computes $\alpha + p\mathcal{O}_K \in \bar{\mathcal{O}}_K$.

The proof of Lemma 3.9 can be found in Appendix B.

Computing embeddings of number fields. Embeddings of a number field in another can be computed efficiently, thanks to the polynomial-time factoring algorithms for number fields [Len83; Lan85].

Theorem 3.6 ([Len83; Lan85]). *There exists a polynomial-time algorithm that given a number field K and a polynomial $g(X) \in K[X]$, factorizes $g(X)$ into irreducible factors over K .*

Let K, K' be number fields and suppose K is encoded with a primitive element $\alpha \in K$ whose minimal polynomial is $g(X) \in \mathbb{Q}[X]$. Each embedding ϕ of K in

K' is determined by the image $\phi(\alpha) \in K'$ which is a root of $g(X)$. These roots can be enumerated by factoring $g(X)$ over K' using Theorem 3.6. So we have:

Lemma 3.10. *There exists a polynomial-time algorithm `ComputeEmbeddings` that given number fields K and K' , computes all the embeddings of K in K' .*

Computing induced ring homomorphisms between $\bar{\mathcal{O}}_K$. Let $\phi : K \hookrightarrow K'$ be an embedding of number fields, which restricts to an inclusion $\mathcal{O}_K \hookrightarrow \mathcal{O}_{K'}$. By passing to the quotient rings $\bar{\mathcal{O}}_K$ and $\bar{\mathcal{O}}_{K'}$, we obtain a ring homomorphism $\bar{\phi} : \bar{\mathcal{O}}_K \rightarrow \bar{\mathcal{O}}_{K'}$. And we say the map $\bar{\phi}$ is *induced from* ϕ . The following lemma states that $\bar{\phi}$ can be efficiently computed from ϕ and some auxiliary data.

Lemma 3.11. *There exists a polynomial-time algorithm `ComputeRingHom` that takes the following data as the input*

- number fields K, K' , an embedding $\phi : K \rightarrow K'$, and a prime number p ,
- the outputs of `ComputeQuotientRing` (see Lemma 3.8) on the inputs (K, p) and (K', p) respectively,⁶

and computes the ring homomorphism $\bar{\phi} : \bar{\mathcal{O}}_K \rightarrow \bar{\mathcal{O}}_{K'}$ induced from ϕ .

The proof of Lemma 3.11 can be found in Appendix B.

3.3 Reduction to computing an idempotent decomposition of $\bar{\mathcal{O}}_F$

Now we start describing the \mathcal{P} -scheme algorithm. Fix the following notations in the remaining sections:

- $f(X)$: the input polynomial in $\mathbb{F}_p[X]$ to be factorized, which is square-free and completely reducible over \mathbb{F}_p ,
- $\tilde{f}(X)$: an irreducible lifted polynomial of $f(X)$ in $\mathbb{Z}[X]$,
- F : the number field $\mathbb{Q}[X]/(\tilde{f}(X))$,
- L : the splitting field of \tilde{f} over \mathbb{Q} ,
- G : the Galois group $\text{Gal}(L/\mathbb{Q}) = \text{Gal}(\tilde{f}/\mathbb{Q})$,

⁶That is, the quotient rings $\bar{\mathcal{O}}_K, \bar{\mathcal{O}}_{K'}$, the maximal p -orders $\mathcal{O}'_K, \mathcal{O}'_{K'}$, the inclusions $\mathcal{O}'_K \hookrightarrow K, \mathcal{O}'_{K'} \hookrightarrow K'$, and the quotient maps $\mathcal{O}'_K \rightarrow \bar{\mathcal{O}}_K, \mathcal{O}'_{K'} \rightarrow \bar{\mathcal{O}}_{K'}$.

- \mathfrak{Q}_0 : a fixed prime ideal of \mathcal{O}_L lying over p .

In this section, we reduce the problem of factoring f to that of computing an idempotent decomposition of $\bar{\mathcal{O}}_F$. For simplicity, we first assume that \tilde{f} is a *monic* polynomial, and then remove the assumption at the end of this section.

Ring isomorphism between $\mathbb{F}_p[X]/(f(X))$ and $\bar{\mathcal{O}}_K$. Let $\alpha := X + (\tilde{f}(X)) \in F$ which is a root of \tilde{f} . As $\tilde{f}(X) \in \mathbb{Z}[X]$ is monic, we know $\alpha \in \mathcal{O}_F$. Define the ring homomorphism $\tilde{\tau} : \mathbb{F}_p[X] \rightarrow \bar{\mathcal{O}}_F$ by letting $\tilde{\tau}(X) = \alpha + p\mathcal{O}_F$, which is well defined since $\bar{\mathcal{O}}_F$ is an \mathbb{F}_p -algebra. Moreover, we have $\tilde{\tau}(f(X)) = \tilde{f}(\alpha) + p\mathcal{O}_F = 0$. So $\tilde{\tau}$ induces a ring homomorphism $\tau : \mathbb{F}_p[X]/(f(X)) \rightarrow \bar{\mathcal{O}}_K$ sending $X + (f(X))$ to $\alpha + p\mathcal{O}_F$.

Let f_1, \dots, f_n be the monic irreducible factors of f over \mathbb{F}_p . As f_i are irreducible and distinct, the ring $\mathbb{F}_p[X]/(f(X))$ is semisimple with the maximal ideals $(f_i(X))$, $i = 1, \dots, n$. Then $\bar{\mathcal{O}}_F$ is also semisimple. Indeed, we have the following lemma:

Lemma 3.12. *The map $\tau : \mathbb{F}_p[X]/(f(X)) \rightarrow \bar{\mathcal{O}}_F$ is a ring isomorphism, and p splits completely in F .*

Proof. The second claim follows from the first since $\mathbb{F}_p[X]/(f(X))$ has n distinct maximal ideals. To prove the first claim, note that the ring homomorphism $\mathbb{F}_p[X]/(f(X)) \rightarrow \mathbb{Z}[\alpha]/p\mathbb{Z}[\alpha]$ sending X to $\alpha + p\mathbb{Z}[\alpha]$ is an isomorphism. So it suffices to show that the natural inclusion $\mathbb{Z}[\alpha] \hookrightarrow \mathcal{O}_F$ induces an isomorphism $\mathbb{Z}[\alpha]/p\mathbb{Z}[\alpha] \rightarrow \bar{\mathcal{O}}_F$.

For $i \in [n]$, choose $\tilde{f}_i(X) \in \mathbb{Z}[X]$ that lifts the factor $f_i(X) \in \mathbb{F}_p[X]$ of f , and define the ideal \mathfrak{P}_i of $\mathbb{Z}[\alpha]$ to be the one generated by $\tilde{f}_i(\alpha)$ and p . As $\mathbb{Z}[\alpha]/p\mathbb{Z}[\alpha] \cong \mathbb{F}_p[X]/(f(X))$ is semisimple, we have $\bigcap_{i=1}^n \mathfrak{P}_i = p\mathbb{Z}[\alpha]$. By [AM69, Theorem 5.10], for each $i \in [n]$, we may choose a prime ideal \mathfrak{Q}_i of \mathcal{O}_F lying over p such that $\mathfrak{Q}_i \cap \mathbb{Z}[\alpha] = \mathfrak{P}_i$. Then we have

$$p\mathcal{O}_F \cap \mathbb{Z}[\alpha] \subseteq \left(\bigcap_{i=1}^n \mathfrak{Q}_i \right) \cap \mathbb{Z}[\alpha] = \bigcap_{i=1}^n \mathfrak{P}_i = p\mathbb{Z}[\alpha].$$

So the map $\mathbb{Z}[\alpha]/p\mathbb{Z}[\alpha] \rightarrow \bar{\mathcal{O}}_F$ is injective. It is in fact an isomorphism since $\mathbb{Z}[\alpha]/p\mathbb{Z}[\alpha]$ and $\bar{\mathcal{O}}_F$ are both vector spaces of dimension n over \mathbb{F}_p . \square

Extracting a factorization from an idempotent decomposition. Let I_F be an idempotent decomposition of $\bar{\mathcal{O}}_F$. By Lemma 3.12, the set $\tau^{-1}(I_F) = \{\tau^{-1}(\delta) : \delta \in I_F\}$ is an idempotent decomposition of $\mathbb{F}_p[X]/(f(X))$. Given $\delta \in I_F$, we can extract a factor $g_\delta(X)$ of $f(X)$ by

$$g_\delta(X) := \gcd(f(X), h_\delta(X)),$$

where $h_\delta(X) \in \mathbb{F}_p[X]$ is a nonzero polynomial of degree at most n lifting $1 - \tau^{-1}(\delta) \in \mathbb{F}_p[X]/(f(X))$. The factor $g_\delta(X)$ is the product of the monic irreducible factors $f_i(X)$ satisfying $\tau^{-1}(\delta) \equiv 1 \pmod{f_i(X)}$. As $f(X) = \tilde{f}(X) \pmod{p}$ is monic and the elements $\tau^{-1}(\delta)$ form an idempotent decomposition of the ring $\mathbb{F}_p[X]/(f(X))$, we have the equality

$$f(X) = \prod_{\delta \in I} g_\delta(X).$$

This gives the following algorithm that computes a factorization of f from I_F :

Algorithm 1 ExtractFactors

Input: $p, f, \tilde{f}, F, \bar{\mathcal{O}}_F$, idempotent decomposition I_F of $\bar{\mathcal{O}}_F$,

p -maximal order \mathcal{O}'_F of F and maps $\mathcal{O}'_F \hookrightarrow F, \mathcal{O}'_F \rightarrow \bar{\mathcal{O}}_F$

Output: factorization of f

- 1: $\alpha \leftarrow X + (\tilde{f}(X)) \in F$
 - 2: call ComputeResidue to compute $\alpha + p\mathcal{O}'_F \in \bar{\mathcal{O}}_F$
 - 3: compute the ring homomorphism $\tau : \mathbb{F}_p[X]/(f(X)) \rightarrow \bar{\mathcal{O}}_F$ sending $X + (f(X))$ to $\alpha + p\mathcal{O}'_F$
 - 4: **for** $\delta \in I_F$ **do**
 - 5: compute nonzero $h_\delta(X) \in \mathbb{F}_p[X]$ of degree at most n lifting $1 - \tau^{-1}(\delta)$
 - 6: $g_\delta(X) \leftarrow \gcd(f(X), h_\delta(X))$
 - 7: **return** the factorization $f(X) = \prod_{\delta \in I_F} g_\delta(X)$
-

For the purpose of computing the map τ , the input contains some auxiliary data (e.g., a p -maximal order \mathcal{O}'_F and the related maps) other than the idempotent decomposition I_F . For now we note that the auxiliary data can be prepared in polynomial time using the subroutines in Section 3.2. Then we have:

Theorem 3.7. *The algorithm ExtractFactors computes the factorization $f(X) = \prod_{\delta \in I_F} g_\delta(X)$ in polynomial time. In particular, it computes the complete factorization (resp. a proper factorization) of $f(X)$ in polynomial time iff the idempotent decomposition I_F of $\bar{\mathcal{O}}_F$ is complete (resp. proper).*

Proof. The algorithm clearly runs in polynomial time: Line 1 is implemented by factoring \tilde{f} over F using Theorem 3.6. The loop in Lines 4–6 iterates $|I_F| \leq n$ times. Line 5 is implemented by solving a system of linear equations over \mathbb{F}_p and Line 6 by the Euclidean algorithm. The fact that the factorization is complete (resp. proper) iff I_F is complete (resp. proper) follows from the fact that $\tau : \mathbb{F}_p[X]/(f(X)) \rightarrow \bar{\mathcal{O}}_F$ is a ring isomorphism. \square

Therefore the problem of computing the complete factorization (resp. a proper factorization) of f reduces to the problem of computing the complete (resp. a proper) idempotent decomposition of $\bar{\mathcal{O}}_F$.

The reduction for non-monic polynomials. After a slight change, the above reduction also works for a possibly non-monic polynomial \tilde{f} . We explain it now.

Suppose $c \in \mathbb{Z} - \{0\}$ is the leading coefficient of \tilde{f} . Its residue $\bar{c} := c \bmod p \in \mathbb{F}_p$ is nonzero since $\deg(\tilde{f}) = \deg(f) = n$. Define $\tilde{f}'(X) := c^{n-1} \cdot \tilde{f}(X/c) \in \mathbb{Z}[X]$ and $f'(X) := \tilde{f}'(X) \bmod p \in \mathbb{F}_p[X]$. The polynomials \tilde{f}' and f' are monic, and $f'(X) = \bar{c}^{n-1} \cdot f(X/\bar{c})$. Let α be a root of \tilde{f} in F as before. Then $\alpha' := c\alpha$ is a root of \tilde{f}' and hence is in \mathcal{O}_F .

Run the algorithm `ExtractFactors` above except that f , \tilde{f} and α are replaced with f' , \tilde{f}' and α' respectively. Then we obtain a factorization $f'(X) = \prod_{\delta \in I_F} g'_\delta(X)$ where the factors $g'_\delta(X) \in \mathbb{F}_p[X]$ are monic. Substituting X with $\bar{c}X$, we obtain a factorization

$$f(X) = \bar{c} \cdot \prod_{\delta \in I_F} g_\delta(X)$$

with the monic factors $g_\delta(X) := \bar{c}^{-\deg(g'_\delta)} \cdot g'_\delta(\bar{c}X) \in \mathbb{F}_p[X]$. Theorem 3.7 then holds for f and \tilde{f} .

3.4 Main algorithm

We present the main body of the \mathcal{P} -scheme algorithm in this section. Its input contains a collection of number fields that are isomorphic to subfields of L . In order to avoid duplicate data, we assume that these number fields are mutually non-isomorphic. This is formalized by the following definition:

Definition 3.3 ((\mathbb{Q}, g) -subfield system). *Let $g(X)$ be a polynomial in $\mathbb{Q}[X]$ with the splitting field $L(g)$ over \mathbb{Q} . Let \mathcal{F} be a collection of number fields such that (1) the fields in \mathcal{F} are mutually non-isomorphic, and (2) each field $K' \in \mathcal{F}$ is isomorphic to a subfield of $L(g)$. We say \mathcal{F} is a (\mathbb{Q}, g) -subfield system.*

Given a (\mathbb{Q}, g) -subfield system, we define a subgroup system over $\text{Gal}(g/\mathbb{Q})$ as follows.

Definition 3.4. *Let $g(X)$ be a polynomial in $\mathbb{Q}[X]$ with the splitting field $L(g)$ over \mathbb{Q} . Let \mathcal{F} be a (\mathbb{Q}, g) -subfield system. Define \mathcal{P}^\sharp to be the poset of subfields of $L(g)$ that includes all the fields isomorphic to those in \mathcal{F} :*

$$\mathcal{P}^\sharp := \{K' \subseteq L(g) : K' \cong K \text{ for some } K \in \mathcal{F}\}.$$

By Galois theory, it corresponds to a poset \mathcal{P} of subgroups of $\text{Gal}(g/\mathbb{Q})$, given by

$$\mathcal{P} := \{H \subseteq \text{Gal}(g/\mathbb{Q}) : (L(g))^H \in \mathcal{P}^\sharp\}$$

which is closed under conjugation in $\text{Gal}(g/\mathbb{Q})$, and hence is a subgroup system over $\text{Gal}(g/\mathbb{Q})$. We say \mathcal{P} and \mathcal{P}^\sharp are associated with \mathcal{F} .

The pseudocode of the algorithm is given in Algorithm 2 below. Its input is the prime number p and a (\mathbb{Q}, \tilde{f}) -subfield system \mathcal{F} . We fix \mathcal{P} to be the subgroup system over $G = \text{Gal}(\tilde{f}/\mathbb{Q})$ associated with \mathcal{F} .

The algorithm outputs, for every $K \in \mathcal{F}$, the ring $\bar{\mathcal{O}}_K$ and an idempotent decomposition I_K of $\bar{\mathcal{O}}_K$, together with the auxiliary data of a p -maximal order \mathcal{O}'_K and the related maps $\mathcal{O}'_K \hookrightarrow K$, $\mathcal{O}'_K \rightarrow \bar{\mathcal{O}}_K$. We will see below that the idempotent decompositions I_K altogether determine a \mathcal{P} -collection, which is guaranteed to be a strongly antisymmetric \mathcal{P} -scheme when the algorithm terminates.

The first half (Lines 1–7) of the algorithm is the preprocessing stage, where we compute $\bar{\mathcal{O}}_K$ for $K \in \mathcal{F}$ and the ring homomorphisms between them that are induced from the field embeddings. For each $K \in \mathcal{F}$, we also initialize the idempotent decomposition I_K of $\bar{\mathcal{O}}_K$ to be the trivial one containing only the unity of $\bar{\mathcal{O}}_K$.

The second half (Lines 8–12) is the “refining” stage. To understand it, we need to associate a \mathcal{P} -collection \mathcal{C} with the idempotent decompositions I_K . By Lemma 3.12, we know p splits completely in F . So it also splits completely in every subfield of L . In particular, for a field K in \mathcal{P}^\sharp or \mathcal{F} , the quotient ring $\bar{\mathcal{O}}_K$ is semisimple.

For each $H \in \mathcal{P}$, we define a partition C_H of the coset space $H \backslash G$ as follows: Let K be the unique field in \mathcal{F} isomorphic to L^H . Fix an isomorphism $\tau_H : K \rightarrow L^H$, which induces a ring isomorphism $\bar{\tau}_H : \bar{\mathcal{O}}_K \rightarrow \bar{\mathcal{O}}_{L^H}$. Define $I_H := \bar{\tau}_H(I_K)$, which is an idempotent decomposition of $\bar{\mathcal{O}}_{L^H}$. By Definition 3.2, it corresponds to a

Algorithm 2 ComputePscheme

Input: prime number p , (\mathbb{Q}, \tilde{f}) -subfield system \mathcal{F}

Output: for each $K \in \mathcal{F}$: $\bar{\mathcal{O}}_K$, idempotent decomposition I_K of $\bar{\mathcal{O}}_K$,
 p -maximal order \mathcal{O}'_K of K and maps $\mathcal{O}'_K \hookrightarrow K$, $\mathcal{O}'_K \rightarrow \bar{\mathcal{O}}_K$

- 1: **for** $K \in \mathcal{F}$ **do**
 - 2: call ComputeQuotientRing to compute $\bar{\mathcal{O}}_F$, a p -maximal order \mathcal{O}'_K of K
 and maps $\mathcal{O}'_K \hookrightarrow K$, $\mathcal{O}'_K \rightarrow \bar{\mathcal{O}}_K$
 - 3: $I_K \leftarrow \{1\}$, where 1 denotes the unity of $\bar{\mathcal{O}}_K$
 - 4: **for** $(K, K') \in \mathcal{F}^2$ **do**
 - 5: call ComputeEmbeddings to compute all the embeddings from K to K'
 - 6: **for** embedding $\phi : K \hookrightarrow K'$ **do**
 - 7: call ComputeRingHom to compute $\bar{\phi} : \bar{\mathcal{O}}_K \rightarrow \bar{\mathcal{O}}_{K'}$ induced from ϕ
 - 8: **repeat**
 - 9: call CompatibilityAndInvarianceTest
 - 10: call RegularityTest
 - 11: call StrongAntisymmetryTest
 - 12: **until** I_K remains the same in the last iteration for all $K \in \mathcal{F}$
 - 13: **return** $\bar{\mathcal{O}}_K$, I_K , \mathcal{O}'_K and the maps $\mathcal{O}'_K \hookrightarrow K$, $\mathcal{O}'_K \rightarrow \bar{\mathcal{O}}_K$ for $K \in \mathcal{F}$
-

partition $P(I_H)$ of $H \setminus G$.⁷ And we define

$$C_H := P(I_H).$$

Finally, define the \mathcal{P} -collection \mathcal{C} by

$$\mathcal{C} := \{C_H : H \in \mathcal{P}\}.$$

We call several subroutines to update I_K in Lines 9–11, whose effects can be understood in terms of \mathcal{C} :

Lemma 3.13. *There exists a subroutine CompatibilityAndInvarianceTest that updates I_K in time polynomial in $\log p$ and the size of \mathcal{F} so that the partitions $C_H \in \mathcal{C}$ are refined, and at least one partition C_H is properly refined if \mathcal{C} is not compatible or invariant.*

⁷Definition 3.2 is made with respect to a fixed prime ideal \mathfrak{Q}_0 of \mathcal{O}_L lying over p . This ideal is chosen at the beginning of Section 3.3.

Lemma 3.14. *There exists a subroutine `RegularityTest` that updates I_K in time polynomial in $\log p$ and the size of \mathcal{F} so that the partitions $C_H \in \mathcal{C}$ are refined, and at least one partition C_H is properly refined if \mathcal{C} is compatible but not regular.*

Lemma 3.15. *Under GRH, there exists a subroutine `StrongAntisymmetryTest` that updates I_K in time polynomial in $\log p$ and the size of \mathcal{F} so that the partitions $C_H \in \mathcal{C}$ are refined, and at least one partition C_H is properly refined if \mathcal{C} is a \mathcal{P} -scheme, but not strongly antisymmetric.*

We will describe these subroutines and prove the lemmas above in the next three sections. For now we just assume them and prove the main result of this section:

Theorem 3.8 (Theorem 3.1 restated). *Under GRH, the algorithm `ComputePscheme` runs in time polynomial in the size of the input, and when it terminates, the \mathcal{P} -collection \mathcal{C} is a strongly antisymmetric \mathcal{P} -scheme.*

Proof. We first analyze the running time. As each field $K \in \mathcal{F}$ is encoded by a rational polynomial of degree $[K : \mathbb{Q}]$, the total degree $N := \sum_{K \in \mathcal{F}} [K : \mathbb{Q}]$ is bounded by the size of \mathcal{F} . The loops in Lines 1–3 and Lines 4–7 iterate $|\mathcal{F}| \leq N$ and $|\mathcal{F}^2| \leq N^2$ times respectively. For each $(K, K') \in \mathcal{F}^2$, there are at most $[K : \mathbb{Q}]$ embeddings from K to K' , and hence the inner loop in Lines 6–7 iterates at most $[K : \mathbb{Q}]$ times for each fixed (K, K') .

For the loop in Lines 8–12, we consider $K \in \mathcal{F}$ and pick $H \in \mathcal{P}$ so that L^H is isomorphic to K . By Lemma 3.5, the number of idempotents in I_K equals the number of blocks in C_H , and this number increases every time I_K is changed by the subroutines. On the other hand, the number of idempotents in I_K is at most $[K : \mathbb{Q}]$. So the loop in Lines 8–12 iterates $O(N)$ times. The claim about the running time easily follows.

Finally, note that the algorithm exits the loop in Lines 8–12 after an iteration iff all of the idempotent decompositions I_K remain the same in that iteration, in which case \mathcal{C} is already a strongly antisymmetric \mathcal{P} -scheme by Lemma 3.13, Lemma 3.14 and Lemma 3.15. \square

Remark. The input of the the algorithm contains \mathcal{F} whose size may be much greater than that of f and \tilde{f} . Therefore, the polynomiality of this algorithm in the size of its input does *not* imply that polynomial factoring over finite fields can be solved in

(deterministic) polynomial time. It does suggest, however, that the total degree of the fields in \mathcal{F} over \mathbb{Q} is the bottleneck of our factoring algorithm.

3.5 Compatibility and invariance test

The subroutine `CompatibilityAndInvarianceTest` is given in Algorithm 3. It has the effect of properly refining at least one partition in \mathcal{C} , unless \mathcal{C} is compatible and invariant.

Algorithm 3 `CompatibilityAndInvarianceTest`

```

1: for  $(K, K') \in \mathcal{F}^2$  and embedding  $\phi : K' \hookrightarrow K$  do
2:   for  $(\delta, \delta') \in I_K \times I_{K'}$  do
3:     if  $\bar{\phi}(\delta')\delta \notin \{0, \delta\}$  then  $\triangleright \bar{\phi} : \bar{O}_{K'} \rightarrow \bar{O}_K$  is induced from  $\phi$ 
4:        $I_K \leftarrow I_K - \{\delta\}$ 
5:        $I_K \leftarrow I_K \cup \{\bar{\phi}(\delta')\delta, (1 - \bar{\phi}(\delta'))\delta\}$ 
6:     return

```

This subroutine attempts to find a ring homomorphism $\bar{\phi} : \bar{O}_{K'} \rightarrow \bar{O}_K$ (induced from a field embedding $\phi : K' \rightarrow K$) and idempotents $\delta \in I_K, \delta' \in I_{K'}$ such that $\bar{\phi}(\delta')\delta$ equals neither δ nor zero. If such δ, δ' , and $\bar{\phi}$ are found, the subroutine updates I_K by replacing $\delta \in I_K$ with two new idempotents $\bar{\phi}(\delta')\delta$ and $(1 - \bar{\phi}(\delta'))\delta$, neither of which is zero. It has the effect of splitting each block $B_{\bar{\tau}_H(\delta)} \in C_H = P(I_H)$ corresponding to $\bar{\tau}_H(\delta) \in I_H$ (see Lemma 3.5) into two blocks, where H ranges over the subgroups in \mathcal{P} satisfying $L^H \cong K$. After the update, the subroutine halts.

Now we prove Lemma 3.13 as promised before.

Proof of Lemma 3.13. Polynomiality of the running time is straightforward. To prove the rest of the claim, we assume that no proper refinement is made, i.e. for all $K, K' \in \mathcal{F}$, $\delta \in I_K, \delta' \in I_{K'}$ and field embeddings $\phi : K' \hookrightarrow K$, we have $\bar{\phi}(\delta')\delta \in \{0, \delta\}$. Then we show that \mathcal{C} is compatible and invariant.

For $H \in \mathcal{P}$, the isomorphism τ_H identifies $L^H \in \mathcal{P}$ with a field $K \in \mathcal{F}$. So the condition above can be reformulated as follows: for all $H, H' \in \mathcal{P}$, $\delta \in I_H, \delta' \in I_{H'}$ and field embeddings $\phi : L^{H'} \hookrightarrow L^H$, we have $\bar{\phi}(\delta')\delta \in \{0, \delta\}$.

Now consider $H, H' \in \mathcal{P}$ satisfying $H \subseteq H'$ and elements $Hg, Hg' \in H \setminus G$ in the same block $B \in C_H = P(I_H)$. We want to show that $\pi_{H, H'}(Hg) = H'g$ and $\pi_{H, H'}(Hg') = H'g'$ are in the same block of $C_{H'}$. By Lemma 3.5, there exists an

idempotent $\delta \in I_H$ for which

$$B = \{Hh \in H \backslash G : {}^{h^{-1}}(i_{L^H, L}(\delta)) \equiv 1 \pmod{\bar{\mathfrak{Q}}_0}\} \quad (3.1)$$

holds, where $i_{L^H, L} : \bar{\mathcal{O}}_{L^H} \hookrightarrow \bar{\mathcal{O}}_L$ is induced from the natural inclusion $L^H \hookrightarrow L$. Choose ϕ to be the natural inclusion $L^{H'} \hookrightarrow L^H$. As $\delta = \sum_{\delta' \in I_{H'}} \bar{\phi}(\delta')\delta$, there exists an idempotent $\delta' \in I_{H'}$ such that $\bar{\phi}(\delta')\delta \neq 0$. By assumption, we have $\bar{\phi}(\delta')\delta = \delta$. Again by Lemma 3.5, the set B' given by

$$B' = \{H'h \in H' \backslash G : {}^{h^{-1}}(i_{L^{H'}, L}(\delta')) \equiv 1 \pmod{\bar{\mathfrak{Q}}_0}\} \quad (3.2)$$

is a block of $C_{H'} = P(I_{H'})$. We claim that $H'g, H'g' \in B'$. To see this, note that as $Hg \in B$, we have

$${}^{g^{-1}}(i_{L^H, L}(\bar{\phi}(\delta')\delta)) = {}^{g^{-1}}(i_{L^H, L}(\delta)) \equiv 1 \pmod{\bar{\mathfrak{Q}}_0}. \quad (3.3)$$

It implies ${}^{g^{-1}}(i_{L^H, L}(\bar{\phi}(\delta'))) \equiv 1 \pmod{\bar{\mathfrak{Q}}_0}$. Note that $i_{L^{H'}, L} = i_{L^H, L} \circ \bar{\phi}$. So we have ${}^{g^{-1}}(i_{L^{H'}, L}(\delta')) \equiv 1 \pmod{\bar{\mathfrak{Q}}_0}$ and hence $H'g \in B'$. Similarly, we have $H'g' \in B'$. So $H'g$ and $H'g'$ are in the same block of $C_{H'}$, as desired. Therefore \mathcal{C} is compatible.

Next consider $H, H' \in \mathcal{P}$ satisfying $H' = hHh^{-1}$ for some $h \in G$ and elements $Hg, Hg' \in H \backslash G$ in the same block B of C_H . We want to show that $c_{H, h}(Hg) = H'hg$ and $c_{H, h}(Hg') = H'hg'$ are in the same block of $C_{H'}$. Again by Lemma 3.5, there exists an idempotent $\delta \in I_H$ for which (3.1) holds. Choose ϕ to be the isomorphism $L^{H'} \rightarrow L^H$ sending $x \in L^{H'}$ to $h^{-1}x \in L^H$. So $\bar{\phi}$ sends $x \in \bar{\mathcal{O}}_{L^{H'}}$ to $h^{-1}x \in \bar{\mathcal{O}}_{L^H}$, or more pedantically, to

$$i_{L^H, L}^{-1} \left({}^{h^{-1}}(i_{L^{H'}, L}(x)) \right) \in L^H.$$

Again, as $\delta = \sum_{\delta' \in I_{H'}} \bar{\phi}(\delta')\delta$, there exists an idempotent $\delta' \in I_{H'}$ such that $\bar{\phi}(\delta')\delta \neq 0$. By assumption, we have $\bar{\phi}(\delta')\delta = \delta$. By Lemma 3.5, the set B' given by (3.2) is a block of $C_{H'} = P(I_{H'})$. We claim that $H'hg, H'hg' \in B'$. To see this, note that (3.3) holds since $Hg \in B$. It implies that

$${}^{(hg)^{-1}}(i_{L^{H'}, L}(\delta')) = {}^{g^{-1}}(i_{L^H, L}(\bar{\phi}(\delta'))) \equiv 1 \pmod{\bar{\mathfrak{Q}}_0}$$

and hence $H'hg \in B'$. Similarly, we have $H'hg' \in B'$. So $H'hg$ and $H'hg'$ are in the same block of $C_{H'}$, as desired. As $c_{H, h}$ is bijective, it maps blocks to blocks. Therefore \mathcal{C} is invariant. \square

3.6 Regularity test

In this section we implement the subroutine `RegularityTest`. It has the effect of properly refining at least one partition in \mathcal{C} if \mathcal{C} is compatible, invariant, but not regular.

A similar algorithm was proposed in [Evd94; Gao01] based on generalizations of the Euclidean algorithm for polynomials over rings. We take an alternative approach developed in [IKS09; Iva+12] based on a “free module test”:

Lemma 3.16 ([IKS09; Iva+12]). *There exists an algorithm `FreeModuleTest` that given a semisimple \mathbb{F}_p -algebra A and a finitely generated A -module M , returns a zero divisor a of A in polynomial time, such that a is zero only if M is a free A -module.*

For completeness, we prove Lemma 3.16 in Appendix B. In addition, we need the following subroutine.

Lemma 3.17. *There exists an algorithm `SplitByZeroDivisor` that given*

- *a semisimple \mathbb{F}_p -algebra R , an idempotent decomposition I of R , and an idempotent $\gamma \in I$,*
- *the ring $\bar{R} := R/(1 - \gamma)$, the quotient map $\pi : R \rightarrow \bar{R}$, and a zero divisor $a \neq 0$ of \bar{R} ,*

replaces $\gamma \in I$ with two nonzero idempotents γ_1, γ_2 satisfying $\gamma = \gamma_1 + \gamma_2$ in polynomial time.

The proof of Lemma 3.17 can be found in Appendix B as well. The subroutine `RegularityTest` is then implemented in Algorithm 4 below.

The subroutine enumerates $(K, K') \in \mathcal{F}^2$, the ring homomorphisms $\bar{\phi} : \bar{\mathcal{O}}_{K'} \rightarrow \bar{\mathcal{O}}_K$ (induced from the field embeddings $\phi : K' \rightarrow K$), and the idempotents $\delta \in I_K$, $\delta' \in I_{K'}$ satisfying $\bar{\phi}(\delta')\delta = \delta$. Line 3 and Line 4 compute the quotient rings $A = \bar{\mathcal{O}}_{K'}/(1 - \delta')$, $M = \bar{\mathcal{O}}_K/(1 - \delta)$ and the corresponding quotient maps. They are quotient rings of semisimple rings and hence also semisimple. By Lemma 3.4, the map $\bar{\phi}$ induces a ring homomorphism $\phi_{\delta, \delta'} : A \rightarrow M$ sending $u + (1 - \delta')$ to $\bar{\phi}(u) + (1 - \delta)$ for $u \in \bar{\mathcal{O}}_{K'}$, which we compute at Line 5. It gives M an A -algebra structure, and in particular an A -module structure. Then we call `FreeModuleTest`

Algorithm 4 RegularityTest

```

1: for  $(K, K') \in \mathcal{F}^2$  and embedding  $\phi : K' \hookrightarrow K$  do
2:   for  $(\delta, \delta') \in I_K \times I_{K'}$  satisfying  $\bar{\phi}(\delta')\delta = \delta$  do
3:     compute  $A = \bar{\mathcal{O}}_{K'}/(1 - \delta')$  and the quotient map  $\bar{\mathcal{O}}_{K'} \rightarrow A$ 
4:     compute  $M = \bar{\mathcal{O}}_K/(1 - \delta)$  and the quotient map  $\bar{\mathcal{O}}_K \rightarrow M$ 
5:     compute  $\phi_{\delta, \delta'} : A \rightarrow M$  sending  $u + (1 - \delta')$  to  $\bar{\phi}(u) + (1 - \delta)$  for
        $u \in \bar{\mathcal{O}}_{K'}$ , making  $M$  an  $A$ -algebra and hence an  $A$ -module
6:     call FreeModuleTest with the input  $A$  and  $M$  to obtain  $a \in A$ 
7:     if  $a \neq 0$  then
8:       call SplitByZeroDivisor to update  $I_{K'}$  using the zero divisor  $a$ 
9:     return

```

at Line 6 which returns a zero divisor a of A by Lemma 3.16. If $a \neq 0$, we call SplitByZeroDivisor (with the input $R = \bar{\mathcal{O}}_{K'}$, $I = I_{K'}$, $\gamma = \delta'$, $\bar{R} = A$, the quotient map $\bar{\mathcal{O}}_{K'} \rightarrow A$, and the zero divisor a) to update $I_{K'}$, so that δ' is replaced with two nonzero idempotents by Lemma 3.17. After the update, the subroutine halts.

Proof of Lemma 3.14. The subroutine obviously runs in time polynomial in $\log p$ and the size of \mathcal{F} . To prove the rest of the lemma, it suffices to show that a zero divisor $a \neq 0$ of A is always found in Line 6 if \mathcal{C} is compatible but not regular.

So assume \mathcal{C} is compatible but not regular. Then there exist $H, H' \in \mathcal{P}$ satisfying $H \subseteq H'$, $B \in C_H$, $B' \in C_{H'}$ and $H'g, H'g' \in B'$ such that

$$|\pi_{H, H'}^{-1}(H'g) \cap B| \neq |\pi_{H, H'}^{-1}(H'g') \cap B|. \quad (3.4)$$

By Lemma 3.5, there exist $\delta \in I_H$ and $\delta' \in I_{H'}$ such that

$$B = \{Hh \in H \setminus G : {}^{h^{-1}}(i_{L^H, L}(\delta)) \equiv 1 \pmod{\bar{\mathfrak{Q}}_0}\}$$

and

$$B' = \{H'h \in H' \setminus G : {}^{h^{-1}}(i_{L^{H'}, L}(\delta')) \equiv 1 \pmod{\bar{\mathfrak{Q}}_0}\}.$$

By (3.4) and compatibility of \mathcal{C} , we have $\pi_{H, H'}(B) \subseteq B'$. Let $\phi : L^{H'} \hookrightarrow L^H$ be the natural inclusion, which induces a ring homomorphism $\bar{\phi} : \bar{\mathcal{O}}_{L^{H'}} \rightarrow \bar{\mathcal{O}}_{L^H}$. We claim that $\bar{\phi}(\delta')\delta = \delta$ holds: assume to the contrary that it does not hold. Then there exists a maximal ideal \mathfrak{m} of $\bar{\mathcal{O}}_L$ such that

$$i_{L^H, L}(\delta) \equiv 1 \pmod{\mathfrak{m}} \quad \text{and} \quad i_{L^H, L}(\bar{\phi}(\delta')) = i_{L^{H'}, L}(\delta') \equiv 0 \pmod{\mathfrak{m}}.$$

Choose $h \in G$ such that $\mathfrak{m} = {}^h\bar{\mathfrak{Q}}_0$. Then we have

$${}^{h^{-1}}(i_{L^H,L}(\delta)) \equiv 1 \pmod{\bar{\mathfrak{Q}}_0} \quad \text{and} \quad {}^{h^{-1}}(i_{L^{H'},L}(\delta')) \equiv 0 \pmod{\bar{\mathfrak{Q}}_0}.$$

It follows that $Hh \in B$ and $\pi_{H,H'}(Hh) = H'h \notin B'$. But this contradicts $\pi_{H,H'}(B) \subseteq B'$. So $\bar{\phi}(\delta')\delta = \delta$ holds.

Define $A := \bar{\mathcal{O}}_{L^{H'}}/(1 - \delta')$ and $M := \bar{\mathcal{O}}_{L^H}/(1 - \delta)$. Let $\phi_{\delta,\delta'} : A \rightarrow M$ be the ring homomorphism sending $u + (1 - \delta')$ to $\bar{\phi}(u) + (1 - \delta)$ for $u \in \bar{\mathcal{O}}_{L^{H'}}$, making M an A -algebra and hence an A -module. We claim that M is not free over A . Assume to the contrary that M is a free A -module. Denote its rank over A by $k \in \mathbb{N}^+$. Define

$$\mathfrak{P} := ({}^g\bar{\mathfrak{Q}}_0 \cap \mathcal{O}_{L^{H'}})/p\mathcal{O}_{L^{H'}} \subseteq \bar{\mathcal{O}}_{L^{H'}} \quad \text{and} \quad \mathfrak{P}' := \mathfrak{P}/(1 - \delta') \subseteq A,$$

which are maximal ideals of $\bar{\mathcal{O}}_{L^{H'}}$ and of A respectively. Then $M/\mathfrak{P}'M$ is a free A/\mathfrak{P}' -module of rank k . On the other hand, we have the isomorphism

$$M/\mathfrak{P}'M \cong \bar{\mathcal{O}}_{L^H}/(\bar{\phi}(\mathfrak{P})\bar{\mathcal{O}}_{L^H} + (1 - \delta)\bar{\mathcal{O}}_{L^H}).$$

It follows from the Chinese remainder theorem that $M/\mathfrak{P}'M$ is isomorphic to $\prod_{\mathfrak{m} \in S} \bar{\mathcal{O}}_{L^H}/\mathfrak{m}$ where S denotes the set of the maximal ideals of $\bar{\mathcal{O}}_{L^H}$ containing both $\bar{\phi}(\mathfrak{P})$ and $1 - \delta$. As p splits completely in L^H , each direct factor $\bar{\mathcal{O}}_{L^H}/\mathfrak{m}$ is isomorphic to \mathbb{F}_p . So $M/\mathfrak{P}'M$ is a vector space of dimension $|S|$ of \mathbb{F}_p . On the other hand, as p splits completely in $L^{H'}$, we have $A/\mathfrak{P}' \cong \mathbb{F}_p$. So rank k of $M/\mathfrak{P}'M$ over A/\mathfrak{P}' equals $|S|$.

By Corollary 3.1, the maximal ideals of $\bar{\mathcal{O}}_{L^H}$ are of the form by $\mathfrak{P}_{Hh} := ({}^h\bar{\mathfrak{Q}}_0 \cap \mathcal{O}_{L^H})/p\mathcal{O}_{L^H}$ which correspond one-to-one to the cosets $Hh \in H \setminus G$. Each maximal ideal \mathfrak{P}_{Hh} contains $\bar{\phi}(\mathfrak{P})$ iff \mathfrak{P} is contained in

$$\bar{\phi}^{-1}(\mathfrak{P}_{Hh}) = ({}^h\bar{\mathfrak{Q}}_0 \cap \mathcal{O}_{L^{H'}})/p\mathcal{O}_{L^{H'}},$$

which, again by Corollary 3.1, holds iff $H'g = H'h$. And \mathfrak{P}_{Hh} contains $1 - \delta$ iff $i_{L^H,L}(1 - \delta) \in {}^h\bar{\mathfrak{Q}}_0$, which holds iff $Hh \in B$. So we have

$$k = |S| = |\{Hh \in B : H'g = H'h\}| = |\pi_{H,H'}^{-1}(H'g) \cap B|.$$

But the same proof shows $k = |\pi_{H,H'}^{-1}(H'g') \cap B|$. This is a contradiction to (3.4). Therefore M is not free over A .

Identify L^H (resp. $L^{H'}$) with a field in \mathcal{F} using the isomorphism τ_H (resp. $\tau_{H'}$) chosen in Section 3.4. By Lemma 3.16, the subroutine is guaranteed to find a nonzero element $a \in A$ in Line 6. It then updates an idempotent decomposition $I_{K'}$ and properly refines some partition in \mathcal{C} by Lemma 3.17, as desired. \square

3.7 Strong antisymmetry test

In this section, we implement the subroutine `StrongAntisymmetryTest`, which has the effect of properly refining at least one partition in \mathcal{C} if \mathcal{C} is a \mathcal{P} -scheme, but not a strongly antisymmetric \mathcal{P} -scheme.

This subroutine is based on an algorithm developed in [Rón92]:

Lemma 3.18 ([Rón92]). *Under GRH, there exists an algorithm `Automorphism` that given a ring A isomorphic to a finite product of \mathbb{F}_p and a nontrivial ring automorphism σ of A , returns a zero divisor $a \neq 0$ of A in polynomial time.*

For completeness, we provide a proof of Lemma 3.18 in Appendix B.

The subroutine `StrongAntisymmetryTest` is implemented in Algorithm 5 below.

Algorithm 5 `StrongAntisymmetryTest`

- 1: construct an edge-labeled directed graph $G = (V, E)$ where $V = \{(K, \delta) : K \in \mathcal{F}, \delta \in I_K\}$ and $E = \emptyset$
 - 2: **for** $(K, \delta) \in V$ **do**
 - 3: compute $A_{K,\delta} := \bar{\mathcal{O}}_K / (1 - \delta)$ and the quotient map $\bar{\mathcal{O}}_K \rightarrow A_{K,\delta}$
 - 4: **for** $((K, \delta), (K', \delta')) \in V^2$ and $\phi : K' \hookrightarrow K$ satisfying $\bar{\phi}(\delta')\delta = \delta$ **do**
 - 5: compute $\phi_{\delta,\delta'} : A_{K',\delta'} \rightarrow A_{K,\delta}$ sending $x + (1 - \delta')$ to $\bar{\phi}(x) + (1 - \delta)$
 - 6: **if** $\phi_{\delta,\delta'}$ is invertible **then**
 - 7: $E \leftarrow E \cup \{e, e'\}$, where the edge e is from (K', δ') to (K, δ) with label $\phi_{\delta,\delta'}$, and e' is from (K, δ) to (K', δ') with label $\phi_{\delta,\delta'}^{-1}$
 - 8: search a nontrivial automorphism σ of $A_{K,\delta}$ for some $(K, \delta) \in V$ that is a composition of maps in $\mathcal{L} := \{\phi_{\delta,\delta'} : \text{there exists an edge } e \in E \text{ with label } \phi_{\delta,\delta'}\}$
 - 9: **if** σ is found at Line 8 **then**
 - 10: call `Automorphism` on $(A_{K,\delta}, \sigma)$ to obtain a zero divisor $a \neq 0$ of $A_{K,\delta}$
 - 11: call `SplitByZeroDivisor` to update I_K using a
 - 12: **return**
-

The subroutine first constructs an edge-labeled directed graph $G = (V, E)$, where the vertex set is

$$V := \{(K, \delta) : K \in \mathcal{F}, \delta \in I_K\}$$

and each edge is labeled by a certain ring isomorphism to be determined later. Initially the edge set E is empty. For every vertex $(K, \delta) \in V$, we compute the ring $A_{K,\delta} := \bar{\mathcal{O}}_K / (1 - \delta)$ and the quotient map $\bar{\mathcal{O}}_K \rightarrow A_{K,\delta}$ at Line 3.

Then we enumerate $((K, \delta), (K', \delta')) \in V^2$ and $\phi : K' \hookrightarrow K$ for which $\bar{\phi}(\delta')\delta = \delta$ holds, and for each of them, we compute a ring homomorphism

$$\phi_{\delta, \delta'} : A_{K', \delta'} \rightarrow A_{K, \delta}$$

that sends $x + (1 - \delta')$ to $\bar{\phi}(x) + (1 - \delta)$ for $x \in \bar{O}_{K'}$. The map $\phi_{\delta, \delta'}$ is well defined by Lemma 3.4. If $\phi_{\delta, \delta'}$ is an isomorphism (i.e., invertible), we add to E an edge e from (K', δ') to (K, δ) with label $\phi_{\delta, \delta'}$, and also an edge e' from (K, δ) to (K', δ') with label $\phi_{\delta, \delta'}^{-1}$.

Next, at Line 8, we search a nontrivial automorphism σ of $A_{K, \delta}$, $(K, \delta) \in V$, such that σ is a composition of maps in \mathcal{L} , where

$$\mathcal{L} := \{\phi_{\delta, \delta'} : \text{there exists an edge } e \in E \text{ with label } \phi_{\delta, \delta'}\}.$$

We sketch a way of implementing this step in time polynomial in $\log p$ and the size of \mathcal{F} : note that the edges whose labels compose into a nontrivial automorphism form a cycle of G . So by computing the strongly connected components of G and restricting to each of them, we reduce to the case that G is strongly connected. Fix a vertex $(K_0, \delta_0) \in V$. For every $(K, \delta) \in V$, compute a ring isomorphism $\psi_{K, \delta} : A_{K_0, \delta_0} \rightarrow A_{K, \delta}$ that is a composition of maps in \mathcal{L} . These isomorphisms exist since we assume G is strongly connected, and they can be computed by, e.g., the breadth-first search algorithm. Then we may find a nontrivial automorphism σ , if it exists, by enumerating the maps $\phi_{\delta, \delta'} : A_{K', \delta'} \rightarrow A_{K, \delta}$ in \mathcal{L} and checking if the automorphism

$$\phi_{\delta, \delta'} \circ \psi_{K', \delta'} \circ \psi_{K, \delta}^{-1} : A_{K, \delta} \rightarrow A_{K, \delta}$$

of $A_{K, \delta}$ is nontrivial.

Finally, if a nontrivial automorphism σ of some ring $A_{K, \delta}$ is successfully found, we use it to update I_K as follows: run the algorithm `Automorphism` on the input $(A_{K, \delta}, \sigma)$ to obtain a zero divisor $a \neq 0$ of $\in A_{K, \delta}$. Then call `SplitByZeroDivisor` (with the input $R = \bar{O}_K$, $I = I_K$, $\gamma = \delta$, $\bar{R} = A_{K, \delta}$, the quotient map $\bar{O}_K \rightarrow A_{K, \delta}$, and the zero divisor a) to update I_K , so that δ is replaced with two nonzero idempotents by Lemma 3.17.

Now we analyze the subroutine. For $H \subseteq G$ and $B \in C_H$, there exists a unique idempotent $\delta = \delta_B \in I_H$ satisfying

$$B = \{Hh \in H \setminus G : h^{-1}(i_{L^H, L}(\delta)) \equiv 1 \pmod{\bar{\mathfrak{Q}}_0}\}.$$

See Definition 3.2 and Lemma 3.5. Write $A_{L^H, \delta}$ for the ring $\bar{\mathcal{O}}_{L^H}/(1 - \delta)$. The maximal ideals of $A_{L^H, \delta}$ are of the form $\mathfrak{m}/(1 - \delta)$ where \mathfrak{m} is a maximal ideal of $\bar{\mathcal{O}}_{L^H}$ containing $1 - \delta$. By Corollary 3.1, the map

$$Hg \mapsto \mathfrak{m}_{Hg} := ({}^g\mathfrak{Q}_0 \cap \mathcal{O}_{L^H})/p\mathcal{O}_{L^H}$$

is a one-to-one correspondence between the right cosets in $H \backslash G$ and the maximal ideals of $\bar{\mathcal{O}}_K$. And \mathfrak{m}_{Hg} contains $1 - \delta$ iff $i_{L^H, L}(1 - \delta) \in {}^g\mathfrak{Q}_0$, which holds iff $Hg \in B$. We conclude that the map

$$Hg \mapsto \mathfrak{m}_{Hg}/(1 - \delta)$$

is a one-to-one correspondence between the right cosets in B and the maximal ideals of $A_{L^H, \delta}$.

We also need the following technical lemma.

Lemma 3.19. *Suppose $H, H' \in \mathcal{P}$, $B \in C_H$, $B' \in C_{H'}$, $\tau : B \rightarrow B'$, and $\phi : L^{H'} \rightarrow L^H$ are in one following cases:*

1. $H \subseteq H'$, $\tau = \pi_{H, H'}|_B : B \rightarrow B'$, and $\phi : L^{H'} \rightarrow L^H$ is the natural inclusion.
2. $H' = hHh^{-1}$ for some $h \in G$, $\tau = c_{H, h}|_B : B \rightarrow B'$, and $\phi : L^{H'} \rightarrow L^H$ sends x to $h^{-1}x$.

Let $\delta := \delta_B \in I_H$ and $\delta' := \delta_{B'} \in I_{H'}$ (see Definition 3.2). Let $\bar{\phi} : \bar{\mathcal{O}}_{L^{H'}} \rightarrow \bar{\mathcal{O}}_{L^H}$ be induced from ϕ . Then $\bar{\phi}(\delta')\delta = \delta$ holds, so that the ring homomorphism

$$\phi_{\delta, \delta'} : A_{L^{H'}, \delta'} \rightarrow A_{L^H, \delta}$$

sending $x + (1 - \delta')$ to $\bar{\phi}(x) + (1 - \delta)$ is well defined by Lemma 3.4. Moreover, for $Hg \in B$, we have

$$\phi_{\delta, \delta'}^{-1}(\mathfrak{m}_{Hg}/(1 - \delta)) = \mathfrak{m}_{\tau(Hg)}/(1 - \delta').$$

Finally, the map $\phi_{\delta, \delta'}$ is an isomorphism if τ is a bijection.

Proof. We claim that for any $Hg \in H \backslash G$, it holds that $\bar{\phi}^{-1}(\mathfrak{m}_{Hg}) = \mathfrak{m}_{\tau(Hg)}$. Fix $Hg \in H \backslash G$. Note that $\bar{\phi}^{-1}(\mathfrak{m}_{Hg})$ is a prime (and hence maximal) ideal of $\bar{\mathcal{O}}_{L^{H'}}$. Therefore to prove the claim, it suffices to show $\bar{\phi}(\mathfrak{m}_{\tau(Hg)}) \subseteq \mathfrak{m}_{Hg}$. In the first case of the lemma, we have $\tau(Hg) = \pi_{H, H'}(Hg) = H'g$, and

$$\mathfrak{m}_{Hg} = ({}^g\mathfrak{Q}_0 \cap \mathcal{O}_{L^H})/p\mathcal{O}_{L^H} \quad \text{and} \quad \mathfrak{m}_{\tau(Hg)} = \mathfrak{m}_{H'g} = ({}^g\mathfrak{Q}_0 \cap \mathcal{O}_{L^{H'}})/p\mathcal{O}_{L^{H'}}.$$

As $\phi : L^{H'} \rightarrow L^H$ is the natural inclusion, we have $\phi(\mathfrak{m}_{\tau(Hg)}) \subseteq \mathfrak{m}_{Hg}$, as desired.

In the second case, we have $\tau(Hg) = c_{H,h}(Hg) = H'hg$, and

$$\mathfrak{m}_{Hg} = ({}^g\mathfrak{Q}_0 \cap \mathcal{O}_{L^H})/p\mathcal{O}_{L^H} \quad \text{and} \quad \mathfrak{m}_{\tau(Hg)} = \mathfrak{m}_{H'hg} = ({}^{hg}\mathfrak{Q}_0 \cap \mathcal{O}_{L^{H'}})/p\mathcal{O}_{L^{H'}}.$$

As $\phi : L^{H'} \rightarrow L^H$ sends x to $h^{-1}x$, again we have $\phi(\mathfrak{m}_{\tau(Hg)}) \subseteq \mathfrak{m}_{Hg}$. This proves the claim.

Next we prove $\bar{\phi}(\delta')\delta = \delta$. As $\bar{\mathcal{O}}_{L^H}$ is semisimple, it suffices to show that for any maximal ideal \mathfrak{m}_{Hg} containing $\bar{\phi}(\delta')$ also contains δ . Fix $Hg \in H \setminus G$ such that $\bar{\phi}(\delta') \in \mathfrak{m}_{Hg}$. Then δ' is contained in $\bar{\phi}^{-1}(\mathfrak{m}_{Hg}) = \mathfrak{m}_{\tau(Hg)}$. As $\delta' = \delta_{B'}$, we have $\tau(Hg) \notin B'$ and hence $Hg \notin B$. Finally, as $\delta = \delta_B$, we have $\delta \in \mathfrak{m}_{Hg}$, as desired.

The next claim that $\phi_{\delta,\delta'}^{-1}(\mathfrak{m}_{Hg}/(1-\delta)) = \mathfrak{m}_{\tau(Hg)}/(1-\delta')$ follows directly from $\bar{\phi}^{-1}(\mathfrak{m}_{Hg}) = \mathfrak{m}_{\tau(Hg)}$. Now assume τ is a bijection. The kernel of $\phi_{\delta,\delta'}$ is

$$\bigcap_{Hg \in B} \phi_{\delta,\delta'}^{-1}(\mathfrak{m}_{Hg}/(1-\delta)) = \bigcap_{Hg \in B} \mathfrak{m}_{\tau(Hg)}/(1-\delta') = \bigcap_{H'g \in B'} \mathfrak{m}_{H'g}/(1-\delta') = 0.$$

So $\phi_{\delta,\delta'}$ is injective. Also note that the dimension of $A_{L^H,\delta}$ (resp. $A_{L^{H'},\delta'}$) over \mathbb{F}_p equals its number of maximal ideals, which is $|B|$ (resp. $|B'|$). As τ is bijective, we have $|B| = |B'|$. So τ is an isomorphism. \square

Now we are ready to prove Lemma 3.15, as promised.

Proof of Lemma 3.15. Assume \mathcal{C} is a \mathcal{P} -scheme but not a strongly antisymmetric \mathcal{P} -scheme. By Lemma 3.18, it suffices to show that some maps in \mathcal{L} compose into a nontrivial automorphism of $A_{K,\delta}$ for some $(K, \delta) \in V$.

As \mathcal{C} is not strongly antisymmetric, there exist $k \in \mathbb{N}^+$, subgroups $H_0, \dots, H_k \in \mathcal{P}$, blocks $B_0 \in C_{H_0}, \dots, B_k \in C_{H_k}$, and maps $\sigma_1, \dots, \sigma_k$ satisfying

- σ_i is a bijective map from B_{i-1} to B_i ,
- σ_i is of the form $c_{H_{i-1},g}|_{B_{i-1}}, \pi_{H_{i-1},H_i}|_{B_{i-1}}$, or $(\pi_{H_i,H_{i-1}}|_{B_i})^{-1}$,
- $H_0 = H_k$ and $B_0 = B_k$,

and the composition $\tau := \sigma_k \circ \dots \circ \sigma_1$ is a nontrivial permutation of $B_0 = B_k$.

Let $\delta_i := \delta_{B_i} \in I_{H_i}$ and $A_{L^{H_i},\delta_i} := \bar{\mathcal{O}}_{L^{H_i}}/(1-\delta_i)$ for $0 \leq i \leq k$. By Lemma 3.19, for $i \in [k]$, there exists a ring isomorphism $\psi_i : A_{L^{H_i},\delta_i} \rightarrow A_{L^{H_{i-1}},\delta_{i-1}}$ such that

$$\psi_i^{-1}(\mathfrak{m}_{H_{i-1}g}/(1-\delta_{i-1})) = \mathfrak{m}_{\sigma_i(H_{i-1}g)}/(1-\delta_i)$$

holds for all $H_{i-1}g \in B_{i-1}$. Moreover, for $i \in [k]$, the map ψ_i is in one of the following two cases:

- ψ_i sends $x + (1 - \delta_i)$ to $\bar{\phi}_i(x) + (1 - \delta_{i-1})$ for $x \in \bar{\mathcal{O}}_{L^{H_i}}$, where ϕ_i is an embedding of L^{H_i} in $L^{H_{i-1}}$.
- ψ_i^{-1} sends $x + (1 - \delta_{i-1})$ to $\bar{\phi}_i(x) + (1 - \delta_i)$ for $x \in \bar{\mathcal{O}}_{L^{H_{i-1}}}$, where ϕ_i is an embedding of $L^{H_{i-1}}$ in L^{H_i} .

Here the first case occurs when σ_i is of the form $c_{H_{i-1},g}|_{B_{i-1}}$ or $\pi_{H_{i-1},H_i}|_{B_{i-1}}$, and the second one occurs when σ_i is of the form $(\pi_{H_i,H_{i-1}}|_{B_i})^{-1}$.

Consider the automorphism $\sigma := \psi_1 \circ \dots \circ \psi_k$ of $A_{L^{H_k},\delta_k} = A_{L^{H_0},\delta_0}$. We have

$$\sigma^{-1}(\mathfrak{m}_{H_0g}/(1 - \delta_0)) = \mathfrak{m}_{\tau(H_0g)}/(1 - \delta_0)$$

for all $H_0g \in B_0$. As τ is a nontrivial permutation of B_0 , there exists $H_0g \in B_0$ satisfying $\tau(H_0g) \neq H_0g$ and hence $\mathfrak{m}_{H_0g}/(1 - \delta_0) \neq \mathfrak{m}_{\tau(H_0g)}/(1 - \delta_0)$. So σ is a nontrivial automorphism.

Finally, identifying each field L^{H_i} with a field $K_i \in \mathcal{F}$ using the isomorphisms $\tau_{H_i} : K_i \rightarrow L^{H_i}$, we see that the ring isomorphisms ψ_i are identified with maps in \mathcal{L} , and they compose into a nontrivial automorphism of $A_{K_0, \bar{\tau}_{H_i}^{-1}(\delta_0)}$. Here K_0 is the unique field in \mathcal{F} isomorphic to L^{H_0} and $\bar{\tau}_{H_i}^{-1}(\delta_0) \in I_{K_0}$. The lemma follows. \square

3.8 Constructing a collection of number fields

The last ingredient of the \mathcal{P} -scheme algorithm is a subroutine that constructs a (\mathbb{Q}, g) -subfield system given a polynomial $g(X) \in \mathbb{Q}[X]$ irreducible over \mathbb{Q} .

This subroutine can be implemented in various ways, leading to algorithms with different running time. We mention two results of this kind: computing the splitting field of g , and computing a (\mathbb{Q}, g) -subfield system whose associated subgroup system is a system of stabilizers. For simplicity, we only state the results, deferring the proofs and the algorithms to Chapter 4 where we discuss the problem of constructing number fields in depth.

Computing the splitting field of a polynomial. The splitting of a polynomial over \mathbb{Q} can be effectively constructed by the following lemma.

Lemma 3.20. *There exists a deterministic algorithm that given a polynomial $g(X) \in \mathbb{Q}[X]$ irreducible over \mathbb{Q} , computes its splitting field $L(g)$ over \mathbb{Q} in time polynomial in $[L(g) : \mathbb{Q}]$ and the size of g .*

The proof is deferred to Chapter 4.

System of stabilizers. We also have an algorithm that computes a (\mathbb{Q}, g) -subfield system whose associated subgroup system is a system of stabilizers:

Lemma 3.21. *There exists a deterministic algorithm that given a polynomial $g(X) \in \mathbb{Q}[X]$ irreducible over \mathbb{Q} and a positive integer $m \leq \deg(g)$, computes a (\mathbb{Q}, g) -subfield system \mathcal{F} , such that the subgroup system associated with \mathcal{F} is the system of stabilizers of depth m over $G(g/\mathbb{Q})$ with respect to the action of $\text{Gal}(g/\mathbb{Q})$ on the set of roots of g in $L(g)$, where $L(g)$ denotes the splitting field of g over \mathbb{Q} . Moreover, the algorithm runs in time polynomial in $(\deg(g))^m$ and the size of g .*

The proof is again deferred to Chapter 4.

3.9 Putting it together

We combine the results in previous sections to obtain the \mathcal{P} -scheme algorithm. The pseudocode is given in Algorithm 6 below.

Algorithm 6 PschemeAlgorithm

Input: $f(X) \in \mathbb{F}_p[X]$ and its irreducible lifted polynomial $\tilde{f}(X) \in \mathbb{Z}[X]$

Output: factorization of f

- 1: call `ComputeNumberFields` to compute a (\mathbb{Q}, \tilde{f}) -subfield system \mathcal{F} such that (1) $F = \mathbb{Q}[X]/(\tilde{f}(X)) \in \mathcal{F}$, and (2) for some $H \in \mathcal{P}$ satisfying $L^H \cong F$, all strongly antisymmetric \mathcal{P} -schemes are discrete (resp. inhomogeneous) on H , where \mathcal{P} is the subgroup system over $G = \text{Gal}(\tilde{f}/\mathbb{Q})$ associated with \mathcal{F}
 - 2: call `ComputePscheme` on the input (p, \mathcal{F}) to obtain I_K for $K \in \mathcal{F}$
 - 3: call `ExtractFactors` to extract a factorization of f from I_F , and output it
-

The subroutine `ComputeNumberFields` at Line 1 is the generic part of the algorithm and can be implemented in various ways. It is supposed to compute a (\mathbb{Q}, \tilde{f}) -subfield system \mathcal{F} such that $F \in \mathcal{F}$, and the associated subgroup system \mathcal{P} over G satisfies a certain combinatorial property (see Theorem 3.9 below). The latter condition is used to show that the factoring algorithm always produces the complete factorization (resp. a proper factorization) of f .

The algorithm `ComputePscheme` (see Section 3.4) at Line 2 takes the input (p, \mathcal{F}) and outputs data that includes the idempotent decompositions I_K for $K \in \mathcal{F}$. Finally, we call the subroutine `ExtractFactors` (see Section 3.3) at Line 3 to extract a factorization of f from I_F .

The following theorem is the main result of this chapter.

Theorem 3.9 (Theorem 3.2 restated). *Suppose there exists a deterministic algorithm that given a polynomial $g(X) \in \mathbb{Z}[X]$ irreducible over \mathbb{Q} , constructs a (\mathbb{Q}, g) -subfield system \mathcal{F} in time $T(g)$ such that*

- $\mathbb{Q}[X]/(g(X))$ is in \mathcal{F} , and
- for some $H \in \mathcal{P}$ satisfying $(L(g))^H \cong \mathbb{Q}[X]/(g(X))$, all strongly antisymmetric \mathcal{P} -schemes are discrete (resp. inhomogeneous) on H , where \mathcal{P} is the subgroup system over $\text{Gal}(g/\mathbb{Q})$ associated with \mathcal{F} , and $L(g)$ is the splitting field of g over \mathbb{Q} .

Then under GRH, there exists a deterministic algorithm that given a polynomial $f(X) \in \mathbb{F}_p[X]$ satisfying Condition 3.1 and an irreducible lifted polynomial $\tilde{f}(X) \in \mathbb{Z}[X]$ of f , outputs the complete factorization (resp. a proper factorization) of f over \mathbb{F}_p in time polynomial in $T(\tilde{f})$ and the size of the input.

Proof. Consider the algorithm `PschemeAlgorithm` above and implement the subroutine `ComputeNumberFields` using the hypothetical algorithm in the theorem. Choose $g = \tilde{f}$. By Theorem 3.8, the \mathcal{P} -collection $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ defined by $C_H = P(\bar{\tau}_H(I_K))$ is a strongly antisymmetric \mathcal{P} -scheme. By the second condition in the theorem, we have $C_H = \infty_{H \setminus G}$ (resp. $C_H \neq 0_{H \setminus G}$) for some $H \in \mathcal{P}$ satisfying $L^H \cong F$. So the corresponding idempotent decomposition I_F is complete (resp. proper). By Theorem 3.7, the algorithm outputs the complete factorization (resp. a proper factorization) of f over \mathbb{F}_p .

The subroutine `ComputeNumberFields` runs in time $T(\tilde{f})$. In particular, the size of \mathcal{F} is bounded by $T(\tilde{f})$. The claim about the running time then follows from Theorem 3.8 and Theorem 3.7. \square

By Theorem 3.9 and Lemma 3.21, we have a deterministic factoring algorithm whose running time is related to the notations $d(G)$ and $d'(G)$ introduced in Definition 2.8:

Corollary 3.2. *Under GRH, there exists a deterministic algorithm that, given a polynomial $f(X) \in \mathbb{F}_p[X]$ of degree $n \in \mathbb{N}^+$ satisfying Condition 3.1 and an irreducible⁸ lifted polynomial $\tilde{f}(X) \in \mathbb{Z}[X]$ of f , computes the complete factorization (resp. a proper factorization) of f over \mathbb{F}_p in time polynomial in $n^{d(G)}$ (resp. $n^{d'(G)}$) and the size of the input, where G is the permutation group $\text{Gal}(f/\mathbb{Q})$ acting on the set of roots of \tilde{f} .*

The unifying framework via the \mathcal{P} -scheme algorithm. The \mathcal{P} -scheme algorithm and the underlying notion of \mathcal{P} -schemes provide a unifying framework for deterministic polynomial factoring over finite fields. To illustrate this point, we show that the main results achieved by known factoring algorithms [Hua91a; Hua91b; Rón88; Rón92; Evd94; IKS09] can be easily derived from Theorem 3.9 or Corollary 3.2 for the special case that the input polynomial satisfies Condition 3.1 (the general case is solved in Chapter 5).

Suppose we want to factorize $f(X) \in \mathbb{F}_p[X]$ given a (possibly reducible) lifted polynomial $\tilde{f}(X) \in \mathbb{Z}[X]$ of f . We reduce to the case that the lifted polynomial is irreducible as follows: first use the factoring algorithm for rational polynomials [LLL82] to factorize \tilde{f} into its irreducible factors $f_1(X), \dots, f_k(X) \in \mathbb{Q}[X]$ over \mathbb{Q} in polynomial time. By Gauss Lemma (see [Lan02, Section IV.2]), we may assume each factor $\tilde{f}_i(X)$ lies in $\mathbb{Z}[X]$. Then the problem of factoring $f(X)$ is reduced to the problem of factoring each $f_i(X) := \tilde{f}_i(X) \bmod p \in \mathbb{F}_p[X]$ with the aid of its irreducible lifted polynomial $\tilde{f}_i(X)$. Moreover, for $i \in [k]$, the Galois group $\text{Gal}(\tilde{f}_i(X)/\mathbb{Q})$ is a quotient group of $\text{Gal}(\tilde{f}/\mathbb{Q})$, and hence $|\text{Gal}(\tilde{f}_i(X)/\mathbb{Q})| \leq |\text{Gal}(\tilde{f}(X)/\mathbb{Q})|$.

So assume \tilde{f} is irreducible over \mathbb{Q} . Choose $\mathcal{F} = \{F, L\}$ where $F = \mathbb{Q}[X]/(\tilde{f}(X))$ and L is the splitting field of \tilde{f} over \mathbb{Q} . Compute \mathcal{F} in time polynomial in $[L : \mathbb{Q}] = |\text{Gal}(\tilde{f}(X)/\mathbb{Q})|$ and the size of \tilde{f} using Lemma 3.20. By Lemma 2.4, all antisymmetric \mathcal{P} -schemes are discrete on H for all $H \in \mathcal{P}$ since the trivial subgroup $\{e\}$ is in \mathcal{P} . Therefore by Theorem 3.9 and the reduction above, we have

Theorem 3.10 ([Rón92]). *Under GRH, there exists a deterministic algorithm that, given a polynomial $f(X) \in \mathbb{F}_p[X]$ satisfying Condition 3.1 and a lifted polynomial $\tilde{f}(X) \in \mathbb{Z}[X]$ of f , computes the complete factorization of f over \mathbb{F}_p in time polynomial in $|\text{Gal}(\tilde{f}/\mathbb{Q})|$ and the size of the input.*

⁸The assumption that \tilde{f} is irreducible is not necessary, and can be avoided by adapting Lemma 3.21. We omit the details.

When $\text{Gal}(\tilde{f}/\mathbb{Q})$ is abelian, it acts semiregularly on the set of roots of \tilde{f} . So we have $|\text{Gal}(\tilde{f}/\mathbb{Q})| \leq \deg(f)$ (the equality holds iff \tilde{f} is irreducible over \mathbb{Q}). Then Theorem 3.10 gives

Corollary 3.3 ([Hua91a; Hua91b]). *Under GRH, there exists a deterministic algorithm that, given a polynomial $f(X) \in \mathbb{F}_p[X]$ satisfying Condition 3.1 and a lifted polynomial of f with an abelian Galois group, computes the complete factorization of f over \mathbb{F}_p in polynomial time.*

Suppose only the polynomial f is known. Let $n = \deg(f)$. We may lift f to a degree- n polynomial $\tilde{f}(X) \in \mathbb{Z}[X]$ such that all coefficients of \tilde{f} are in the interval $[0, p-1]$. So the size of \tilde{f} is $O(n \log p)$. Reduce to the case that \tilde{f} is irreducible over \mathbb{Q} as above. As $\text{Gal}(\tilde{f}(X)/\mathbb{Q})$ is a subgroup of $\text{Sym}(n)$, we derive the following theorem from Theorem 3.10.

Theorem 3.11 ([Rón88; Rón92]). *Under GRH, there exists a deterministic algorithm that, given a polynomial $f(X) \in \mathbb{F}_p[X]$ of degree $n \in \mathbb{N}^+$ that satisfies Condition 3.1, computes the complete factorization of f in time polynomial in $n!$ and $\log p$.*

Alternatively, Theorem 3.11 can be derived from Corollary 3.2 by noting $d(G) \leq n-1$ (where $G = \text{Gal}(\tilde{f}/\mathbb{Q})$). Similarly, using the bound $d(G) = O(\log n)$ in Lemma 2.6, we derive the following theorem from Corollary 3.2.

Theorem 3.12 ([Evd94; IKS09]). *Under GRH, there exists a deterministic algorithm that, given a polynomial $f(X) \in \mathbb{F}_p[X]$ of degree $n \in \mathbb{N}^+$ satisfying Condition 3.1, computes the complete factorization of f over \mathbb{F}_p in time polynomial in $n^{\log n}$ and $\log p$.*

By Corollary 3.2 and Lemma 2.18, we have

Theorem 3.13 ([Rón88; IKS09]). *Under GRH, there exists a deterministic algorithm that, given a polynomial $f(X) \in \mathbb{F}_p[X]$ of degree $n > 1$ satisfying Condition 3.1, computes a proper factorization of f over \mathbb{F}_p in time polynomial in n^ℓ and $\log p$, where ℓ is the least prime factor of n .*

In latter chapters, we also prove (and generalize) the main result of [Evd92] using the \mathcal{P} -scheme algorithm. It states that polynomial factoring over finite fields can be solved in deterministic polynomial time under GRH given a lifted polynomial that has a solvable Galois group. For more details, see Theorem 4.3 and Theorem 5.13.

Chapter 4

CONSTRUCTING NUMBER FIELDS

In this chapter, we discuss the problem of constructing number fields using a polynomial $g(X) \in \mathbb{Q}[X]$ irreducible over \mathbb{Q} . In particular, we prove Lemma 3.20 and Lemma 3.21 as promised before.

In fact, we consider the more general problem of constructing *relative number fields*, which we explain now.

Relative number fields. Recall that a number field K is encoded using the minimal polynomial $h(X) \in \mathbb{Q}[X]$ of a primitive element α of K over \mathbb{Q} , i.e., $K = \mathbb{Q}(\alpha)$. Suppose K_0 is a number field encoded in this way. A *relative number field* K over K_0 is a number field containing K_0 , encoded by the minimal polynomial $h(X) \in K_0[X]$ of a primitive element α of K over K_0 (i.e. $K = K_0(\alpha)$). We regard K as a K_0 -algebra by maintaining its structure constants in the standard K_0 -basis

$$\{1 + (h(X)), X + (h(X)), \dots, X^{d-1} + (h(X))\},$$

where $d = [K : K_0]$. Note that when $K_0 = \mathbb{Q}$, this is the usual way we encode a number field.

Given a number field K_0 , we discuss various techniques of constructing relative number fields over K_0 given a polynomial $g(X) \in K_0[X]$ irreducible over K_0 . In particular, we discuss the technique of adjoining roots of polynomials and use it to prove Lemma 3.20 and Lemma 3.21.

Motivated by the \mathcal{P} -scheme algorithm in Chapter 2, we consider the problem of constructing a collection of (relative) number fields using $g(X)$, such that for the associated subgroup system \mathcal{P} , all strongly antisymmetric \mathcal{P} -schemes are discrete (resp. inhomogeneous) on a distinguished subgroup $H \in \mathcal{P}$. We describe a reduction of this problem to the case that the Galois group of $g(X)$ is a *primitive* permutation group. The idea was essentially introduced in [LM85], leading to a polynomial-time algorithm that determines if a given rational polynomial is solvable.¹ It was also used in [Evd92] to obtain a polynomial-time factoring algorithm

¹A rational polynomial $g(X) \in \mathbb{Q}[X]$ is solvable if its roots are expressible in the field operations and radicals. It is equivalent to the solvability of the Galois group $\text{Gal}(g/\mathbb{Q})$.

for $f(X) \in \mathbb{F}_p[X]$, provided that a solvable polynomial $\tilde{f}(X) \in \mathbb{Z}[X]$ lifting $f(X)$ is given. We reproduce the main result of [Evd92] for the case that f satisfies Condition 3.1. For the general case, see Chapter 5.

We note that most results in this chapter are essentially known in the literature, except that we state them in a relative setting or in the terminology of \mathcal{P} -schemes. In particular, the discussion about algebraic numbers in Section 4.1 follows [WR76], and the techniques of constructing number fields are mostly folklore or from [Lan84; LM85; Evd92].

Outline of the chapter. Notations and preliminaries are given in Section 4.1. In particular, we define the *complexity* of a subgroup system, which is used to bound the size of a collection of (relative) number fields and the running time of the algorithms. This notion also plays a role in subsequent chapters. In Section 4.2, we discuss the technique of constructing (relative) number fields by adjoining roots of a polynomial, and use it to prove Lemma 3.20 and Lemma 3.21. In Section 4.3, we establish the reduction to primitive Galois groups and use it to prove the main result of [Evd92] for the special case that $f(X) \in \mathbb{F}_p[X]$ satisfies Condition 3.1. Finally, we discuss some other techniques in Section 4.4. These techniques are not directly used in the thesis, but may still have their own interest.

4.1 Preliminaries

Let K and K' be relative number fields over a number field K_0 . We say an embedding (resp. isomorphism) $\tau : K \rightarrow K'$ is an embedding (resp. isomorphism) *over* K_0 if τ is K_0 -linear, i.e., $\tau(ax) = a\tau(x)$ for all $a \in K_0$ and $x \in K$. By choosing $x = 1$, we see that this is equivalent to $\tau(a) = a$ for all $a \in K_0$. We write $K \cong_{K_0} K'$ for the statement that K is isomorphic to K' over K_0 .

(K_0, g) -subfield systems and the associated subgroup systems. We generalize the notion of (\mathbb{Q}, g) -subfield systems (Definition 3.3) and the associated subgroup systems (Definition 3.4) as follows:

Definition 4.1 ((K_0, g) -subfield system). *Let K_0 be a number field. Let $g(X)$ be a polynomial in $K_0[X]$ with the splitting field L over K_0 . Let \mathcal{F} be a collection of relative number fields over K_0 such that (1) the fields in \mathcal{F} are mutually non-isomorphic over K_0 , and (2) each field $K' \in \mathcal{F}$ is isomorphic to a subfield of L over K_0 . We say \mathcal{F} is a (K_0, g) -subfield system.*

Definition 4.2. Let $g(X)$ be a polynomial in $K_0[X]$ with the splitting field L over K_0 . Let \mathcal{F} be a (K_0, g) -subfield system. Define \mathcal{P}^\sharp to be the poset of subfields of L that includes all the fields isomorphic to those in \mathcal{F} over K_0 :

$$\mathcal{P}^\sharp := \{K' \subseteq L : K' \cong_{K_0} K \text{ for some } K \in \mathcal{F}\}.$$

By Galois theory, it corresponds to a poset \mathcal{P} of subgroups of $\text{Gal}(g/K_0)$, given by

$$\mathcal{P} := \{H \subseteq \text{Gal}(g/K_0) : L^H \in \mathcal{P}^\sharp\},$$

which is closed under conjugation in $\text{Gal}(g/K_0)$, and hence is a subgroup system over $\text{Gal}(g/K_0)$. We say \mathcal{P} and \mathcal{P}^\sharp are associated with \mathcal{F} .

The complexity of a subgroup system. The size of a (K_0, g) -subfield system \mathcal{F} is primarily controlled by the total degree of the fields in \mathcal{F} over K_0 , which is the number of coefficients in K_0 we need to maintain. We relate this quantity to the complexity of a subgroup system, defined as follows.

Definition 4.3 (complexity of a subgroup system). Suppose \mathcal{P} is a subgroup system over a finite group G . Then G acts on \mathcal{P} by conjugation, i.e., $g \in G$ sends $H \in \mathcal{P}$ to $gHg^{-1} \in \mathcal{P}$. Let $\mathcal{P}_0 \subseteq \mathcal{P}$ be a complete set of representatives of the G -orbits under this action. Define the complexity of \mathcal{P} to be

$$c(\mathcal{P}) := \sum_{H \in \mathcal{P}_0} [G : H].$$

As conjugate subgroups have the same order, the complexity $c(\mathcal{P})$ is well defined. And we have

Lemma 4.1. For a (K_0, g) -subfield system \mathcal{F} , the total degree of the fields in \mathcal{F} over K_0 equals $c(\mathcal{P})$, where \mathcal{P} is the subgroup system associated with \mathcal{F} .

Proof. Conjugate subgroups correspond to conjugate subfields under the Galois correspondence. So for $K \in \mathcal{F}$ there exists a unique subgroup $H \in \mathcal{P}_0$ satisfying $L^H \cong_{K_0} K$. And the map $K \mapsto H$ is a one-to-one correspondence between \mathcal{F} and \mathcal{P}_0 . Finally note that $[K : K_0] = [G : H]$ for H corresponding to K . \square

The following lemma bounds the complexity of a system of stabilizers.

Lemma 4.2. *Let G be a finite group acting on a finite set S . Let $m \in \mathbb{N}^+$ and $m' = \min\{|S|, m\}$. Let \mathcal{P} be the system of stabilizers of depth m' with respect to the action of G on S . Then*

$$c(\mathcal{P}) \leq \sum_{k=1}^{m'} \prod_{i=1}^k (|S| - i) = O(|S|^{m'}).$$

Proof. Replacing m with m' does not change \mathcal{P} . So we may assume $m = m' \leq |S|$. When $|S| \geq 2$, we have

$$\sum_{k=1}^m \prod_{i=0}^{k-1} (|S| - i) \leq \sum_{k=1}^m |S|^k = O(|S|^m).$$

The same holds trivially when $|S| = 1$.

Next we prove $c(\mathcal{P}) \leq \sum_{k=1}^m \prod_{i=1}^k (|S| - i)$. Let $\mathcal{P}_0 \subseteq \mathcal{P}$ be as in Definition 4.3. It suffices to find an injective map

$$\tau : \coprod_{H \in \mathcal{P}_0} H \backslash G \hookrightarrow \prod_{k=1}^m S^{(k)},$$

since the cardinality of $\coprod_{H \in \mathcal{P}_0} H \backslash G$ is $c(\mathcal{P})$, whereas the cardinality of $\prod_{k=1}^m S^{(k)}$ is $\sum_{k=1}^m \prod_{i=1}^k (|S| - i)$.

For each $k \in [m]$, the group G acts diagonally on $S^{(k)}$. For each $H \in \mathcal{P}_0$, we pick $k = k(H) \leq m$ and $x = x(H) \in S^{(k)}$ such that $H = G_x$ with respect to the diagonal action. By Lemma 2.1, we have an injective map $H \backslash G \rightarrow S^{(k)}$ whose image is the G -orbit of x . These maps altogether give the map τ . To show τ is injective, it suffices to show that for different $H, H' \in \mathcal{P}_0$, the coset spaces $H \backslash G$ and $H' \backslash G$ are mapped to different G -orbits. Assume to the contrary that they are mapped to the the same G -orbit O . So $x(H), x(H') \in O$. Then $k(H) = k(H')$ and $x(H') = {}^g(x(H))$ for some $g \in G$. But then we have

$$H' = G_{x(H')} = G_{g x(H)} = g G_{x(H)} g^{-1} = g H g^{-1},$$

which is a contradiction to the choice of \mathcal{P}_0 . So τ is injective. \square

Algebraic numbers. The fields in a (K_0, g) -subfield system \mathcal{F} are encoded by polynomials in $K_0[X]$. So to bound the size of \mathcal{F} , we also need to bound the size of the coefficients of these polynomials, which are algebraic numbers in K_0 . This is closely related to the following definition, introduced in [WR76].

Definition 4.4. For an algebraic number α , define $\|\alpha\|$ to be the greatest absolute value of $i(\alpha) \in \mathbb{C}$ where i ranges over the embeddings of $\mathbb{Q}(\alpha)$ in \mathbb{C} .²

For algebraic numbers α, β , we clearly have $\|\alpha + \beta\| \leq \|\alpha\| + \|\beta\|$ and $\|\alpha \cdot \beta\| \leq \|\alpha\| \cdot \|\beta\|$.

The following lemma relates the size of an algebraic number $\alpha \in K_0$ (i.e., the number of bits used to encode α in K_0) to $\|\alpha\|$.

Lemma 4.3. Suppose K_0 is a number field encoded by a polynomial $h(X) \in \mathbb{Q}[X]$ irreducible over \mathbb{Q} of degree n and size s_0 . Let α be an algebraic number in K_0 of size s . Let D be the smallest positive integer such that $D\alpha$ is an algebraic integer. Then s is polynomial in $\log \|\alpha\|$, $\log D$ and s_0 . Conversely, $\log \|\alpha\|$ and $\log D$ are polynomial in s and s_0 .

Proof. Suppose $h(X) = \sum_{i=0}^n c_i X^i$ where $n = \deg(h)$ and $c_i \in \mathbb{Q}$ for all i . By substituting X with X/k for some large enough $k \in \mathbb{N}^+$ and clearing the denominators, we may assume $h(X) \in \mathbb{Z}[X]$ and $c_n = 1$. Both the encoding of h and that of α use at least n coefficients in \mathbb{Q} . So we have $s, s_0 \geq n$.

The algebraic number $\alpha \in K_0$ is encoded by the constants $d_0, \dots, d_{n-1} \in \mathbb{Q}$ satisfying

$$\alpha = \sum_{i=0}^{n-1} d_i \beta^i, \quad (4.1)$$

where β is a root of h in K_0 . So we have $\|\alpha\| \leq \sum_{i=0}^{n-1} |d_i| \|\beta\|^i$. It was shown in [WR76] that $\|\beta\| \leq \sum_{i=0}^{n-1} |c_i|$. And we clearly have $\log |c_i| \leq s_0$ and $\log |d_i| \leq s$ for $0 \leq i \leq n-1$. It follows that $\log \|\alpha\|$ is polynomial in s and s_0 .

Let $D' \in \mathbb{N}^+$ be the least common multiple of the denominators of d_i . As $h(X) \in \mathbb{Z}[X]$ and $c_n = 1$, we know β is an algebraic integer. Then $D'\alpha$ is also an algebraic integer by (4.1). So D is bounded by D' . It follows that $\log D$ is polynomial in s and s_0 . Then the second claim of the lemma is proved.

For the first claim, it suffices to show that the size of each d_i is polynomial in $\log \|\alpha\|$, $\log D$ and s_0 . This follows from [WR76, Section 7 and Lemma 8.3]. \square

The following lemma relates the size of the minimal polynomial of an algebraic number α over a number field K_0 to $\|\alpha\|$.

² $\|\alpha\|$ is called the size of α in [WR76]. We reserve the term *size* (of an object) for the number of bits used to encode an object in an algorithm.

Lemma 4.4. *Suppose K_0 is a number field encoded by a rational polynomial irreducible over \mathbb{Q} of size s_0 (let $s_0 = 1$ if $K_0 = \mathbb{Q}$). Let α be an algebraic number, and let D be the smallest positive integer such that $D\alpha$ is an algebraic integer. Let $h(X) \in K_0[X]$ be the minimal polynomial of α whose size is s and degree is n . Then s is polynomial in $\log \|\alpha\|$, $\log D$, s_0 and n . Conversely, $\log \|\alpha\|$ and $\log D$ are polynomial in s and s_0 .*

Proof. We clearly have $n \leq s$. Suppose $h(X) = \sum_{i=0}^n c_i X^i$, where $c_i \in K_0$ and $c_n = 1$. It was as shown in [WR76] that $\|\alpha\| \leq \sum_{i=0}^{n-1} \|c_i\|$. It follows from Lemma 4.3 that $\log \|\alpha\|$ is polynomial in s and s_0 .

Note that for sufficiently large $k \in \mathbb{N}^+$ that is polynomial in s and s_0 , the coefficients of the polynomial $k^n h(X/k)$ are all algebraic integers. It follows that $k\alpha$ is an algebraic integer (cf. [AM69, Corollary 5.4]). So D is bounded by k and hence is polynomial in s and s_0 . Then the second claim of the lemma is proved.

For the first claim, we may assume α is an algebraic integer by replacing α with $D\alpha$ and c_i with $D^{n-i}c_i$. Then any conjugate α' of α over \mathbb{Q} is also an algebraic integer, and $\|\alpha'\| = \|\alpha\|$. For $0 \leq i \leq n-1$, the coefficient c_i of h is (up to sign) given by the i th elementary symmetric polynomial in a subset of conjugates of α over \mathbb{Q} . It follows from Lemma 4.3 that the size of each c_i is polynomial in $\log \|\alpha\|$, $\log D$, s_0 and n . So s is polynomial in $\log \|\alpha\|$, $\log D$, s_0 and n as well. \square

Finding a primitive element over \mathbb{Q} . Suppose $K_0 = \mathbb{Q}(\alpha)$ is a number field encoded by the minimal polynomial of a primitive element α over \mathbb{Q} , and $K = K_0(\beta)$ is a relative number field over K_0 , encoded by the minimal polynomial of a primitive element β over K_0 . We would like to represent K directly in the form $\mathbb{Q}(\gamma)$, encoded by the minimal polynomial of a primitive element γ over \mathbb{Q} . The first step is to find such an element γ , which can be achieved using a constructive version of the primitive element theorem (see, e.g., [Wae91]). For completeness, we give the details as follows.

Lemma 4.5. *Suppose K_0 is a number field and α, β are algebraic numbers. Let $d = [K_0(\alpha, \beta) : K_0]$. Then $k\alpha + \beta$ is a primitive element of $K_0(\alpha, \beta)$ over K_0 for some integer $k \in [1, d+1]$.*

Proof. Consider a “bad” nonzero integer k for which $K_0(k\alpha + \beta)$ is a proper subfield of $K_0(\alpha, \beta)$. Let L be the Galois closure of $K_0(\alpha, \beta)/K_0$. Then by the fundamental

theorem of Galois theory, there exists an automorphism ϕ of L fixing $K_0(k\alpha + \beta)$ but not $K_0(\alpha, \beta)$. Then either $\phi(\alpha) \neq \alpha$ or $\phi(\beta) \neq \beta$. As ϕ fixes $k\alpha + \beta$, we have $k\phi(\alpha) + \phi(\beta) = \phi(k\alpha + \beta) = k\alpha + \beta$, from which we see that actually $\phi(\alpha) \neq \alpha$ and $\phi(\beta) \neq \beta$ both hold. Then k is determined by $\phi(\alpha)$ and $\phi(\beta)$ via $k = (\phi(\beta) - \beta)/(\alpha - \phi(\alpha))$. So the number of bad choices of k is bounded by the number of $(\phi(\alpha), \phi(\beta))$ where ϕ ranges over the automorphisms of L fixing K_0 . The later is the cardinality of the orbit of (α, β) under the action of $\text{Gal}(L/K_0)$. By the orbit-stabilizer theorem, it equals

$$[\text{Gal}(L/K_0) : \text{Gal}(L/K_0(\alpha, \beta))] = [K_0(\alpha, \beta) : K_0] = d.$$

So there are at most d bad choices of k . The lemma follows since $[1, d + 1]$ contains more than d integers. \square

This gives an efficient algorithm of finding a primitive element over \mathbb{Q} :

Lemma 4.6. *There exists a polynomial-time algorithm that given a number field K_0 and a relative number field K over K_0 , find a primitive element γ of K over \mathbb{Q} and its minimal polynomial $h(X) \in \mathbb{Q}[X]$ over \mathbb{Q} .*

Proof. Suppose K_0 is encoded by a polynomial $g(X) \in \mathbb{Q}[X]$ irreducible over \mathbb{Q} , and K is encoded by a polynomial $g'(X) \in K_0[X]$ irreducible over K_0 . Then we are explicitly given a root α of $g(X)$ and a root β of $g'(X)$ in K , and $K = \mathbb{Q}(\alpha, \beta)$.

Enumerate the integers $k \in [1, d + 1]$, where $d = [K : \mathbb{Q}]$. For each k , we compute $\gamma = k\alpha + \beta \in K$, and then compute its minimal polynomial $h(X) \in \mathbb{Q}[X]$ over \mathbb{Q} by solving linear equations over \mathbb{Q} . This step runs in polynomial time by Lemma 4.4. Output γ and h whenever $\deg(h) = [K : \mathbb{Q}]$. By Lemma 4.5, a primitive element γ is guaranteed to be found. \square

By computing a primitive element over \mathbb{Q} , we can efficiently turn a relative number field into an ordinary number field:

Corollary 4.1. *There exists a polynomial-time algorithm that given a number field K_0 and a relative number field K over K_0 , computes an ordinary number field K' , a \mathbb{Q} -basis B of K , and an isomorphism $\phi : K \rightarrow K'$ encoded by $\phi(x) \in K'$ for $x \in B$.*

Proof. Find a primitive element γ of K over \mathbb{Q} and its minimal polynomial $h(X) \in \mathbb{Q}[X]$ over \mathbb{Q} using Lemma 4.6. Compute $K' := \mathbb{Q}[X]/(h(X))$ and $B = \{1, \gamma, \gamma^2, \dots, \gamma^{d-1}\}$, where $d = [K : \mathbb{Q}]$. Then compute the isomorphism $\phi : K \rightarrow K'$, which sends γ^i to $X^i + (h(X))$ for $i = 0, 1, \dots, d-1$. \square

As an application, we generalize Lemma 3.10 to obtain an efficient algorithm that computes embeddings of relative number fields over a given number field.

Lemma 4.7. *There exists a polynomial-time algorithm `ComputeRelEmbeddings` that given a number field K_0 and relative number fields K and K' over K_0 , computes all the embeddings of K in K' over K_0 .*

Proof. Identify K and K' with ordinary number fields using Corollary 4.1. Run the algorithm `ComputeEmbeddings` in Lemma 3.10 to compute all the embeddings of K in K' , and ignore those not fixing K_0 . \square

4.2 Adjoining roots of polynomials

One of the most basic techniques of constructing number fields is adjoining roots of polynomials. It can be efficiently performed by the following lemma.

Lemma 4.8. *There exists a polynomial-time algorithm `AdjoinRoot` that given a number field K_0 , a relative number field K over K_0 , and a polynomial $h(X) \in K[X]$ irreducible over K , computes the relative number field $K' = K(\alpha)$ over K_0 (up to isomorphism over K_0), where α is an arbitrary root of $h(X)$. Moreover, suppose K is encoded by the minimal polynomial of a primitive element $\beta \in K$ over K_0 . Then K' is encoded by the minimal polynomial of an element of the form $\beta + k\alpha$ over K_0 , where $1 \leq k \leq [K' : K_0] + 1$.*

Proof. Form the K -algebra $K'' := K[X]/(h(X))$ which is a field. We need to encode K'' as a relative number field over K_0 . Let $\alpha := X + (h(X)) \in K''$ which is a root of $h(X)$. Then α and β are explicitly known in K'' . Let $d := [K'' : K_0] + 1$. By Lemma 4.5, there exists $k \in [1, d+1]$ such that $\gamma = \beta + k\alpha$ is a primitive element of K'' over K_0 . Compute such an element γ by enumerating k and checking if the degree of the minimal polynomial of γ over K_0 equals d . Once γ is found, compute the relative number field $K' := K_0[X]/(g(X))$ over K_0 , where $g(X)$ is the minimal polynomial of γ over K_0 . It is isomorphic to $K'' = K(\alpha)$ over K_0 via the K_0 -linear map sending $X + (g(X))$ to γ . \square

By repeatedly adjoining roots, we obtain an algorithm that computes the splitting field of a given irreducible polynomial over a number field K_0 . See Algorithm 7.

Algorithm 7 SplittingField

Input: number field K_0 and $g(X) \in K_0[X]$ irreducible over K_0

Output: the splitting field of g over K_0 as a relative number field over K_0

- 1: $K \leftarrow K_0$, regarded as a relative number field over K_0
 - 2: factorize g over K
 - 3: **while** g has an irreducible non-linear factor over K **do**
 - 4: pick an irreducible non-linear factor g_0 of g over K
 - 5: run AdjoinRoot on (K_0, K, g_0) to obtain K'
 - 6: $K \leftarrow K'$
 - 7: factorize g over K
 - 8: **return** K
-

Line 2 and Line 7 are implemented using the polynomial-time factoring algorithms for number fields [Len83; Lan85].³ And we have

Lemma 4.9. *Given a number field K_0 and a polynomial $g(X) \in K_0[X]$ irreducible over K_0 , the algorithm SplittingField computes the splitting field K of g over K_0 in time polynomial in $[K : K_0]$ and the size of the input.*

Proof. The algorithm initializes K to K_0 and keeps adjoining roots of g to K until it contains all these roots. The resulting field K is by definition the splitting field of g over K_0 . At most $t := \log[K : K_0]$ intermediate fields are constructed other than K_0 . By induction and Lemma 4.8, each intermediate field is encoded by the minimal polynomial of a primitive element $k_1\alpha_1 + \cdots + k_s\alpha_s$ over K_0 where $s \leq t$, all α_i are roots of g and $1 \leq k_i \leq [K : K_0] + 1$. The claim about the running time then follows from Lemma 4.4 and Lemma 4.8. \square

Choosing $K_0 = \mathbb{Q}$ proves Lemma 3.20. Similarly, we have an algorithm constructing a (K_0, g) -subfield system whose associated subgroup system is a system of stabilizers. See Algorithm 8 below.

Again, Line 8 is implemented using the polynomial-time factoring algorithms for number fields [Len83; Lan85]. The condition at Line 11 is checked using the algorithm ComputeRelEmbeddings in Lemma 4.7.

³Here we factorize g over the relative number field K . It can be reduced to the problem of factoring polynomials over an ordinary number field by Corollary 4.1.

Algorithm 8 Stabilizers

Input: number field K_0 , $m \in \mathbb{N}$, and $g(X) \in K_0[X]$ irreducible over K_0

Output: (K_0, g) -subfield system \mathcal{F}

```

1: if  $m = 0$  then
2:   return  $\emptyset$ 
3:  $m \leftarrow \min(\deg(g), m)$ 
4:  $\mathcal{F} \leftarrow \{K_0[X]/(g(X))\}$ 
5: for  $i \leftarrow 2$  to  $m$  do
6:    $\mathcal{F}_{\text{old}} \leftarrow \mathcal{F}$ 
7:   for  $K \in \mathcal{F}_{\text{old}}$  do
8:     factorize  $g$  over  $K$ 
9:     for irreducible non-linear factor  $g_0$  of  $g$  over  $K$  do
10:      run AdjoinRoot on  $(K_0, K, g_0)$  to obtain  $K'$ 
11:      if  $K'$  is non-isomorphic to all fields in  $\mathcal{F}$  over  $K_0$  then
12:         $\mathcal{F} \leftarrow \mathcal{F} \cup \{K'\}$ 
13: return  $\mathcal{F}$ 

```

We have the following lemma.

Lemma 4.10. *Given a number field K_0 , an integer $m \in \mathbb{N}$, and a polynomial $g(X) \in K_0[X]$ irreducible over K_0 , the algorithm `Stabilizers` computes a (K_0, g) -subfield system \mathcal{F} , such that the subgroup system \mathcal{P} associated with \mathcal{F} is the system of stabilizers of depth m over $G(g/K_0)$ with respect to the action of $\text{Gal}(g/K_0)$ on the set of roots of g in L , where L denotes the splitting field of g over K_0 . Moreover, the algorithm runs in time polynomial in $c(\mathcal{P})$ and the size of the input.*

Proof. If $m = 0$, the algorithm simply returns $\mathcal{F} = \emptyset$. It replaces m with $\min(\deg(g), m)$ at Line 3, which does not change the desired subgroup system. So we may assume $m \leq \deg(g)$. The condition at Line 11 guarantees that the fields in \mathcal{F} are mutually non-isomorphic over K_0 . For $k \in [m]$, let \mathcal{P}_k be the the system of stabilizers of depth k over $G(g/K_0)$ with respect to the action of $\text{Gal}(g/K_0)$ on the set of roots of g in L , and let \mathcal{P}_k^\sharp be the corresponding poset of subfields of L determined by the Galois correspondence. Then \mathcal{P}_k^\sharp consists of the fields of the form $K_0(\alpha_1, \dots, \alpha_i)$, where $i \in [k]$ and $\alpha_1, \dots, \alpha_i$ are roots of g in L .

We want to show that at the end of the algorithm, the subgroup system \mathcal{P} associated with \mathcal{F} equals \mathcal{P}_m . And it suffices to prove that for $k \in [m]$, after the k th iteration

of the loop in Lines 5–12, every field in \mathcal{F} is isomorphic to some field in \mathcal{P}_k^\sharp over K_0 and vice versa. This follows from a simple induction on k .

Denote by d the maximum degree of the fields in \mathcal{F} over K_0 . Then d and $|\mathcal{F}|$ are bounded by $c(\mathcal{P})$. By induction and Lemma 4.8, each field in \mathcal{F} is encoded by the minimal polynomial of a primitive element $k_1\alpha_1 + \cdots + k_s\alpha_s$ over K_0 where $s \leq m \leq \deg(g)$, all α_i are roots of g , and $1 \leq k_i \leq d + 1$. The claim about the running time then follows from Lemma 4.4 and Lemma 4.8. \square

By Lemma 4.2, the complexity $c(\mathcal{P})$ of the subgroup system \mathcal{P} in Lemma 4.10 is bounded by $(\deg(g))^{m'}$, where $m' = \min\{\deg(g), m\}$. Lemma 3.21 then follows by choosing $K_0 = \mathbb{Q}$.

4.3 Reduction to primitive group actions

Suppose K_0 is a number field, $g(X) \in K_0[X]$ is irreducible over K_0 , and L is the splitting field of g over K_0 . The Galois group $\text{Gal}(g/K_0) = \text{Gal}(L/K_0)$ acts faithfully and transitively on the set of roots of g in L , and hence is a transitive permutation group on this set.

Motivated by Theorem 3.9, we are interested in the problem of constructing a (K_0, g) -subfield system \mathcal{F} such that all strongly antisymmetric \mathcal{P} -schemes are discrete (resp. inhomogeneous) on H , where \mathcal{P} is the subgroup system over $\text{Gal}(g/K_0)$ associated with \mathcal{F} , and H is a subgroup in \mathcal{P} satisfying $L^H \cong_{K_0} K_0[X]/(g(X))$. In this section, we describe a reduction, based on the work [LM85; Evd92], that reduces the problem to the special case that $\text{Gal}(g/K_0)$ is a *primitive* permutation group.

Definition 4.5 (primitive permutation group). *Suppose G is a permutation group on a finite set S . A nonempty subset B of S is called a set of imprimitivity⁴ of G if for all $g \in G$, either ${}^gB = B$ or $B \cap {}^gB = \emptyset$. A set of imprimitivity is trivial if it is a singleton or the whole set S . We say G is primitive if it only has trivial sets of imprimitivity. Otherwise G is imprimitive.*

It is well known that for transitive permutation groups, primitivity is equivalent to maximality of stabilizers.

⁴A set of imprimitivity is also called a *block* by some authors. We reserve the term *block* to denote a set in a partition instead.

Lemma 4.11. *Let S be a finite set where $|S| > 1$, and let $x \in S$. A transitive permutation group G on S is primitive iff G_x is maximal in G .*

See, e.g., [Wie64] for the proof of Lemma 4.11. We also need the following result, proved in [LM85].

Theorem 4.1 ([LM85]). *There exists a polynomial-time algorithm `Tower` that given a number field K_0 and a polynomial $g(X) \in K_0[X]$ irreducible over K_0 ,⁵ computes a tower of relative number fields over K_0*

$$K_0 \subseteq K_1 \subseteq \cdots \subseteq K_{k-1} \subseteq K_k$$

together with the inclusions $K_{i-1} \hookrightarrow K_i$ and the polynomials $g_i(X) \in K_{i-1}[X]$ irreducible over K_{i-1} for $i \in [k]$, such that $K_k \cong_{K_0} K_0[X]/(g(X))$, and the following conditions are satisfied for $i \in [k]$:

1. K_i is isomorphic to $K_{i-1}[X]/(g_i(X))$ over K_{i-1} , and
2. the Galois group $G_i := \text{Gal}(L_i/K_{i-1})$ acts primitively on the set of roots of g_i in L_i , where L_i is the Galois closure of K_i/K_{i-1} .

For $i \in [k]$, let $H_i := \text{Gal}(L_i/K_i) \subseteq G_i$. See Figure 4.1 for an illustration. Note that the first condition above is equivalent to $K_i = K_{i-1}(\alpha)$ for some root α_i of g_i in L_i . So H_i is the stabilizer of α_i . Then the second condition is equivalent to maximality of H_i in G_i .

The following theorem is the main result of this section.

Theorem 4.2. *Suppose there exists an algorithm `PrimitiveAction` that, given a number field K_0 and a polynomial $g(X) \in K_0[X]$ irreducible over K_0 with $\text{Gal}(g/K_0)$ acting primitively on the set of roots of g in L , where L is the splitting field of g over K_0 , computes a (K_0, g) -subfield system in time $T(K_0, g)$. Then there exists an algorithm `GeneralAction` that given K_0 and g as above, but without the assumption that $\text{Gal}(g/K_0)$ acts primitively on S , computes*

- a (K_0, g) -subfield system \mathcal{F} , and,

⁵The paper [LM85] presented their algorithm only for $K_0 = \mathbb{Q}$, but it easily extends to a general base field K_0 .




Figure 4.1: The tower of fields and Galois groups in Theorem 4.1

- a tower of relative number fields $K_0 \subseteq K_1 \subseteq \cdots \subseteq K_{k-1} \subseteq K_k$ over K_0 and $g_i(X) \in K_{i-1}[X]$ for $i \in [k]$ satisfying the conditions in Theorem 4.1, such that $K_k \cong_{K_0} K_0[X]/(g(X))$ and the sizes of the polynomials g_i are polynomial in the size of the input

in time polynomial in $\sum_{i=1}^k T(K_{i-1}, g_i)$ and the size of the input. Moreover, if for each $i \in [k]$, the (K_{i-1}, g_i) -subfield system \mathcal{F}_i computed by `PrimitiveAction` on the input (K_{i-1}, g_i) satisfies

1. $K_{i-1}[X]/(g_i(X)) \in \mathcal{F}_i$,
2. All strongly antisymmetric \mathcal{P} -schemes are discrete (resp. inhomogeneous) on H , where \mathcal{P} is the subgroup system over $\text{Gal}(g_i/K_{i-1})$ associated with \mathcal{F}_i and H is a subgroup in \mathcal{P} whose fixed field is isomorphic to K_i over K_{i-1} .

Then \mathcal{F} satisfies

1. $K_0[X]/(g(X)) \in \mathcal{F}$,
2. All strongly antisymmetric \mathcal{P} -schemes are discrete (resp. inhomogeneous) on H , where \mathcal{P} is the subgroup system over $\text{Gal}(g/K_0)$ associated with \mathcal{F} and H is a subgroup in \mathcal{P} satisfying $L^H \cong_{K_0} K_0[X]/(g(X))$.

See Algorithm 9 for the pseudocode of the algorithm `GeneralAction`. It proceeds as follows: maintain \mathcal{F} , which initially only contains $K_0[X]/(g(X))$. Then we

call the algorithm `Tower` to compute a tower $K_0 \subseteq K_1 \subseteq \cdots \subseteq K_{k-1} \subseteq K_k$ and $g_i(X) \in K_{i-1}[X]$ for $i \in [k]$ as in Theorem 4.1. Next, run the hypothetical algorithm `PrimitiveAction` in Theorem 4.2 on (K_{i-1}, g_i) for each $i \in [k]$ to obtain a (K_{i-1}, g_i) -subfield system \mathcal{F}_i . For $i \in [k]$, add the fields in \mathcal{F}_i to \mathcal{F} , but encode them as relative number fields over K_0 (using Lemma 4.6). In addition, avoid adding fields to \mathcal{F} that are isomorphic to some existent field $K \in \mathcal{F}$ over K_0 , so that all the fields in \mathcal{F} are mutually non-isomorphic over K_0 . After all \mathcal{F}_i are processed, output \mathcal{F} .

Algorithm 9 `GeneralAction`

Input: number field K_0 and $g(X) \in K_0[X]$ irreducible over K_0

Output: (K_0, g) -subfield system \mathcal{F}

- 1: $\mathcal{F} \leftarrow \{K_0[X]/(g(X))\}$
 - 2: run `Tower` on (K_0, g) to obtain a tower $K_0 \subseteq K_1 \subseteq \cdots \subseteq K_{k-1} \subseteq K_k$ and $g_i(X) \in K_{i-1}[X]$ irreducible over K_{i-1} for $i \in [k]$
 - 3: **for** $i \leftarrow 1$ **to** k **do**
 - 4: run `PrimitiveAction` on (K_{i-1}, g_i) to obtain \mathcal{F}_i
 - 5: **for** $K \in \mathcal{F}_i$ **do**
 - 6: compute a relative number field K' over K_0 such that $K' \cong_{K_0} K$
 - 7: **if** K' is non-isomorphic to all fields in \mathcal{F} over K_0 **then**
 - 8: $\mathcal{F} \leftarrow \mathcal{F} \cup \{K'\}$
 - 9: **return** \mathcal{F}
-

The proof of Theorem 4.2 is based on the following lemma.

Lemma 4.12. *Let $k \in \mathbb{N}^+$ and $G_k \subseteq G_{k-1} \subseteq \cdots \subseteq G_1 \subseteq G_0$ be a chain of finite groups. For $i \in [k]$, let N_i be a subgroup of G_i that is normal in G_{i-1} , $\pi_i : G_{i-1} \rightarrow G_{i-1}/N_i$ be the corresponding quotient map, and \mathcal{P}_i be a subgroup system over G_{i-1}/N_i that contains G_i/N_i . Define*

$$\mathcal{P} = \{g\pi_i^{-1}(H)g^{-1} : 1 \leq i \leq k, H \in \mathcal{P}_i, g \in G_0\},$$

which is a subgroup system over G_0 and contains $\pi_i^{-1}(G_i/N_i) = G_i$ for all $i \in [k]$. Then we have

1. *If for all $i \in [k]$, all strongly antisymmetric \mathcal{P}_i -schemes are discrete on G_i/N_i , then all strongly antisymmetric \mathcal{P} -schemes are discrete on G_k .*

2. If for some $i \in [k]$, all strongly antisymmetric \mathcal{P}_i -schemes are inhomogeneous on G_i/N_i , then all strongly antisymmetric \mathcal{P} -schemes are inhomogeneous on G_k .

The same holds if strong antisymmetry is replaced by antisymmetry.

We defer the proof of Lemma 4.12 to Section 6.1.

Proof of Theorem 4.2. The claims about K_i and g_i follow from Theorem 4.1. Use the following notations for $i \in [k]$:

- L_i : the splitting field of g_i over K_{i-1} , which is a subfield of L .
- $G_i := \text{Gal}(L_i/K_{i-1})$ and $N_i := \text{Gal}(L/L_i)$.
- π_i : the natural projection $\text{Gal}(L/K_{i-1}) \rightarrow \text{Gal}(L/K_{i-1})/N_i \cong G_i$.
- \mathcal{P}_i : the subgroup system over G_i associated with \mathcal{F}_i .

Then by construction, the subgroup system over $\text{Gal}(L/K_0)$ associated with \mathcal{F} is

$$\mathcal{P} := \{g\pi_i^{-1}(H)g^{-1} : 1 \leq i \leq k, H \in \mathcal{P}_i, g \in G\}.$$

Assume the conditions on \mathcal{F}_i in Theorem 4.2 are satisfied. Then for all $i \in [k]$, all strongly antisymmetric \mathcal{P}_i -schemes are discrete (resp. inhomogeneous) on $\text{Gal}(L_i/K_i) \in \mathcal{P}_i$. Applying Lemma 4.12 to the chain

$$\text{Gal}(L/K_k) \subseteq \text{Gal}(L/K_{k-1}) \subseteq \cdots \subseteq \text{Gal}(L/K_1) \subseteq \text{Gal}(L/K_0)$$

and $N_i, \pi_i, \mathcal{P}_i$, we conclude that all strongly antisymmetric \mathcal{P} -schemes are discrete (resp. inhomogeneous) on the subgroup $\text{Gal}(L/K_k) \in \mathcal{P}$. And the corresponding fixed field K_k is isomorphic to $K_0[X]/(g(X))$ over K_0 , as desired.

The total running time of the algorithm `PrimitiveAction` and the total size of \mathcal{F}_i are both bounded by $\sum_{i=1}^k T(K_{i-1}, g_i)$. The other operations take time polynomial in the total size of \mathcal{F}_i and the size of the input. The claim about the running time follows. \square

As an application, we prove the main result of [Evd92] for the special case that the input polynomial satisfies Condition 3.1 (i.e., it is defined over \mathbb{F}_p , square free, and complete reducible over \mathbb{F}_p).

Theorem 4.3 ([Evd92]). *Under GRH, there exists a deterministic polynomial-time algorithm that, given a polynomial $f(X) \in \mathbb{F}_p[X]$ satisfying Condition 3.1 and a lifted polynomial $\tilde{f}(X) \in \mathbb{Z}[X]$ of f whose Galois group $\text{Gal}(\tilde{f}/\mathbb{Q})$ is solvable, computes the complete factorization of f over \mathbb{F}_p .*

The proof relies on the following bound for the orders of primitive solvable permutation groups, proved by Pálffy [Pál82].

Theorem 4.4 ([Pál82]). *Let G be a primitive solvable permutation group on a set of cardinality $n \in \mathbb{N}^+$. Then $|G| \leq 24^{-1/3}n^c$ for a constant $c = 3.24399\dots$*

Proof of Theorem 4.3. As in Section 6, we factorize \tilde{f} into its irreducible factors $f_1(X), \dots, f_k(X) \in \mathbb{Z}[X]$ over \mathbb{Q} in polynomial time using the factoring algorithm in [LLL82]. The Galois groups $\text{Gal}(\tilde{f}_i(X)/\mathbb{Q})$ are quotient groups of $\text{Gal}(\tilde{f}/\mathbb{Q})$, and hence are solvable as well. By replacing $\tilde{f}(X)$ with $\tilde{f}_i(X)$ and $f(X)$ with $f_i(X) := \tilde{f}_i(X) \bmod p \in \mathbb{F}_p[X]$ for each $i \in [k]$, we reduce to the case that \tilde{f} is irreducible over \mathbb{Q} .

Let L be the splitting field of \tilde{f} over \mathbb{Q} . When $\text{Gal}(\tilde{f}/\mathbb{Q})$ acts primitively on the set of roots of \tilde{f} in L , its order is bounded by a polynomial in $\deg(f)$ by Theorem 4.4. Then by Theorem 4.9, we can construct \mathcal{F} in polynomial time such that $\mathbb{Q}[X]/(\tilde{f}(X)) \in \mathcal{F}$ and all strongly antisymmetric \mathcal{P} -schemes are discrete on H , where \mathcal{P} is the subgroup system over $\text{Gal}(\tilde{f}/\mathbb{Q})$ associated with \mathcal{F} and H is a subgroup in \mathcal{P} satisfying $L^H \cong \mathbb{Q}[X]/(\tilde{f}(X))$. By Theorem 4.2, we also have a polynomial-time algorithm of constructing such \mathcal{F} in the general case. The theorem then follows from Theorem 3.9. \square

In Chapter 5, we prove a generalization of Theorem 4.3 (see Theorem 5.13), which implies the main result of [Evd92] in its general form. In particular, the assumption that \tilde{f} satisfies Condition 3.1 is no longer required.

4.4 Other techniques of constructing number fields

In this section, we survey some other techniques of constructing number fields. While we do not use these techniques directly in the thesis, they are worth mentioning because of their own interest and their applications to other problems [Lan84; LM85; Len92; Coh93].

Taking the compositum of number fields. Note that the fields computed in the two algorithms `SplittingField` and `Stabilizers` in Section 4.2 are (up to isomorphism over K_0) compositums of conjugates of the field $K_0[X]/(g(X))$. The general problem of constructing the compositum of (relative) number fields is solved by the following lemma.

Lemma 4.13. *There exists a polynomial-time algorithm that given a number field K_0 and relative number fields K, L over K_0 , constructs all the compositums $K'L'$ up to isomorphism over K_0 where K' (resp. L') ranges over the conjugates of K (resp. L) over K_0 in the algebraic closure \bar{K}_0 of K_0 .⁶*

Proof. Take the irreducible polynomial $g(X) \in K_0[X]$ that encodes L , i.e., $L \cong_{K_0} K_0[X]/(g(X))$. Factorize $g(X)$ into irreducible polynomials $g_1(X), \dots, g_k(X)$ over K . Then compute and output the fields $K[X]/(g_1(X)), \dots, K[X]/(g_k(X))$.

To see that this gives the desired output, note that we may fix $K = K'$ as fields are constructed only up to isomorphism over K_0 . Let $\alpha_1, \dots, \alpha_n$ be the roots of g in \bar{K}_0 , where $n = \deg(g)$. Then the conjugates of L in \bar{K}_0 over K_0 are precisely $K_0(\alpha_1), \dots, K_0(\alpha_n)$. For $i \in [n]$, there exists a unique $j_i \in [k]$ such that α_i is the root of g_{j_i} , and the compositum of K and $K_0(\alpha_i)$ is just $K(\alpha_i) \cong_{K_0} K[X]/(g_{j_i}(X))$. \square

Taking the intersection of number fields. The intersection of two number fields can be computed efficiently, as shown in [LM85].

Theorem 4.5 ([LM85]). *There exists a polynomial-time algorithm that given*

- *number fields $K = \mathbb{Q}(\alpha)$, $K' = \mathbb{Q}(\beta)$ encoded by the minimal polynomials of primitive elements $\alpha \in K$ and $\beta \in K'$ over \mathbb{Q} respectively, and*
- *the minimal polynomial $h_0(X) \in K[X]$ of β over K ,⁷*

computes the number field $K \cap K'$ up to isomorphism.

The algorithm in [LM85] also extends to relative number fields. We omit the details.

⁶Here K and L are embedded in \bar{K}_0 via some K_0 -linear embeddings. The choices of these embeddings do not matter as we construct $K'L'$ for all the conjugates K' and L' over K_0 .

⁷The polynomial h_0 is needed for the problem to be well defined.

Adjoining a square root of the discriminant. Suppose K is a relative number field over K_0 encoded by the minimal polynomial $h(X) \in K_0[X]$ of a primitive element $\alpha \in K$ over K_0 . Let L be the Galois closure of K/K_0 and let $G = \text{Gal}(L/K_0)$. Then G acts on the set S of roots of h in L and hence can be identified with a subgroup of $\text{Sym}(S)$.

Suppose $S = \{\alpha_1, \dots, \alpha_n\}$. Define the *discriminant* of h to be

$$\Delta_h := \prod_{1 \leq i < j \leq n} (\alpha_i - \alpha_j)^2.$$

We have ${}^g\Delta_h = \Delta_h$ for all $g \in G$. So $\Delta_h \in L^G = K_0$.

Now consider the subfield $K'_0 := K_0(\sqrt{\Delta_h})$ of L , where $\sqrt{\Delta_h} := \prod_{1 \leq i < j \leq n} (\alpha_i - \alpha_j)$ is a square root of Δ_h in L . A permutation $g \in G$ fixes $\sqrt{\Delta_h}$ precisely when g is an even permutation of S , which implies

$$\text{Gal}(L/K'_0) = G \cap \text{Alt}(S).$$

With this observation, we have

Lemma 4.14. *There exists a polynomial-time algorithm that given a number field K_0 and a relative number field K over K_0 encoded by $h(X) \in K_0[X]$, computes $L^{G \cap \text{Alt}(S)}$ up to isomorphism over K_0 , where L is the Galois closure of K/K_0 , $G = \text{Gal}(L/K_0)$, and S is the set of roots of h in L .*

Proof. We have $L^{G \cap \text{Alt}(S)} = K_0(\sqrt{\Delta_h})$ by the above discussion. Let $n = \deg(h)$. Then discriminant Δ_h satisfies the identity

$$\Delta_h = (-1)^{n(n-1)/2} \text{Res}(h, h'),$$

where $\text{Res}(h, h')$ denotes the *resultant* of h and its derivative h' . and is given by the determinant of the *Sylvester matrix* associated with h and h' [Lan02]. Thus we can compute Δ_h in polynomial time. Then we test if Δ_h is a square in K_0 by factoring $X^2 - \Delta_h$ over K_0 . If Δ_h is a square, we have $K_0(\sqrt{\Delta_h}) = K_0$ and correspondingly $G \subseteq \text{Alt}(S)$. In this case we just output K_0 . Otherwise we output $K_0[X]/(X^2 - \Delta_h)$. \square

Remark. The technique above was used in [Lan84] for the determination of the Galois groups of number field extensions. It is not clear, however, if it helps for the problem of polynomial factoring over finite fields. We note that replacing K_0

with $K'_0 = K_0(\sqrt{\Delta_h})$ and K with $K'_0 K$ has the effect of reducing the Galois group G to $G \cap \text{Alt}(S)$, but the order of G is reduced by at most a factor of two. This does not help in the case that $G = \text{Sym}(S)$ and \mathcal{P} is a system of stabilizers of depth $m \leq |S| - 2$ (with respect to the natural action of G): As both $\text{Sym}(S)$ and $\text{Alt}(S)$ are k -transitive for $k = |S| - 2$, \mathcal{P} -schemes for $\text{Sym}(S)$ and those for $\text{Alt}(S)$ both correspond to m -schemes on S (see Theorem 2.1), and hence they are the equivalent.

Computing the fixed field of the automorphism group. The following lemma gives a characterization of the fixed field of an automorphism subgroup.

Lemma 4.15. *Suppose K/K_0 is a field extension and α is a primitive element of K over K_0 . For a subgroup $U \subseteq \text{Aut}(K/K_0)$, the field K^U is generated by elementary symmetric polynomials in the elements ${}^g\alpha$ (indexed by $g \in U$) over K_0 .*

Proof. Let K' be the subfield of K generated by elementary symmetric polynomials in ${}^g\alpha$, $g \in U$ over K_0 . We obviously have $K' \subseteq K^U$. By Galois theory, it holds that $[K : K^U] = |U|$ (see, e.g., [Lan02, Section VI.1, Theorem 1.8]). So it suffices to prove $[K : K'] \leq |U|$.

Consider the polynomial $\phi(X) = \prod_{g \in U} (X - {}^g\alpha)$. The coefficients of $\phi(X)$ are, up to sign, given by elementary symmetric polynomials in ${}^g\alpha$, $g \in U$ and hence $\phi(X) \in K'[X]$. As $\phi(\alpha) = 0$, the minimal polynomial of α over K' divides $\phi(X)$, and its degree is at most $\deg(\phi) = |U|$. So we have $[K'(\alpha) : K'] \leq |U|$. The claim follows by noting that $K'(\alpha) = K$. \square

Lemma 4.15 provides a method of computing the fixed field of the automorphism group $\text{Aut}(K/K_0)$:

Theorem 4.6. *There exists a polynomial-time algorithm that given a number field K_0 and a relative number field K over K_0 , computes the fixed field $K^{\text{Aut}(K/K_0)} \subseteq K$.*

Proof. Suppose K is encoded by the minimal polynomial of a primitive element α over K_0 . We compute all the automorphisms of K in $\text{Aut}(K/K_0)$ using Lemma 4.7. Then we adjoining to K_0 the first k elementary symmetric functions in ${}^g\alpha$, $g \in \text{Aut}(K/K_0)$ where $k = |\text{Aut}(K/K_0)|$. The resulting field is exactly $K^{\text{Aut}(K/K_0)}$ by Lemma 4.15. \square

More generally, given K_0, K and a subgroup $U \subseteq \text{Aut}(K/K_0)$ of automorphisms of K , the same proof shows that K^U can be constructed in polynomial time.

Now suppose L is a Galois extension of K_0 that contains K . Let $G = \text{Gal}(L/K_0)$ and $H = \text{Gal}(L/K)$. Then $\text{Aut}(K/K_0)$ is identified with $N_G(H)/H$, and we have $K^{\text{Aut}(K/K_0)} = K^{N_G(H)/H} = L^{N_G(H)}$. So Theorem 4.6 states that $L^{N_G(H)}$ can be constructed in polynomial time given $K = L^H$ and K_0 . In the context of polynomial factoring using the \mathcal{P} -scheme algorithm, this means that we can efficiently enlarge a subgroup system \mathcal{P} by including the normalizers $N_G(H)$ of $H \in \mathcal{P}$.

A natural question arising from this observation is whether adding the normalizers (or more generally subgroups between $N_G(H)$ and H) to the subgroup system helps a \mathcal{P} -scheme algorithm obtain the complete factorization (resp. a proper factorization). By Theorem 3.9, this reduces to the question whether it helps for proving all strongly antisymmetric \mathcal{P} -schemes are discrete (resp. inhomogeneous) on a distinguished subgroup $H \in \mathcal{P}$.

For discreteness of strongly antisymmetric \mathcal{P} -schemes, we give an affirmative answer in general: we show that for some subgroup system \mathcal{P} and $H \in \mathcal{P}$, there exist strongly antisymmetric \mathcal{P} -schemes that are not discrete on H , but adding normalizers to the subgroup system rules out their existence.

Example 4.1. Choose a finite group G and a subgroup $H \subseteq G$ such that $N_G(H)$ is a proper normal subgroup of G .⁸ Choose $\mathcal{P} = \{gHg^{-1} : g \in G\}$ which is a subgroup system over G . Define a \mathcal{P} -collection $\mathcal{C} = \{C_{H'} : H' \in \mathcal{P}\}$ as follows: the group $N_G(H)$ acts on $H \backslash G$ by left translation ${}^gHh = Hgh$ and $H \backslash G$ is partitioned into $N_G(H)$ -orbits. Choose a complete set of representatives $B \subseteq H \backslash G$ for these orbits. Define $C_H = \{{}^gB : g \in N_G(H)\}$. For any other subgroup H' in \mathcal{P} , choose $g \in G$ such that $H' = gHg^{-1}$, and define $C_{H'} = \{c_{H,g}(B) : B \in C_H\}$. It is easy to see that \mathcal{C} is a well defined strongly antisymmetric \mathcal{P} -scheme. Moreover, it is not discrete on H since $N_G(H)$ does not act transitively on $H \backslash G$.

Now define $\mathcal{P}' = \mathcal{P} \cup \{N_G(H)\}$ which is also a subgroup system over G . We claim that any antisymmetric \mathcal{P}' -schemes \mathcal{C}' must be discrete on any subgroup in \mathcal{P}' . To see this, note that \mathcal{C}' is discrete on $N_G(H) \in \mathcal{P}'$ since $N_G(H)$ is normal in G . Then \mathcal{C}' is also discrete on all the other subgroups $H' \in \mathcal{P}'$ by compatibility, and the claim follows. In particular, it is impossible to extend \mathcal{C} to an antisymmetric \mathcal{P}' -scheme.

⁸For example, we may choose G to be the semidirect product $(K \times K) \rtimes C_2$, where K is a nontrivial finite group and C_2 permutes the two direct factors of $K \times K$. Let $H = K \times \{e\}$. Then $N_G(H) = K \times K \trianglelefteq G$.

Despite the example above, adding normalizers to the subgroup system seems not helpful for attacking the most difficult cases in polynomial factoring: for a subgroup system \mathcal{P} over a finite group G , define

$$\mathcal{P}_+ = \{U : H \subseteq U \subseteq N_G(H), H \in \mathcal{P}\},$$

which is also a subgroup system over G . For several important families of permutation groups, we show that if \mathcal{P} is the corresponding system of stabilizers of certain depth m (where m is not too large), any \mathcal{P} -scheme \mathcal{C} can be extended to a \mathcal{P}_+ -scheme \mathcal{C}' with antisymmetry and strong antisymmetry preserved. In particular, if \mathcal{C} is not discrete or inhomogeneous on some subgroup $H \in \mathcal{P}$, then neither is \mathcal{C}' .

Lemma 4.16. *Let S be a finite set and let G be $\text{Sym}(S)$ or $\text{Alt}(S)$ acting naturally on S . Let \mathcal{P} be the system of stabilizers of depth m over G with respect to this action where $m < |S|/2$. Then any \mathcal{P} -scheme \mathcal{C} can be extended to a \mathcal{P}_+ -scheme \mathcal{C}' such that \mathcal{C}' is antisymmetric (resp. strongly antisymmetric) if so is \mathcal{C} .*

Lemma 4.17. *Let V be a finite dimensional vector space over a finite field and let G be $\text{GL}(V)$ acting naturally on $S := V - \{0\}$. Let \mathcal{P} be the system of stabilizers of depth m over G with respect to this action where $m < \dim_F V$. Then any \mathcal{P} -scheme \mathcal{C} can be extended to a \mathcal{P}_+ -scheme \mathcal{C}' such that \mathcal{C}' is antisymmetric (resp. strongly antisymmetric) if so is \mathcal{C} .*

We defer the proofs of Lemma 4.16 and Lemma 4.17 to Section 6.4 . There we define the *closure* \mathcal{P}_{cl} of a subgroup system \mathcal{P} , and then show that \mathcal{P} -schemes can always be extended to \mathcal{P}_{cl} -schemes with antisymmetry and strong antisymmetry preserved. Lemma 4.16 and Lemma 4.17 then follow immediately once we verify that $\mathcal{P}_{\text{cl}} = \mathcal{P}_+$ in these cases.

THE GENERALIZED \mathcal{P} -SCHEME ALGORITHM

In Chapter 3, we developed the \mathcal{P} -scheme algorithm that factorizes polynomials satisfying Condition 3.1, i.e., they are defined over a prime field \mathbb{F}_p , square-free, and completely reducible over \mathbb{F}_p . In this chapter, we extend this algorithm to factorize general polynomials $f(X) \in \mathbb{F}_q[X]$ over a finite field \mathbb{F}_q of characteristic p . The generality is reflected in the following three aspects: (1) \mathbb{F}_q may be a non-prime field, (2) the degrees of the irreducible factors of f may be greater than one, and (3) the multiplicities of the irreducible factors of f may be greater than one.

Motivation. Techniques like Berlekamp's reduction [Ber70], square-free factorization [Yun76; Knu98] and distinct-degree factorization [CZ81] were commonly used in literature to reduce the problem to the special case that the input polynomial satisfies Condition 3.1. However, these reductions do not preserve the information of the lifted polynomial \tilde{f} employed by the \mathcal{P} -scheme algorithm. Therefore, it is desirable to avoid these reductions and extend the \mathcal{P} -scheme algorithm to the general setting instead.

As a concrete example, consider the following polynomial $\tilde{f}(X) \in \mathbb{Z}[X]$ irreducible over \mathbb{Q} , taken from [KM00]:

$$\begin{aligned} \tilde{f}(X) = & X^{14} + 28X^{11} + 28X^{10} - 28X^9 + 140X^8 + 360X^7 + 147X^6 \\ & + 196X^5 + 336X^4 - 546X^3 - 532X^2 + 896X + 823. \end{aligned}$$

For $p = 43$, the reduced polynomial $f(X) = \tilde{f}(X) \bmod p$ has seven distinct linear factors and one irreducible factor of degree 7 over \mathbb{F}_p :

$$\begin{aligned} f(X) = & (X + 2)(X + 4)(X + 9)(X + 19)(X + 23)(X + 30)(X + 42) \\ & (X^7 + 14X^4 + 15X^3 + 31X^2 + 15X + 38). \end{aligned}$$

The standard way of factoring f over \mathbb{F}_p is first applying distinct-degree factorization [CZ81] to obtain a partial factorization $f = f_0 f_1$, where

$$f_0(X) = (X + 2)(X + 4)(X + 9)(X + 19)(X + 23)(X + 30)(X + 42)$$

is the product of the linear factors and satisfies Condition 3.1. Then we factorize f_0 over \mathbb{F}_p . To achieve this goal deterministically, we pick a lifted polynomial

$\tilde{f}_0(X) \in \mathbb{Z}[X]$ of f , which we may assume to be irreducible, and run the \mathcal{P} -scheme algorithm in Chapter 3. Suppose the (\mathbb{Q}, \tilde{f}) -subfield system in the algorithm is constructed by Lemma 3.21 and the associated subgroup system \mathcal{P} is the system of stabilizers of depth m , where $m \in \mathbb{N}^+$ is sufficiently large. In the worst case, the action of $\text{Gal}(\tilde{f}_0/\mathbb{Q})$ on the set of roots of \tilde{f} is permutation isomorphic to the natural action of the symmetric group $\text{Sym}(7)$ on $[7]$. Then we need $m \geq 3$ to obtain a proper factorization of f , since by Theorem 2.1 and Lemma 2.19, there exists a strongly antisymmetric \mathcal{P} -scheme homogeneous on a stabilizer if $m \leq 2$.¹

On the other hand, the action of the Galois group of \tilde{f} on the set of roots of \tilde{f} is permutation isomorphic to the action of the wreath product² $C_7 \wr C_2$ on $[7] \times [2]$, where C_7 permutes $[7]$ cyclically and C_2 permutes the two copies of $[7]$. This action has a base of size two, which suggests that choosing $m = 2$ is sufficient for completely factoring f , provided that we have a generalization of Theorem 3.2 that employs the polynomial \tilde{f} . The goal of this chapter is to establish such a generalization.

The example above generalizes to an infinite family of instances: for every $k \in \mathbb{N}^+$, there exists $\tilde{f}(X) \in \mathbb{Z}[X]$ irreducible over \mathbb{Q} of degree $2k$ such that the action of the Galois group on the set of roots of \tilde{f} is permutation isomorphic to the action of $C_k \wr C_2$ on $[k] \times [2]$.³ And for such \tilde{f} , there exists infinitely many prime numbers p such that $f(X) = \tilde{f}(X) \pmod{p}$ has k distinct linear factors and one irreducible factor of degree k .⁴ Using the generalized \mathcal{P} -scheme algorithm developed in this chapter, it is sufficient to choose $m = 2$ in order to completely factorize $\tilde{f} \pmod{p}$, leading to a polynomial-time factoring algorithm for such instances (f, \tilde{f}) . On the other hand, using distinct-degree factorization and the \mathcal{P} -scheme algorithm in Chapter 3, the best known general upper bound for m is $O(\log k)$ (see Theorem 3.12), and the resulting algorithm takes superpolynomial time.

Lifted polynomial. To formulate the main result of this chapter, we first need to generalize the notion of *lifted polynomials* (see Definition 1.1). Recall that a lifted polynomial of $f(X) \in \mathbb{F}_p[X]$ is a polynomial $\tilde{f}(X) \in \mathbb{Z}[X]$ of degree

¹For the same reason, one needs to choose $m \geq 3$ if the m -scheme algorithm [IKS09] is used.

²For the definition of the wreath product of groups, see Definition 6.11.

³Shafarevich's theorem on solvable Galois groups [Sha54; ILF97] implies that the existence of integral polynomials realizing the family of groups $C_k \wr C_2$ as Galois groups. For an algorithm explicitly computing such a polynomial, see [KM00].

⁴This follows from Chebotarëv's density theorem. See, e.g., [Neu99].

$\deg(f)$ satisfying $\tilde{f}(X) \bmod p = f(X)$. For the general case $\mathbb{F}_q = \mathbb{F}_{p^d}$, we fix the following notations: assume \mathbb{F}_q is encoded by a monic irreducible polynomial $h(Y) \in \mathbb{F}_p[Y]$ of degree d , i.e., it is identified with $\mathbb{F}_p[Y]/(h(Y))$ via an isomorphism $\psi_0 : \mathbb{F}_p[Y]/(h(Y)) \rightarrow \mathbb{F}_q$ which we can efficiently compute. Lift h to a monic polynomial $\tilde{h}(Y) \in \mathbb{Z}[Y]$ of degree d which is necessarily irreducible over \mathbb{Q} . Define $A_0 := \mathbb{Z}[Y]/(\tilde{h}(Y))$ and $K_0 := \mathbb{Q}[Y]/(\tilde{h}(Y))$. Composing ψ_0 with the natural projection $A_0 \rightarrow \mathbb{F}_p[Y]/(h(Y))$ sending x to $x \bmod p$, we obtain a surjective ring homomorphism $\tilde{\psi}_0 : A_0 \rightarrow \mathbb{F}_q$. Finally extend $\tilde{\psi}_0$ to the ring $A_0[X]$ by applying it to each coefficient:

$$\tilde{\psi}_0 : A_0[X] \rightarrow \mathbb{F}_q[X].$$

With these notations, we generalize the definition of lifted polynomials as follows.

Definition 5.1 (lifted polynomial). *Suppose $f(X) \in \mathbb{F}_q[X]$ is a polynomial of degree $n \in \mathbb{N}^+$. A lifted polynomial of f (with respect to \tilde{h} and ψ_0) is a polynomial $\tilde{f}(X) \in A_0[X]$ of degree n satisfying $\tilde{\psi}_0(\tilde{f}) = f$. An irreducible lifted polynomial of f is a lifted polynomial of f that is irreducible over K_0 .*

Given $f(X) \in \mathbb{F}_q[X]$, we can choose a lifted polynomial \tilde{f} of f efficiently. Furthermore, we argue that \tilde{f} can be assumed to be irreducible over K_0 . To see this, we need the following lemma.

Lemma 5.1. *There exists a polynomial-time algorithm that given p and a polynomial $\tilde{f}(X) \in A_0[X]$ satisfying $\tilde{\psi}_0(\tilde{f}) \neq 0$, computes an integer D satisfying $D \equiv 1 \pmod{p}$ and a factorization of $D \cdot \tilde{f}$ into irreducible factors \tilde{f}_i over K_0 . Furthermore all of the factors $\tilde{f}_i(X)$ are in $A_0[X]$.*

The proof can be found in Appendix C. Compute D and f_i using the lemma above. We have $\tilde{\psi}_0(D \cdot \tilde{f}) = \tilde{\psi}_0(\tilde{f}) = f$ since $D \equiv 1 \pmod{p}$. So the polynomials $\tilde{\psi}_0(\tilde{f}_i)$ are factors of f , and we have reduced the problem to factoring each $\tilde{\psi}_0(\tilde{f}_i) \in \mathbb{F}_q[X]$ using its irreducible lifted polynomial \tilde{f}_i .

The discussion above justifies the assumption that an irreducible lifted polynomial \tilde{f} of f is given, with respect to \tilde{h} and ψ_0 . The notations \tilde{h} , ψ_0 , A_0 , and K_0 are fixed throughout this chapter.

Main result. The main result of this chapter is a generalization of Theorem 3.2:

Theorem 5.1 (informal). *Suppose there exists a deterministic algorithm that given a polynomial $g(X) \in A_0[X]$ irreducible over K_0 , constructs in time $T(g)$ a collection \mathcal{F} of subfields of the splitting field L of g over K_0 such that*

- $F = K_0[X]/(g(X))$ is in \mathcal{F} , and
- all strongly antisymmetric \mathcal{P} -schemes are discrete on $\text{Gal}(L/F) \in \mathcal{P}$, where \mathcal{P} is the subgroup system associated with \mathcal{F} .

Then under GRH, there exists a deterministic algorithm that given $f(X) \in \mathbb{F}_q[X]$ and an irreducible lifted polynomial $\tilde{f}(X) \in A_0[X]$ of f , outputs the complete factorization of f over \mathbb{F}_q in time polynomial in $T(\tilde{f})$ and the size of the input.

See Theorem 5.9 for the formal statement. For simplicity, here we only state the result for computing the complete factorization of f . The results for computing a proper factorization are slightly more complicated to state, and we refer the reader to Section 5.10 for details.

Overview of the generalized \mathcal{P} -scheme algorithm

Recall that a \mathcal{P} -scheme algorithm in Chapter 3 consists of three parts: (1) a reduction to the problem of computing an idempotent decomposition of the ring $\bar{\mathcal{O}}_F$, where $F = \mathbb{Q}[X]/(\tilde{f}(X))$, (2) computing idempotent decompositions for a collection of number fields, and (3) constructing the collection of number fields used in the previous part. The factoring algorithm in this chapter has the same structure but with some differences: we generalize the reduction in Part (1), where F now denotes the number field $K_0[X]/(\tilde{f}(X))$. And in Part (3), we construct a collection of relative number fields over K_0 instead of ordinary number fields. The main difference is in Part (2), which we now explain.

\mathcal{P} -schemes of double cosets. In Chapter 3, we proved that for a subfield K of the splitting field L of \tilde{f} , G the Galois group of \tilde{f} , and $H = \text{Gal}(L/K)$, an idempotent decomposition of the ring $\bar{\mathcal{O}}_K$ corresponds to a partition of the right coset space $H \backslash G$. The crucial condition for this claim to hold is that p splits completely in the splitting field L of \tilde{f} , which in turn relies on the assumption that f is square-free and completely reducible over the field of definition. In general, one can prove that an idempotent decomposition of $\bar{\mathcal{O}}_K$ corresponds to a partition of the *double coset space* $H \backslash G / \mathcal{D}_{\Omega_0}$ instead of the right coset space $H \backslash G$, where $\mathcal{D}_{\Omega_0} \subseteq G$ is known

as the *decomposition group* (of a fixed prime ideal \mathfrak{Q}_0 of \mathcal{O}_L over K_0). For the special case studied in Chapter 3, the decomposition group $\mathcal{D}_{\mathfrak{Q}_0}$ is trivial, and hence the double coset space $H \backslash G / \mathcal{D}_{\mathfrak{Q}_0}$ coincides with the right coset space $H \backslash G$.

To address the general case, we define the notion of \mathcal{P} -collections (resp. \mathcal{P} -schemes) of double cosets, generalizing (ordinary) \mathcal{P} -collections (resp. \mathcal{P} -schemes). Various properties including (strong) antisymmetry, discreteness and homogeneity can be extended to \mathcal{P} -schemes of double cosets. In addition, as the rings $\bar{\mathcal{O}}_K$ are not necessarily semisimple in general, we replace them with the rings R_K , defined by

$$R_K := \{x \in \bar{\mathcal{O}}_K / \text{Rad}(\bar{\mathcal{O}}_K) : x^p = x\},$$

where $\text{Rad}(\bar{\mathcal{O}}_K)$ denotes the *radical* of $\bar{\mathcal{O}}_K$. These rings have the advantage of being finite products of \mathbb{F}_p , so that we can directly use the results in Chapter 3. Then we generalize the algorithm in Chapter 3 to compute a collection of idempotent decompositions of the rings R_K so that they correspond to a strongly antisymmetric \mathcal{P} -schemes of double cosets.

In addition, we introduce the following notations concerning partitions of double coset spaces: for every double coset $Hg\mathcal{D}_{\mathfrak{Q}_0} \in H \backslash G / \mathcal{D}_{\mathfrak{Q}_0}$ where $H \subseteq G$, we associate two positive integers $f(Hg\mathcal{D}_{\mathfrak{Q}_0})$ and $e(Hg\mathcal{D}_{\mathfrak{Q}_0})$, called the *inertia degree* and the *ramification index* of $Hg\mathcal{D}_{\mathfrak{Q}_0}$ respectively.⁵ Then we say a partition P of $H \backslash G / \mathcal{D}_{\mathfrak{Q}_0}$ has *locally constant inertia degrees* (resp. *ramification indices*) if for every block B in P , all the double cosets in B have the same inertia degree (resp. ramification index). We design efficient algorithms that force the partitions in a \mathcal{P} -collection of double cosets to have locally constant inertia degrees and ramification indices. These algorithms may be regarded as the analogues of distinct-degree factorization [CZ81] and square-free factorization [Yun76; Knu98] that factorize a polynomial according to the degrees and the multiplicities of the irreducible factors.

The discussion above is summarized by the following theorem, which generalizes Theorem 3.1 in Chapter 3.

Theorem 5.2 (informal). *Under GRH, there exists a deterministic algorithm that given a poset $\mathcal{P}^\#$ of number fields between K_0 and L corresponding to a poset \mathcal{P} of subgroups of G , outputs idempotent decompositions of R_K for $K \in \mathcal{P}^\#$*

⁵These names come from the fact that $f(Hg\mathcal{D}_{\mathfrak{Q}_0})$ (resp. $e(Hg\mathcal{D}_{\mathfrak{Q}_0})$) is the inertia degree (resp. ramification index) of the prime ideal of \mathcal{O}_{L^H} lying over p corresponding to $Hg\mathcal{D}_{\mathfrak{Q}_0}$. See Definition 5.2 for details.

corresponding to a strongly antisymmetric \mathcal{P} -scheme of double cosets \mathcal{C} with respect to \mathcal{D}_{Ω_0} . Moreover, all the partitions in \mathcal{C} have locally constant inertia degrees and ramification indices. The running time is polynomial in the size of the input.

From a \mathcal{P} -scheme of double cosets to an ordinary \mathcal{P} -scheme. Theorem 5.2 is still not enough for proving our main result (Theorem 5.1), since the algorithm in Theorem 5.2 only produces a strongly antisymmetric \mathcal{P} -scheme of double cosets rather than an (ordinary) \mathcal{P} -scheme. While strongly antisymmetric \mathcal{P} -schemes of double cosets are interesting objects, we do not know if their existence implies the existence of strongly antisymmetric (ordinary) \mathcal{P} -schemes.

To overcome this problem, we strengthen the algorithm by maintaining not only idempotent decompositions of a collection of rings R_K , but also elements in rings of the form $\bar{\mathcal{O}}_K$ or $(\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)) \otimes_{\mathbb{F}_q} \mathbb{F}_{q^i}$, $i \in \mathbb{N}^+$. More specifically, we compute auxiliary elements $s_\delta \in \bar{\mathcal{O}}_K$ (resp. $t_\delta \in (\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)) \otimes_{\mathbb{F}_q} \mathbb{F}_{q^i}$) for number fields K and idempotents δ . Then we define a \mathcal{P} -collection $\tilde{\mathcal{C}}$ based on these auxiliary elements and the \mathcal{P} -scheme of double cosets \mathcal{C} computed in Theorem 5.2. Moreover, we describe subroutines that properly refines the partitions in \mathcal{C} unless $\tilde{\mathcal{C}}$ is a strongly antisymmetric \mathcal{P} -scheme. This allows us to strengthen Theorem 5.2 so that the algorithm produces a strongly antisymmetric (ordinary) \mathcal{P} -scheme $\tilde{\mathcal{C}}$ in addition to a \mathcal{P} -scheme of double cosets. See Theorem 5.8 for the formal statement. Our main result (Theorem 5.1) then follows easily.

Outline of the chapter. Notations and mathematical preliminaries are given in Section 5.1, and algorithmic preliminaries are given in Section 5.2. In Section 5.3, we reduce the problem of factoring f to that of computing an idempotent decomposition of R_F . In Section 5.4, we give (a preliminary version of) the main body of the algorithm that computes idempotent decompositions corresponding to a strongly antisymmetric \mathcal{P} -scheme of double cosets. This \mathcal{P} -scheme also has the property that all of its partitions have locally constant inertia degrees and ramification indices, as guaranteed by the subroutines described in Section 5.5 and Section 5.6.

The next three sections address the problem of producing an (ordinary) \mathcal{P} -scheme from the above \mathcal{P} -scheme of double cosets. More specifically, in Section 5.7, we give a subroutine that computes the auxiliary elements s_δ and t_δ , and use these elements to define a \mathcal{P} -collection $\tilde{\mathcal{C}}$. In Section 5.8, we introduce a property about \mathcal{P} -collections called $(\mathcal{C}, \mathcal{D})$ -separatedness, and use it to give a criterion for $\tilde{\mathcal{C}}$ being a strongly antisymmetric \mathcal{P} -scheme. In Section 5.9, we modify the algorithm in

Section 5.4 to produce a strongly antisymmetric \mathcal{P} -scheme, based on the results in Section 5.7 and Section 5.8.

Finally, in Section 5.10, we combine the results in previous sections to obtain the generalized \mathcal{P} -scheme algorithm, and use it to prove the main result of this chapter (Theorem 5.1). Using the algorithm, we also obtain generalizations of the main results in [Hua91a; Hua91b; Rón88; Rón92; Evd92; Evd94; IKS09].

5.1 Preliminaries

For a number field K , denote by $\bar{\mathcal{O}}_K$ the quotient ring $\mathcal{O}_K/p\mathcal{O}_K$. For $K_0 = \mathbb{Q}[Y]/(\tilde{h}(Y))$, we have

Lemma 5.2. *The ideal $p\mathcal{O}_{K_0}$ is a prime ideal of \mathcal{O}_{K_0} . And $\bar{\mathcal{O}}_{K_0} \cong \mathbb{F}_q$.*

Proof. Let $\bar{Y} := Y + (\tilde{h}(Y)) \in \mathcal{O}_{K_0}$. Consider the ring homomorphism $i : \mathbb{F}_p[Y]/(h(Y)) \rightarrow \bar{\mathcal{O}}_{K_0}$ sending $Y + (h(Y))$ to $\bar{Y} + p\mathcal{O}_{K_0}$. Clearly i is a nonzero map since $i(1) = 1$. As $\mathbb{F}_p[Y]/(h(Y))$ is a field, the map i is injective. As $\mathbb{F}_p[Y]/(h(Y))$ and $\bar{\mathcal{O}}_{K_0}$ both have dimension $\deg(h)$ over \mathbb{F}_p , the map i is an isomorphism. So $\bar{\mathcal{O}}_{K_0} \cong \mathbb{F}_p[Y]/(h(Y)) \cong \mathbb{F}_q$ and $p\mathcal{O}_{K_0}$ is prime. \square

In the following, we give some notations and facts from algebraic number theory. The proofs can be found in standard references like [Neu99].

Splitting of prime ideals. Let K be a finite extension of K_0 . The ideal $p\mathcal{O}_K$ splits in the unique way into a product of prime ideals of \mathcal{O}_K , up to the ordering:

$$p\mathcal{O}_K = \prod_{i=1}^k \mathfrak{P}_i^{e(\mathfrak{P}_i)} = \bigcap_{i=1}^k \mathfrak{P}_i^{e(\mathfrak{P}_i)},$$

where $\mathfrak{P}_1, \dots, \mathfrak{P}_k$ are distinct and $e(\mathfrak{P}_i) \in \mathbb{N}^+$. We say $\mathfrak{P}_1, \dots, \mathfrak{P}_k$ are the prime ideals of \mathcal{O}_K lying over p . For $i \in [k]$, define $\kappa_{\mathfrak{P}_i} := \mathcal{O}_K/\mathfrak{P}_i$ which is a finite field, called the *residue field* of \mathfrak{P}_i . The inclusion $\mathcal{O}_{K_0} \hookrightarrow \mathcal{O}_K$ induces an embedding of $\bar{\mathcal{O}}_{K_0} \cong \mathbb{F}_q$ in $\kappa_{\mathfrak{P}_i}$, making $\kappa_{\mathfrak{P}_i}$ an extension field of $\bar{\mathcal{O}}_{K_0}$. Let $f(\mathfrak{P}_i) := [\kappa_{\mathfrak{P}_i} : \bar{\mathcal{O}}_{K_0}]$. We call $e(\mathfrak{P}_i)$ and $f(\mathfrak{P}_i)$ the *ramification index* and the *inertia degree* of \mathfrak{P}_i (over $p\mathcal{O}_{K_0}$) respectively. It holds that

$$\sum_{i=1}^k e(\mathfrak{P}_i) f(\mathfrak{P}_i) = [K : K_0].$$

Vector spaces $\mathfrak{P}^i/\mathfrak{P}^{i+1}$. We also use the following facts implicitly:

For a number field K , $i \in \mathbb{N}$ and a nonzero prime ideal \mathfrak{P} of \mathcal{O}_K , the abelian group $\mathfrak{P}^i/\mathfrak{P}^{i+1}$ is an one-dimensional vector space over the field $\kappa_{\mathfrak{P}} = \mathcal{O}_K/\mathfrak{P}$, where the scalar multiplication is defined by

$$(u + \mathfrak{P}) \cdot (v + \mathfrak{P}^{i+1}) = uv + \mathfrak{P}^{i+1} \quad \text{for } u \in \mathcal{O}_K, v \in \mathfrak{P}^i.$$

For $i, j \in \mathbb{N}$ and $u \in \mathfrak{P}^i - \mathfrak{P}^{i+1}$, the map

$$x + \mathfrak{P}^{j+1} \mapsto ux + \mathfrak{P}^{i+j+1}$$

is an isomorphism from $\mathfrak{P}^j/\mathfrak{P}^{j+1}$ to $\mathfrak{P}^{i+j}/\mathfrak{P}^{i+j+1}$, both regarded as vector spaces over $\kappa_{\mathfrak{P}}$. In particular, for $i, j \in \mathbb{N}$ and $u \in \mathfrak{P}^i - \mathfrak{P}^{i+1}$, we have $u^j \in \mathfrak{P}^{ij} - \mathfrak{P}^{ij+1}$.

Now suppose K, K' are finite extensions of K_0 and $K \subseteq K'$. And $\mathfrak{P}, \mathfrak{Q}$ are prime ideals of \mathcal{O}_K and $\mathcal{O}_{K'}$ respectively, both lying over p , such that $\mathfrak{Q} \cap \mathcal{O}_K = \mathfrak{P}$. Then $e(\mathfrak{P})$ divides $e(\mathfrak{Q})$ and $f(\mathfrak{P})$ divides $f(\mathfrak{Q})$. And for $i \in \mathbb{N}$, the inclusion $\mathcal{O}_K \hookrightarrow \mathcal{O}_{K'}$ induces an inclusion $\mathfrak{P}^i/\mathfrak{P}^{i+1} \hookrightarrow \mathfrak{Q}^{i'}/\mathfrak{Q}^{i'+1}$ where $i' = i \cdot e(\mathfrak{Q})/e(\mathfrak{P})$.

The decomposition group and the inertia group. Let L/K_0 be a Galois extension of number fields with the Galois group $G = \text{Gal}(L/K_0)$. Let \mathfrak{P} be a prime ideal of \mathcal{O}_L lying over p . The group

$$\mathcal{D}_{\mathfrak{P}} := \{g \in G : {}^g\mathfrak{P} = \mathfrak{P}\} \subseteq G$$

is called the *decomposition group* of \mathfrak{P} over K_0 . And the group

$$\mathcal{I}_{\mathfrak{P}} := \{g \in G : {}^g x \equiv x \pmod{\mathfrak{P}} \text{ for all } x \in \mathcal{O}_L\}$$

is a normal subgroup of $\mathcal{D}_{\mathfrak{P}}$, called the *inertia group* of \mathfrak{P} over K_0 . Each automorphism $g \in \mathcal{D}_{\mathfrak{P}}$ of L restricts to an automorphism of \mathcal{O}_L fixing \mathcal{O}_{K_0} and satisfying ${}^g\mathfrak{P} = \mathfrak{P}$, and hence induces an automorphism \bar{g} of the residue field $\kappa_{\mathfrak{P}}$ fixing $\bar{\mathcal{O}}_{K_0}$, defined by

$$\bar{g}(x + \mathfrak{P}) = {}^g x + \mathfrak{P}.$$

The map $\pi : g \mapsto \bar{g}$ is a surjective group homomorphism from $\mathcal{D}_{\mathfrak{P}}$ to $\text{Gal}(\kappa_{\mathfrak{P}}/\bar{\mathcal{O}}_{K_0})$ whose kernel is precisely $\mathcal{I}_{\mathfrak{P}}$, i.e., we have a short exact sequence

$$1 \rightarrow \mathcal{I}_{\mathfrak{P}} \rightarrow \mathcal{D}_{\mathfrak{P}} \xrightarrow{\pi} \text{Gal}(\kappa_{\mathfrak{P}}/\bar{\mathcal{O}}_{K_0}) \rightarrow 1.$$

The Galois group $\text{Gal}(\kappa_{\mathfrak{P}}/\bar{\mathcal{O}}_{K_0})$ is cyclic and is generated by the *Frobenius automorphism* $x \mapsto x^q$ of $\kappa_{\mathfrak{P}}$ over $\bar{\mathcal{O}}_{K_0} \cong \mathbb{F}_q$.

The wild inertia group. Let L , G and \mathfrak{P} be as above. The group

$$\mathcal{W}_{\mathfrak{P}} := \{g \in G : {}^g x \equiv x \pmod{\mathfrak{P}^2} \text{ for all } x \in \mathcal{O}_L\}.$$

is a normal subgroup of $\mathcal{I}_{\mathfrak{P}}$, called the *wild inertia group* of \mathfrak{P} over K_0 .

Choose $\pi_L \in \mathfrak{P} - \mathfrak{P}^2$. We have a group homomorphism $\mathcal{I}_{\mathfrak{P}} \rightarrow \kappa_{\mathfrak{P}}^{\times}$ sending $g \in \mathcal{I}_{\mathfrak{P}}$ to the unique element $c_g \in \kappa_{\mathfrak{P}}^{\times}$ satisfying ${}^g \pi_L + \mathfrak{P}^2 = c_g(\pi_L + \mathfrak{P}^2)$. This map is independent of the choice of π_L , and its kernel is precisely $\mathcal{W}_{\mathfrak{P}}$. It is also known that $\mathcal{W}_{\mathfrak{P}}$ is a p -group. See [Neu99, Section II.10].

In our factoring algorithm, the group G is a subgroup of $\text{Sym}(n)$ where n is the degree of the input polynomial $f(X) \in \mathbb{F}_q[X]$. We can always assume $p > n$, since the case $p \leq n$ is solved in polynomial time by Berlekamp's algorithm in [Ber70]. Under this assumption, the p -subgroup $\mathcal{W}_{\mathfrak{P}}$ of G is trivial, and hence the map $\mathcal{I}_{\mathfrak{P}} \rightarrow \kappa_{\mathfrak{P}}^{\times}$ above is injective. In particular, the inertia group $\mathcal{I}_{\mathfrak{P}}$ is cyclic.

Prime ideals vs. double cosets. We have the following generalization of Theorem 3.4, which gives a one-to-one correspondence between prime ideals lying over p and double cosets. See [Neu99] for its proof.

Theorem 5.3. *Let L/K_0 be a Galois extension of number fields and let $G = \text{Gal}(L/K_0)$. Fix a prime ideal \mathfrak{Q}_0 of \mathcal{O}_L lying over p . For any subgroup $H \subseteq G$ and its fixed field $K = L^H$, the map $Hg\mathcal{D}_{\mathfrak{Q}_0} \mapsto {}^g \mathfrak{Q}_0 \cap \mathcal{O}_K$ is a one-to-one correspondence between the double cosets in $H \backslash G / \mathcal{D}_{\mathfrak{Q}_0}$ and the prime ideals of \mathcal{O}_K lying over p .⁶ Moreover, for $g \in G$ and the prime ideal $\mathfrak{P} = {}^g \mathfrak{Q}_0 \cap \mathcal{O}_K$ corresponding to $Hg\mathcal{D}_{\mathfrak{Q}_0}$, define*

$$n(\mathfrak{P}) := |\{Hh \in H \backslash G : Hh\mathcal{D}_{\mathfrak{Q}_0} = Hg\mathcal{D}_{\mathfrak{Q}_0}\}|.$$

Then

$$e(\mathfrak{P}) = |\{Hh \in H \backslash G : Hh\mathcal{I}_{\mathfrak{Q}_0} = Hg\mathcal{I}_{\mathfrak{Q}_0}\}| \quad \text{and} \quad f(\mathfrak{P}) = \frac{n(\mathfrak{P})}{e(\mathfrak{P})}.$$

Motivated by Theorem 5.3, we define the ramification index and the inertia degree of a double coset:

⁶Note that this map is well defined: for another representative $hgh' \in G$ of $Hg\mathcal{D}_{\mathfrak{Q}_0}$, where $h \in H$ and $h' \in \mathcal{D}_{\mathfrak{Q}_0}$, we have ${}^{hgh'} \mathfrak{Q}_0 \cap \mathcal{O}_K = {}^{hg} \mathfrak{Q}_0 \cap \mathcal{O}_K = {}^h({}^g \mathfrak{Q}_0 \cap \mathcal{O}_K) = {}^g \mathfrak{Q}_0 \cap \mathcal{O}_K$ since ${}^{h'} \mathfrak{Q}_0 = \mathfrak{Q}_0$ and \mathcal{O}_K is fixed by H .

Definition 5.2. Let G be a finite group, H, \mathcal{D} subgroups of G , and \mathcal{I} a normal subgroup of \mathcal{D} . Define the ramification index of a double coset $Hg\mathcal{D} \in H \backslash G / \mathcal{D}$ with respect to $(\mathcal{D}, \mathcal{I})$ to be

$$e(Hg\mathcal{D}) := |\{Hh \in H \backslash G : Hh\mathcal{I} = Hg\mathcal{I}\}|,$$

which is well defined.⁷ And define the inertia degree of $Hg\mathcal{D}$ with respect to $(\mathcal{D}, \mathcal{I})$ to be

$$f(Hg\mathcal{D}) := \frac{|\{Hh \in H \backslash G : Hh\mathcal{D} = Hg\mathcal{D}\}|}{e(Hg\mathcal{D})}.$$

Suppose L/K_0 is a Galois extension of number fields with the Galois group G . Fix a prime ideal \mathfrak{Q}_0 of \mathcal{O}_L lying over p . Let H be a subgroup of G and $K = L^H$. Then by Theorem 5.3, the ramification index (resp. inertia degree) of a double coset $Hg\mathcal{D}_{\mathfrak{Q}_0} \in H \backslash G / \mathcal{D}_{\mathfrak{Q}_0}$ with respect to $(\mathcal{D}_{\mathfrak{Q}_0}, \mathcal{I}_{\mathfrak{Q}_0})$ is precisely the ramification index (resp. inertia degree) of the corresponding prime ideal ${}^g\mathfrak{Q}_0 \cap \mathcal{O}_K$ of \mathcal{O}_K .

We also introduce the following notations concerning partitions of a double coset space.

Definition 5.3. Let $G, H, \mathcal{D}, \mathcal{I}$ be as in Definition 5.2. We say a partition P of $H \backslash G / \mathcal{D}$ has locally constant ramification indices (resp. inertia degrees) with respect to $(\mathcal{D}, \mathcal{I})$ if for every $B \in P$, all the double cosets in B have the same ramification index (resp. inertia degree) with respect to $(\mathcal{D}, \mathcal{I})$. For such a partition P and any $B \in P$, denote by $e(B)$ (resp. $f(B)$) the ramification index (resp. inertia degree) of any double coset in B .

Radicals of rings and polynomials. Let A be a (commutative) ring. An element $x \in A$ is *nilpotent* if $x^k = 0$ for some $k \in \mathbb{N}^+$. The *radical* (or *nilradical*) of A , denoted by $\text{Rad}(A)$, is the ideal consisting of the nilpotent elements of A . It equals the intersection of all the prime ideals of A (see [AM69]).

Let $g(X) \in \mathbb{F}_q[X]$ be a non-constant polynomial with the following factorization

$$g(X) = c \cdot \prod_{i=1}^k (g_i(X))^{m_i}$$

⁷To see that $e(Hg\mathcal{D})$ is well defined, consider two representatives g and g' of $Hg\mathcal{D}$. Then $g' = sgt$ for some $s \in H$ and $t \in \mathcal{D}$. Note that $Hht\mathcal{I} = Hh\mathcal{I}t$ for all $h \in G$. It follows that the map $Hh \mapsto Hht$ is a bijection from $\{Hh \in H \backslash G : Hh\mathcal{I} = Hg\mathcal{I}\}$ to $\{Hh \in H \backslash G : Hh\mathcal{I} = Hg'\mathcal{I}\}$.

over \mathbb{F}_q , where $c \in \mathbb{F}_q$ is the leading coefficient of g and g_1, \dots, g_k are distinct monic irreducible polynomials over \mathbb{F}_q . Define the *radical* $\text{Rad}(g)$ of g to be the monic polynomial $\prod_{i=1}^k g_i(X) \in \mathbb{F}_q[X]$. For $A = \mathbb{F}_q[X]/(g(X))$, the ideal of A generated by $\text{Rad}(g) + (g(X)) \in A$ is precisely $\text{Rad}(A)$.

The ring R_K . Suppose K is a finite extension of K_0 and $p\mathcal{O}_K$ splits into the product of prime ideals

$$p\mathcal{O}_K = \prod_{i=1}^k \mathfrak{P}_i^{e(\mathfrak{P}_i)},$$

where $\mathfrak{P}_1, \dots, \mathfrak{P}_k$ are distinct. The radical of $\bar{\mathcal{O}}_K$ is given by

$$\text{Rad}(\bar{\mathcal{O}}_K) = \bigcap_{i=1}^k \mathfrak{P}_i/p\mathcal{O}_K = \left(\bigcap_{i=1}^k \mathfrak{P}_i \right) / p\mathcal{O}_K.$$

By the Chinese remainder theorem, we have the isomorphism

$$\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K) \rightarrow \prod_{i=1}^k \mathcal{O}_K/\mathfrak{P}_i = \prod_{i=1}^k \kappa_{\mathfrak{P}_i},$$

sending $x + \text{Rad}(\bar{\mathcal{O}}_K) \in \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ to $(\tilde{x} \bmod \mathfrak{P}_1, \dots, \tilde{x} \bmod \mathfrak{P}_k)$, where $\tilde{x} \in \mathcal{O}_K$ is an arbitrary element lifting $x \in \bar{\mathcal{O}}_K$. In particular, the ring $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ is semisimple.

Define R_K to be the subring of $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ consisting of elements fixed by the Frobenius automorphism $x \mapsto x^p$ over \mathbb{F}_p , i.e.,

$$R_K := \{x \in \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K) : x^p = x\}.$$

The isomorphism $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K) \rightarrow \prod_{i=1}^k \kappa_{\mathfrak{P}_i}$ above identifies R_K with the subring $\prod_{i=1}^k \mathbb{F}_p$ of $\prod_{i=1}^k \kappa_{\mathfrak{P}_i}$. So R_K is a finite product of copies of \mathbb{F}_p and in particular is semisimple.

Observe that the map $\mathfrak{m} \mapsto (\mathfrak{m}/\text{Rad}(\bar{\mathcal{O}}_K)) \cap R_K$ is a one-to-one correspondence between the maximal ideals of $\bar{\mathcal{O}}_K$ and those of R_K . Combining this fact with Theorem 5.3, we obtain

Lemma 5.3. *Let L, G, \mathfrak{Q}_0 be as in Theorem 5.3. For any subgroup $H \subseteq G$ and its fixed field $K = L^H$, the map*

$$Hg\mathcal{D}_{\mathfrak{Q}_0} \mapsto \frac{({}^g\mathfrak{Q}_0 \cap \mathcal{O}_K)/p\mathcal{O}_K}{\text{Rad}(\bar{\mathcal{O}}_K)} \cap R_K$$

is a one-to-one correspondence between the double cosets in $H \backslash G / \mathcal{D}_{\mathfrak{Q}_0}$ and the maximal ideals of R_K .

Idempotent decompositions vs. partitions of a double coset space. In the following, we establish a one-to-one correspondence between the idempotent decompositions of R_K and the partitions of a certain double coset space.

For a number field extension L/K , the inclusion $\mathcal{O}_K \hookrightarrow \mathcal{O}_L$ induces an inclusion $\bar{\mathcal{O}}_K \hookrightarrow \bar{\mathcal{O}}_L$. So we may regard $\bar{\mathcal{O}}_K$ as a subring of $\bar{\mathcal{O}}_L$. Note that $\text{Rad}(\bar{\mathcal{O}}_L) \cap \bar{\mathcal{O}}_K = \text{Rad}(\bar{\mathcal{O}}_K)$. Passing to the quotient rings yields an inclusion $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K) \hookrightarrow \bar{\mathcal{O}}_L/\text{Rad}(\bar{\mathcal{O}}_L)$. Restricting to the subring R_K , we obtain an inclusion

$$i_{K,L} : R_K \hookrightarrow R_L.$$

Also note that if L/K_0 is a Galois extension with the Galois group G , the action of G on \mathcal{O}_L induces an action on R_L that permutes the maximal ideals of R_L .

Fix the following notations: let L be a Galois extension of K_0 with the Galois group $G = \text{Gal}(L/K_0)$. For a (nonzero) prime ideal $\bar{\mathfrak{Q}}$ of $\bar{\mathcal{O}}_L$ lying over p , define

$$\bar{\mathfrak{Q}} := \frac{\bar{\mathfrak{Q}}/p\bar{\mathcal{O}}_L}{\text{Rad}(\bar{\mathcal{O}}_L)} \cap R_L,$$

which is a maximal ideal of R_L , and let $\delta_{\bar{\mathfrak{Q}}}$ be the primitive idempotent of $\bar{\mathcal{O}}_L$ satisfying $\delta_{\bar{\mathfrak{Q}}} \equiv 1 \pmod{\bar{\mathfrak{Q}}}$ and $\delta_{\bar{\mathfrak{Q}}} \equiv 0 \pmod{\bar{\mathfrak{Q}}'}$ for all maximal ideal $\bar{\mathfrak{Q}}' \neq \bar{\mathfrak{Q}}$ of $\bar{\mathcal{O}}_L$. Finally, fix a prime ideal $\bar{\mathfrak{Q}}_0$ of $\bar{\mathcal{O}}_L$ lying over p .

Definition 5.4. *Suppose H is a subgroup of G and $K = L^H$. Then*

- *for an idempotent decomposition I of R_K , define $P(I)$ to be the partition of $H \backslash G / \mathcal{D}_{\bar{\mathfrak{Q}}_0}$ where $Hg\mathcal{D}_{\bar{\mathfrak{Q}}_0}, Hg'\mathcal{D}_{\bar{\mathfrak{Q}}_0}$ are in the same block iff $g^{-1}(i_{K,L}(\delta)) \equiv g'^{-1}(i_{K,L}(\delta)) \pmod{\bar{\mathfrak{Q}}_0}$ holds for all $\delta \in I$, and*
- *for a partition P of $H \backslash G / \mathcal{D}_{\bar{\mathfrak{Q}}_0}$, define $I(P)$ to be the idempotent decomposition of R_K consisting of the idempotents*

$$\delta_B := i_{K,L}^{-1} \left(\sum_{g\mathcal{D}_{\bar{\mathfrak{Q}}_0} \in G/\mathcal{D}_{\bar{\mathfrak{Q}}_0} : Hg\mathcal{D}_{\bar{\mathfrak{Q}}_0} \in B} {}^g\delta_{\bar{\mathfrak{Q}}_0} \right),$$

where B ranges over the blocks in P .

We have the following two lemmas that generalize Lemma 3.5 and Lemma 3.6 respectively. In particular, Lemma 5.5 establishes a one-to-one correspondence between the idempotent decompositions of R_K and the partitions of $H \backslash G / \mathcal{D}_{\bar{\mathfrak{Q}}_0}$.

Lemma 5.4. *The partitions $P(I)$ and the idempotent decompositions $I(P)$ are well defined. And for any idempotent decomposition I of \bar{O}_K , the idempotents $\delta \in I$ correspond one-to-one to the blocks of $P(I)$ via the map*

$$\delta \mapsto B_\delta := \{Hg\mathcal{D}_{\Omega_0} \in H \backslash G / \mathcal{D}_{\Omega_0} : g^{-1}(i_{K,L}(\delta)) \equiv 1 \pmod{\bar{\Omega}_0}\}$$

with the inverse map $B \mapsto \delta_B$.

Lemma 5.5. *The map $I \mapsto P(I)$ is a one-to-one correspondence between the idempotent decompositions of R_K and the partitions of $H \backslash G / \mathcal{D}_{\Omega_0}$, with the inverse map $P \mapsto I(P)$.*

Their proofs are similar to those of Lemma 3.5 and Lemma 3.6, and can be found in Appendix C.

\mathcal{P} -collections and \mathcal{P} -schemes of double cosets. Let G be a finite group and $\mathcal{D} \subseteq G$ a subgroup. We generalize projections and conjugations introduced in Chapter 2 so that they are defined between double coset spaces:

- (projection) for $H \subseteq H' \subseteq G$, define the *projection* $\pi_{H,H'}^{\mathcal{D}} : H \backslash G / \mathcal{D} \rightarrow H' \backslash G / \mathcal{D}$ to be the map sending $Hg\mathcal{D} \in H \backslash G / \mathcal{D}$ to $H'g\mathcal{D} \in H' \backslash G / \mathcal{D}$, and
- (conjugation) for $H \subseteq G$ and $g \in G$, define the *conjugation* $c_{H,g}^{\mathcal{D}} : H \backslash G / \mathcal{D} \rightarrow gHg^{-1} \backslash G / \mathcal{D}$ to be the map sending $Hh\mathcal{D} \in H \backslash G / \mathcal{D}$ to $(gHg^{-1})gh\mathcal{D} \in gHg^{-1} \backslash G / \mathcal{D}$.

Next we define \mathcal{P} -collections and \mathcal{P} -schemes of double cosets.

Definition 5.5. *Let \mathcal{P} be a subgroup system over a finite group G . Then a \mathcal{P} -collection of double cosets with respect to a subgroup \mathcal{D} of G is a family $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ indexed by \mathcal{P} where each C_H is a partition of $H \backslash G / \mathcal{D}$. Moreover, \mathcal{C} is a \mathcal{P} -scheme of double cosets with respect to \mathcal{D} if it has the following properties:*

- (compatibility) for $H, H' \in \mathcal{P}$ with $H \subseteq H'$ and $x, x' \in H \backslash G / \mathcal{D}$ in the same block of C_H , the images $\pi_{H,H'}^{\mathcal{D}}(x)$ and $\pi_{H,H'}^{\mathcal{D}}(x')$ are in the same block of $C_{H'}$.
- (invariance) for $H \in \mathcal{P}$ and $g \in G$, the map $c_{H,g}^{\mathcal{D}} : H \backslash G / \mathcal{D} \rightarrow gHg^{-1} \backslash G / \mathcal{D}$ maps any block of C_H to a block of $C_{gHg^{-1}}$.

- (regularity) for $H, H' \in \mathcal{P}$ with $H \subseteq H'$, $B \in C_H$, $B' \in C_{H'}$, the number of $x \in B$ satisfying $\pi_{H,H'}^{\mathcal{D}}(x) = y$ is a constant when y ranges over the elements of B' .

We also define the following optional properties for a \mathcal{P} -scheme of double cosets \mathcal{C} with respect to \mathcal{D} :

- (homogeneity and discreteness) \mathcal{C} is homogeneous on $H \in \mathcal{P}$ if $C_H = 0_{H \backslash G / \mathcal{D}}$, and otherwise inhomogeneous on H . It is discrete on H if $C_H = \infty_{H \backslash G / \mathcal{D}}$, and otherwise non-discrete on H .
- (antisymmetry) \mathcal{C} is antisymmetric if for $H \in \mathcal{P}$, $g \in N_G(H)$, $B \in C_H$ and $Hg\mathcal{D} \in B$, either $c_{H,g}^{\mathcal{D}}(Hg\mathcal{D}) = Hg\mathcal{D}$ or $c_{H,g}^{\mathcal{D}}(Hg\mathcal{D}) \notin B$.
- (strong antisymmetry) \mathcal{C} is strongly antisymmetric if for any sequence of subgroups $H_0, \dots, H_k \in \mathcal{P}$, $B_0 \in C_{H_0}, \dots, B_k \in C_{H_k}$, and maps $\sigma_1, \dots, \sigma_k$ satisfying

- σ_i is a bijective map from B_{i-1} to B_i ,
- σ_i is of the form $c_{H_{i-1},g}^{\mathcal{D}}|_{B_{i-1}}$, $\pi_{H_{i-1},H_i}^{\mathcal{D}}|_{B_{i-1}}$, or $(\pi_{H_i,H_{i-1}}^{\mathcal{D}}|_{B_i})^{-1}$,
- $H_0 = H_k$ and $B_0 = B_k$,

the composition $\sigma_k \circ \dots \circ \sigma_1$ is the identity map on $B_0 = B_k$.

The notions of \mathcal{P} -collections and \mathcal{P} -schemes introduced in Chapter 2 correspond to the special case that \mathcal{D} is trivial.

Extension of scalars of $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$. In Section 5.7–5.8, we need a family of rings $A_{K,i}$ that are obtained from $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ via “extension of scalars”, whose definitions are given below.

Let K be a finite extension of K_0 . The inclusion $A_0 \subseteq \mathcal{O}_{K_0} \hookrightarrow \mathcal{O}_K$ induces an embedding of $\mathbb{F}_q \cong A_0/pA_0$ in $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$, endowing $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ the structure of an \mathbb{F}_q -algebra. For $i \in \mathbb{N}^+$, we define the tensor product

$$A_{K,i} := (\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)) \otimes_{\mathbb{F}_q} \mathbb{F}_{q^i},$$

which is an \mathbb{F}_{q^i} -algebra and is spanned by tensors $a \otimes b$ over \mathbb{F}_q where $a \in \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ and $b \in \mathbb{F}_{q^i}$ (see [AM69] for the definition of tensor products of

rings). Intuitively, the ring $A_{K,i}$ is obtained from $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ by extending the scalars from \mathbb{F}_q to \mathbb{F}_{q^i} . And $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ is naturally identified with a subring of $A_{K,i}$ via $a \mapsto a \otimes 1$. As $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ is semisimple, so is $A_{K,i}$.⁸ The Frobenius automorphism $x \mapsto x^q$ of $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ over \mathbb{F}_q induces an automorphism of $A_{K,i}$ over \mathbb{F}_{q^i} sending $a \otimes b$ to $a^q \otimes b$. We denote this automorphism by $\sigma_{K,i}$.

The following lemma is also needed, whose proof is deferred to Appendix C.

Lemma 5.6. *For any maximal ideal \mathfrak{m} of $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$, the group $\langle \sigma_{K,i} \rangle$ generated by $\sigma_{K,i}$ acts transitively on the set of the maximal ideal of $A_{K,i}$ containing \mathfrak{m} .*

Suppose K, K' are extensions of K_0 and $K \subseteq K'$. Then the inclusion $\mathcal{O}_K \hookrightarrow \mathcal{O}_{K'}$ induces an embedding $\iota : \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K) \hookrightarrow \bar{\mathcal{O}}_{K'}/\text{Rad}(\bar{\mathcal{O}}_{K'})$, which in turn induces a ring homomorphism $\iota' : A_{K,i} \hookrightarrow A_{K',i}$ sending $a \otimes b$ to $\iota(a) \otimes b$. The map ι' is injective since \mathbb{F}_{q^i} is a *flat* \mathbb{F}_q -module (see, e.g., [AM69, Proposition 2.19 and Exercise 2.4]). This allows us to regard $A_{K,i}$ as a subring of $A_{K',i}$. Note that $\iota' \circ \sigma_{K,i} = \sigma_{K',i} \circ \iota'$.

Finally, suppose L/K_0 is a finite Galois extension with the Galois group G . The action of G on L induces an action on $\bar{\mathcal{O}}_L/\text{Rad}(\bar{\mathcal{O}}_L)$, which in turn induces an action on $A_{L,i}$ via ${}^g(a \otimes b) := {}^g a \otimes b$. This action commutes with $\sigma_{L,i}$.⁹

5.2 Algorithmic preliminaries

In this section, we discuss some basic procedures used in the algorithm.

Computation of radicals, and square-free factorization. We need to compute the radical of a finite dimensional (commutative) \mathbb{F}_p -algebra. This problem was studied in [FR85; Rón90] and solved in polynomial time in the more general setting of associative algebras. We state their result but restrict to the special case of commutative algebras.

Theorem 5.4 ([FR85; Rón90]). *There exists a polynomial-time algorithm that given a finite dimensional (commutative) \mathbb{F}_p -algebra A , computes an \mathbb{F}_p -basis of $\text{Rad}(A)$ in A .*

⁸We use the fact that $\mathbb{F}_{q^d} \otimes_{\mathbb{F}_q} \mathbb{F}_{q^i}$ is semisimple for $d, i \in \mathbb{N}^+$: suppose $\mathbb{F}_{q^d} \cong \mathbb{F}_q[X]/(g(X))$ where $g(X) \in \mathbb{F}_q[X]$ is irreducible over \mathbb{F}_q . Then $\mathbb{F}_{q^d} \otimes_{\mathbb{F}_q} \mathbb{F}_{q^i} \cong \mathbb{F}_{q^i}[X]/(g(X)) \cong \prod_{j=1}^k \mathbb{F}_{q^i}[X]/(g_j(X))$ where g_1, \dots, g_k are the irreducible factors of g over \mathbb{F}_{q^i} .

⁹This follows from the fact that the action of G on $\bar{\mathcal{O}}_L/\text{Rad}(\bar{\mathcal{O}}_L)$ respects the multiplication and hence commutes with the automorphism $x \mapsto x^q$.

See, e.g., [Rón90, Theorem 2.7].

Next we discuss the problem of computing the radical of a nonzero polynomial $g(X) \in \mathbb{F}_q[X]$. This is solved via *square-free factorization*.

Definition 5.6. A square-free factorization of a nonzero polynomial $g(X) \in \mathbb{F}_q[X]$ over \mathbb{F}_q is a factorization

$$g(X) = c \cdot \prod_{i=1}^k (g_i(X))^{m_i},$$

where $c \in \mathbb{F}_q$ is the leading coefficient of g and the factors $g_1(X), \dots, g_k(X) \in \mathbb{F}_q[X]$ are monic, square-free, and pairwise coprime.

Theorem 5.5 ([Yun76; Knu98]). *There exists a polynomial-time algorithm that computes a square-free factorization of a given nonzero polynomial $g(X) \in \mathbb{F}_q[X]$.*

Given the square-free factorization $g(X) = c \cdot \prod_{i=1}^k (g_i(X))^{m_i}$, the radical $\text{Rad}(g)$ is simply the product of $g_i(X)$. So we have

Corollary 5.1. *There exists a polynomial-time algorithm that given a nonzero polynomial $g(X) \in \mathbb{F}_q[X]$, computes its radical $\text{Rad}(g)$.*

Alternatively, we can compute $\text{Rad}(g)$ by computing the radical of $\mathbb{F}_q[X]/(g(X))$ and then its generator. The details are left to the reader.

Computation of annihilators. Let R be a (commutative) ring. For a set $S \subseteq R$, define the *annihilator* $\text{Ann}_R(S)$ of S to be the ideal

$$\text{Ann}_R(S) := \{r \in R : rs = 0 \text{ for all } s \in S\}$$

of R . When S is a singleton $\{s\}$, we also write $\text{Ann}_R(s)$ instead of $\text{Ann}_R(\{s\})$ and call it the annihilator of s .

When R is an finite dimensional \mathbb{F}_p -algebra, we can efficiently compute the annihilator $\text{Ann}_R(s)$ of an element $s \in R$ by solving the system of \mathbb{F}_p -linear equations given by $xs = 0$. Similarly, when S is an \mathbb{F}_p -subspace of R (in particular, when S is an ideal of R), we can compute $\text{Ann}_R(S)$ efficiently given R and an \mathbb{F}_p -basis B of S by solving the system of \mathbb{F}_p -linear equations $xs = 0$, where s ranges over the basis B .

Computation of various rings and ring homomorphisms. The algorithm uses relative number fields over K_0 rather than ordinary number fields, i.e., every number field is an extension of K_0 and is encoded as a K_0 -algebra $K_0[X]/(g(X))$ where $g(X) \in K_0[X]$ is irreducible over K_0 .

Given a relative number field K over K_0 , we can identify K_0 with an ordinary number field \tilde{K} by Corollary 4.1. It allows us to efficiently compute a p -maximal order \mathcal{O}'_K as well as the quotient ring $\bar{\mathcal{O}}_K$ as in Chapter 3. We can also efficiently compute the rings $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ and R_K , which are used in the generalized \mathcal{P} -scheme algorithm developed in this chapter. This is summarized by the following lemma, whose proof is deferred to Appendix C.

Lemma 5.7. *There exists a polynomial-time algorithm `ComputeRings` that given p and a relative number field K over K_0 , computes the following data*

- a p -maximal order \mathcal{O}'_K of K and the inclusion $\mathcal{O}'_K \hookrightarrow K$,
- $\bar{\mathcal{O}}_K$ and the quotient map $\mathcal{O}'_K \rightarrow \bar{\mathcal{O}}_K$,
- $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ and the quotient map $\bar{\mathcal{O}}_K \rightarrow \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$,
- R_K and the inclusion $R_K \hookrightarrow \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$,

where $\bar{\mathcal{O}}_K$, $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$, and R_K are encoded as algebras over \mathbb{F}_p and \mathcal{O}'_K is encoded as an algebra over \mathbb{Z} .

Suppose K and K' are relative number fields over K_0 and $\phi : K \rightarrow K'$ is a field embedding over K_0 . The map ϕ induces a ring homomorphism $\bar{\phi} : \bar{\mathcal{O}}_K \rightarrow \bar{\mathcal{O}}_{K'}$ sending $x + p\mathcal{O}_K \in \bar{\mathcal{O}}_K$ to $\phi(x) + p\mathcal{O}_{K'} \in \bar{\mathcal{O}}_{K'}$. As the image of a nilpotent element (resp. an element fixed by the automorphism $x \mapsto x^p$) under $\bar{\phi}$ is also nilpotent (resp. fixed by $x \mapsto x^p$), the map $\bar{\phi}$ induces a ring homomorphism $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K) \rightarrow \bar{\mathcal{O}}_{K'}/\text{Rad}(\bar{\mathcal{O}}_{K'})$, and we denote this map by $\hat{\phi}$. Finally, the map $\hat{\phi}$ restricts to a ring homomorphism $R_K \rightarrow R_{K'}$, which we denote by $\tilde{\phi}$. The maps $\bar{\phi}$, $\hat{\phi}$ and $\tilde{\phi}$ can be efficiently computed from ϕ (and some auxiliary data) by the following lemma.

Lemma 5.8. *There exists a polynomial-time algorithm `ComputeRingHoms` that given p , relative number fields K, K' over K_0 , a field embedding $\phi : K \rightarrow K'$ over K_0 , and the outputs of `ComputeRings` (see Lemma 5.7) on the inputs (K, p) and (K', p) respectively, computes the maps $\bar{\phi} : \bar{\mathcal{O}}_K \rightarrow \bar{\mathcal{O}}_{K'}$, $\hat{\phi} : \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K) \rightarrow \bar{\mathcal{O}}_{K'}/\text{Rad}(\bar{\mathcal{O}}_{K'})$ and $\tilde{\phi} : R_K \rightarrow R_{K'}$.*

See Appendix C for its proof.

5.3 Reduction to computing an idempotent decomposition of R_F

Now we start describing the generalized \mathcal{P} -scheme algorithm. It is always implicitly assumed that the prime number p , $\tilde{h}(Y) \in \mathbb{Z}[Y]$ and $h(Y) = \tilde{h}(Y) \bmod p \in \mathbb{F}_p[Y]$ are known to the algorithm, so that $\mathbb{F}_p[Y]/(h(Y))$, $A_0 = \mathbb{Z}[Y]/(\tilde{h}(Y))$ and $K_0 = \mathbb{Q}[X]/(\tilde{h}(Y))$ are also known. And $\mathbb{F}_p[Y]/(h(Y))$ is identified with a finite field \mathbb{F}_q via an isomorphism $\psi_0 : \mathbb{F}_p[Y]/(h(Y)) \rightarrow \mathbb{F}_q$ that we can efficiently compute.

In addition, we fix the following notations in the remaining sections:

- $f(X)$: the input polynomial in $\mathbb{F}_q[X]$ to be factorized,
- $\tilde{f}(X)$: an irreducible lifted polynomial of $f(X)$ in $A_0[X]$,
- F : the number field $K_0[X]/(\tilde{f}(X))$,
- L : the splitting field of \tilde{f} over K_0 ,
- G : the Galois group $\text{Gal}(L/K_0) = \text{Gal}(\tilde{f}/K_0)$,
- \mathfrak{Q}_0 : a fixed prime ideal of \mathcal{O}_L lying over p .

In this section, we reduce the problem of factoring f to computing an idempotent decomposition of $\bar{\mathcal{O}}_F$, generalizing the result in Section 3.3. For simplicity, we assume that \tilde{f} is a *monic* polynomial, and remove this assumption at the end of this section.

Ring homomorphisms τ and $\bar{\tau}$. Let $\alpha := X + (\tilde{f}(X)) \in F$, which is a root of \tilde{f} in F . As $\tilde{f}(X)$ is a monic polynomial in $A_0[X]$ and $A_0 \subseteq \mathcal{O}_{K_0}$, we have $\alpha \in \mathcal{O}_F$ (see [AM69, Corollary 5.4]).

Consider the natural inclusion $A_0[X]/(\tilde{f}(X)) = A_0[\alpha] \hookrightarrow \mathcal{O}_F$. Taking the quotients of both sides of this map mod p and identify $A_0/pA_0 = \mathbb{F}_p[Y]/(h(Y))$ with \mathbb{F}_q via ψ_0 , we obtain a ring homomorphism

$$\tau : \mathbb{F}_q[X]/(f(X)) \rightarrow \bar{\mathcal{O}}_F.$$

Let $g := \text{Rad}(f)$. Then the radical of $\mathbb{F}_q[X]/(f(X))$ is generated by $g(X) + (f(X))$. we obtain a ring homomorphism

$$\bar{\tau} : \mathbb{F}_q[X]/(g(X)) \rightarrow \bar{\mathcal{O}}_F/\text{Rad}(\bar{\mathcal{O}}_F),$$

which sends an element $h(X) + (g(X))$ to $\tau(h(X)) + \text{Rad}(\bar{\mathcal{O}}_F)$. Note that both $\mathbb{F}_q[X]/(g(X))$ and $\bar{\mathcal{O}}_F/\text{Rad}(\bar{\mathcal{O}}_F)$ are semisimple rings.

We can efficiently compute $\bar{\tau}$ by the following lemma.

Lemma 5.9. *There exists a polynomial-time algorithm that given f, \tilde{f}, F and the outputs of `ComputeRings` (see Lemma 5.7) on the input (F, p) , computes the \mathbb{F}_q -algebra $\mathbb{F}_q[X]/(g(X))$ (encoded in the standard \mathbb{F}_q -basis $\{1, X, \dots, X^{\deg(g)-1}\}$) and the map $\bar{\tau} : \mathbb{F}_q[X]/(g(X)) \rightarrow \bar{\mathcal{O}}_F/\text{Rad}(\bar{\mathcal{O}}_F)$.*

Proof. Compute g using Corollary 5.1 and form the \mathbb{F}_q -algebra $\mathbb{F}_q[X]/(g(X))$. To compute $\bar{\tau}$, we first compute $\alpha = X + (\tilde{f}(X)) \in F$ and $\bar{Y} := Y + (\tilde{h}(Y)) \in K_0 \subseteq F$. Then compute $\alpha + p\mathcal{O}_F, \bar{Y} + p\mathcal{O}_F \in \bar{\mathcal{O}}_F$ by identifying F with an ordinary number field (see Corollary 4.1) and running the algorithm `ComputeResidue` in Lemma 3.9 on $\alpha, \bar{Y} \in F$. Next, compute $\tau : \mathbb{F}_q[X]/(f(X)) \rightarrow \bar{\mathcal{O}}_F$ as the unique \mathbb{F}_p -linear map sending $X + (f(X))$ to $\alpha + p\mathcal{O}_F$ and $Y + (h(Y)) \in \mathbb{F}_p[Y]/(h(Y)) \cong \mathbb{F}_q$ to $\bar{Y} + p\mathcal{O}_F$. Finally compute $\bar{\tau}$ from τ by passing to the quotients modulo radicals using the given map $\bar{\mathcal{O}}_F \rightarrow \bar{\mathcal{O}}_F/\text{Rad}(\bar{\mathcal{O}}_F)$. \square

Extracting a factorization from an idempotent decomposition. We extract a factorization of f from an idempotent decomposition of R_F . This is achieved by the algorithm `ExtractFactorsV2` below (see Algorithm 10), extending the algorithm in Section 3.3.

The algorithm first computes $g = \text{Rad}(f)$, the ring $\mathbb{F}_q[X]/(g(X))$, and the map $\bar{\tau}$ at Line 1 using Lemma 5.9. It also maintains an idempotent decomposition I of the ring $\mathbb{F}_q[X]/(g(X))$ which initially only contains the unity.

The loop in Lines 3–8 enumerates idempotents $\delta' \in I_F$. For each δ' , we compute an ideal $J = \bar{\tau}^{-1}((1 - \delta')\bar{\mathcal{O}}_F/\text{Rad}(\bar{\mathcal{O}}_F))$ of $\mathbb{F}_q[X]/(g(X))$ and an element $\delta_0 \in J$ satisfying $(1 - \delta_0)J = \{0\}$ by solving systems of linear equations. As $\mathbb{F}_q[X]/(g(X))$ is semisimple, the element δ_0 is the unique idempotent of $\mathbb{F}_q[X]/(g(X))$ that generates J . And we use it to refine I .

The loop in Lines 9–12 extracts, for each idempotent $\delta \in I$, a monic factor g_δ of f . Furthermore, we compute a square-free factorization $g_\delta(X) = \prod_{i=1}^{k_\delta} g_{\delta,i}(X)$ for each factor g_δ . Finally, the algorithm returns the factorization

$$f(X) = \prod_{\delta \in I} \prod_{i=1}^{k_\delta} g_{\delta,i}(X).$$

Algorithm 10 ExtractFactorsV2

Input: f, \tilde{f}, F , the outputs of ComputeRings (see Lemma 5.7) on the input (F, p) , and an idempotent decomposition I_F of R_F

Output: factorization of f

- 1: compute $g = \text{Rad}(f)$, $\mathbb{F}_q[X]/(g(X))$ and $\bar{\tau} : \mathbb{F}_q[X]/(g(X)) \rightarrow \bar{\mathcal{O}}_F/\text{Rad}(\bar{\mathcal{O}}_F)$
- 2: $I \leftarrow \{1\}$, where 1 denotes the unity of $\mathbb{F}_q[X]/(g(X))$
- 3: **for** $\delta' \in I_F$ **do**
- 4: $J \leftarrow \bar{\tau}^{-1}((1 - \delta')\bar{\mathcal{O}}_F/\text{Rad}(\bar{\mathcal{O}}_F))$
- 5: compute $\delta_0 \in J$ satisfying $(1 - \delta_0)J = \{0\}$
- 6: **for** $\delta \in I$ satisfying $\delta_0\delta \notin \{0, \delta\}$ **do**
- 7: $I \leftarrow I - \{\delta\}$
- 8: $I \leftarrow I \cup \{\delta_0\delta, (1 - \delta_0)\delta\}$
- 9: **for** $\delta \in I$ **do**
- 10: compute nonzero $h_\delta(X) \in \mathbb{F}_q[X]$ of degree at most $\deg(g)$ lifting $1 - \delta$
- 11: $g_\delta(X) \leftarrow \text{gcd}(f(X), (h_\delta(X))^n) \quad \triangleright n = \deg(f)$
- 12: compute a square-free factorization $g_\delta(X) = \prod_{i=1}^{k_\delta} g_{\delta,i}(X)$
- 13: **return** the factorization $f(X) = \prod_{\delta \in I} \prod_{i=1}^{k_\delta} g_{\delta,i}(X)$

The following theorem is the main result of this section.

Theorem 5.6. *The algorithm ExtractFactorsV2 computes a factorization of f over \mathbb{F}_q in polynomial time, such that*

1. *the factorization is complete if I_F is a complete idempotent decomposition,*
2. *the factorization is proper if I_F is a proper idempotent decomposition, and*
3. *at least one factor in the factorization is irreducible over \mathbb{F}_q if I_F contains a primitive idempotent.*

Analysis of the algorithm. To prove Theorem 5.6, we introduce the following notations: let S (resp. S_F) denote the set of the maximal ideals of $\mathbb{F}_q[X]/(g(X))$ (resp. $\bar{\mathcal{O}}_F/\text{Rad}(\bar{\mathcal{O}}_F)$). For a maximal ideal \mathfrak{m} of $\bar{\mathcal{O}}_F/\text{Rad}(\bar{\mathcal{O}}_F)$, the preimage $\bar{\tau}^{-1}(\mathfrak{m})$ is a prime (and hence maximal) ideal of $\mathbb{F}_q[X]/(g(X))$. So we obtain a map

$$\pi : S_F \rightarrow S,$$

sending \mathfrak{m} to $\tau^{-1}(\mathfrak{m})$. It can be shown that π is surjective.¹⁰

Suppose $f(X) = \prod_{i=1}^k (f_i(X))^{m_i}$ where f_1, \dots, f_k are distinct monic irreducible factors of f over \mathbb{F}_q . For $i \in [k]$, let \mathfrak{m}_i be the (maximal) ideal of $\mathbb{F}_q[X]/(g(X))$ generated by $f_i(X) + (g(X))$. Then we have

$$S = \{\mathfrak{m}_1, \dots, \mathfrak{m}_k\} \quad \text{and} \quad g(X) = \prod_{i=1}^k f_i(X).$$

The proof of Theorem 5.6 is based on the following lemma.

Lemma 5.10. *Let I be the idempotent decomposition of $\mathbb{F}_q[X]/(g(X))$ given at the end of the algorithm `ExtractFactorsV2`. Define the partition P of S by*

$$P := \{B_\delta : \delta \in I\}, \quad \text{where} \quad B_\delta := \{\mathfrak{m} \in S : \delta \equiv 1 \pmod{\mathfrak{m}}\}$$

and the partition P' of S_F by

$$P' := \{B'_\delta : \delta \in I_F\}, \quad \text{where} \quad B'_\delta := \{\mathfrak{m} \in S_F : \delta \equiv 1 \pmod{\mathfrak{m}}\}.$$

Then P is the coarsest common refinement of the partitions $\{\pi(B), S - \pi(B)\}$, where B ranges over the blocks in P' . Moreover, for each $\delta \in I$, the polynomial g_δ in the algorithm is given by

$$g_\delta(X) = \prod_{i \in [k]: \mathfrak{m}_i \in B_\delta} (f_i(X))^{m_i}.$$

Proof. For the last claim, it suffices to prove, for all $i \in [k]$, that h_δ is divisible by f_i iff $\mathfrak{m}_i \in B_\delta$. By the choice of h_δ , it holds for all $i \in [k]$ that h_δ is divisible by f_i iff $1 - \delta \in \mathfrak{m}_i$. The claim then follows from the definition of B_δ .

For the first claim, it suffices to show that for every $\delta' \in I_F$ enumerated at Line 3 and δ_0 computed at Line 5 in the same iteration, it holds that $B_{\delta_0} \in \{\pi(B'_{\delta'}), S - \pi(B'_{\delta'})\}$. We claim that $B_{\delta_0} = S - \pi(B'_{\delta'})$. As the ideal J computed at Line 4 is generated by δ_0 , this claim is equivalent to $J = \bigcap_{\mathfrak{m} \in \pi(B'_{\delta'})} \mathfrak{m}$. Note that for $\mathfrak{m} \in S_F$, it holds that $1 - \delta' \in \mathfrak{m}$ iff $\mathfrak{m} \in B'_{\delta'}$ by the definition of $B'_{\delta'}$. So we have

$$(1 - \delta')\bar{\mathcal{O}}_F / \text{Rad}(\bar{\mathcal{O}}_F) = \bigcap_{\mathfrak{m} \in B'_{\delta'}} \mathfrak{m}$$

¹⁰To prove this, it suffices to show that any prime ideal of $A_0[X]/(\tilde{f}(X)) = A_0[\alpha] \subseteq \mathcal{O}_F$ is contained in a prime ideal of \mathcal{O}_F , which follows from [AM69, Theorem 5.10].

and hence

$$J = \bar{\tau}^{-1} \left(\bigcap_{\mathfrak{m} \in B'_{\delta'}} \mathfrak{m} \right) = \bigcap_{\mathfrak{m} \in B'_{\delta'}} \bar{\tau}^{-1}(\mathfrak{m}) = \bigcap_{\mathfrak{m} \in B'_{\delta'}} \pi(\mathfrak{m}) = \bigcap_{\mathfrak{m} \in \pi(B'_{\delta'})} \mathfrak{m}$$

as desired. \square

We also need the following lemma.

Lemma 5.11. $\pi : S_F \rightarrow S$ is bijective if f is square-free, i.e., $m_i = 1$ for $i \in [k]$.

Proof. Suppose $p\mathcal{O}_F$ splits into the product of prime ideals by

$$p\mathcal{O}_F = \prod_{i=1}^{\ell} \mathfrak{P}_i^{e(\mathfrak{P}_i)},$$

where $\mathfrak{P}_1, \dots, \mathfrak{P}_{\ell}$ are distinct prime ideals lying over p . For $j \in [\ell]$, let $\mathfrak{m}'_j := \frac{\mathfrak{P}_j/p\mathcal{O}_F}{\text{Rad}(\bar{\mathcal{O}}_F)}$. Then $S_F = \{\mathfrak{m}'_1, \dots, \mathfrak{m}'_{\ell}\}$. Let $n = \deg(f)$. Assume f is square-free. Then we have

$$\sum_{i=1}^k \deg(f_i) = \sum_{i=1}^k m_i \deg(f_i) = n = \sum_{j=1}^{\ell} e(\mathfrak{P}_j) f(\mathfrak{P}_j). \quad (5.1)$$

Fix $i \in [k]$. We know $\pi^{-1}(i) \neq \emptyset$ since π is surjective. Consider $j \in \pi^{-1}(i)$. As $\bar{\tau}(\mathfrak{m}_i) \subseteq \mathfrak{m}'_j$, the map $\bar{\tau} : \mathbb{F}_q[X]/(g(X)) \rightarrow \bar{\mathcal{O}}_F/\text{Rad}(\bar{\mathcal{O}}_F)$ induces a field embedding

$$\frac{\mathbb{F}_q[X]/(g(X))}{\mathfrak{m}_i} \hookrightarrow \frac{\bar{\mathcal{O}}_F/\text{Rad}(\bar{\mathcal{O}}_F)}{\mathfrak{m}'_j}.$$

The left hand side is isomorphic to $\mathbb{F}_q[X]/(f_i(X))$ whereas the right hand side is isomorphic to $\mathcal{O}_F/\mathfrak{P}_j = \kappa_{\mathfrak{P}_j}$. Therefore $\deg(f_i)$ divides $f(\mathfrak{P}_j)$.

Note that $e(\mathfrak{P}_j) \geq 1$ holds for all $j \in [\ell]$. It follows from (5.1) that in fact $e(\mathfrak{P}_j) = 1$ holds for all $j \in [\ell]$. Moreover, for all $i \in [k]$, the set $\pi^{-1}(i)$ contains only one element $j_i \in [\ell]$, and $\deg(f_i) = f(\mathfrak{P}_{j_i})$. In particular, the map π is bijective. \square

Now we are ready to prove Theorem 5.6.

Proof of Theorem 5.6. Polynomiality of the algorithm is straightforward. Suppose I_F is a complete idempotent decomposition of R_F . It is also a complete idempotent decomposition of $\bar{\mathcal{O}}_F/\text{Rad}(\bar{\mathcal{O}}_F)$ since the maximal ideals of $\bar{\mathcal{O}}_F/\text{Rad}(\bar{\mathcal{O}}_F)$ correspond one-to-one to those of R_F via $\mathfrak{m} \mapsto \mathfrak{m} \cap R_F$. So the partition P' in

Lemma 5.10 is ∞_{S_F} . By Lemma 5.10 and surjectivity of π , the partition P equals ∞_S , and the algorithm outputs the complete factorization $f(X) = \prod_{i \in [k]} (f_i(X))^{m_i}$.

Similarly, if I_F contains a primitive idempotent δ . Then P' contains a singleton B'_δ . By Lemma 5.10, the partition P contains a singleton $\pi(B'_\delta)$, and algorithm outputs a factorization of $f(X)$ in which the irreducible factors $f_i(X)$ appear m_i times, where i is the unique index in $[k]$ satisfying $\pi(B'_\delta) = \{m_i\}$.

Finally, suppose I_F is a proper idempotent decomposition of R_F , and hence a proper idempotent decomposition of $\bar{\mathcal{O}}_F/\text{Rad}(\bar{\mathcal{O}}_F)$. Then $P' \neq 0_{S_F}$. If π is bijective, then by Lemma 5.10, we have $P \neq 0_S$, and the algorithm outputs a proper factorization of f . Now suppose π is not bijective. Then f is not square-free by Lemma 5.11. As we compute a square-free factorization for each g_δ , the algorithm still outputs a proper factorization of f . \square

The reduction for non-monic polynomials. The same trick in Section 3.3 can be applied to make the above reduction work for a possibly non-monic polynomial \tilde{f} : let $c \in A_0$ be the leading coefficient of $\tilde{f}(X) \in A_0[X]$, and let $\bar{c} := \tilde{\psi}_0(c) \in \mathbb{F}_q^\times$. Compute the monic polynomials $\tilde{f}'(X) := c^{n-1} \cdot \tilde{f}(X/c) \in A_0[X]$ and $f'(X) := \bar{c}^{n-1} f(X/\bar{c}^{n-1}) \in \mathbb{F}_q[X]$. Run the algorithm `ExtractFactorsV2` on f' and \tilde{f}' instead of f and \tilde{f} , and obtain a factorization of f' . Finally, we recover a factorization of f from that of f' by substituting X with $\bar{c}X$ in each factor.

Remark. The reduction in this section exploits the well known connection between factorization of polynomials over finite fields and the splitting of prime ideals in number field extensions, which dates back to the classical work of Kummer and Dedekind (see, e.g., [Neu99, Proposition I.8.3]). The Kummer-Dedekind theorem, however, requires the map $\mathbb{F}_q[X]/(f(X)) \rightarrow \bar{\mathcal{O}}_F$ to be an isomorphism. For this reason, known factoring algorithms that use an irreducible lifted polynomial \tilde{f} often assume p is *regular* with respect to \tilde{f} . See, e.g., [Hua84; Hua91a; Hua91b; Rón92].¹¹ This assumption is *not* needed in our algorithm. The key observation is that we can always employ the surjective map π from the set of prime ideals of $\bar{\mathcal{O}}_F/\text{Rad}(\bar{\mathcal{O}}_F)$ to that of $\mathbb{F}_q[X]/(g(X))$, where $g = \text{Rad}(f)$. In algebro-geometric terminology, the map π is interpreted as the morphism of reduced affine schemes

$$\pi : \text{Spec}(\bar{\mathcal{O}}_F)_{\text{red}} \rightarrow \text{Spec}(A_0[\alpha]/pA_0[\alpha])_{\text{red}}$$

¹¹We say p is regular with respect to \tilde{f} if $pA_0[\alpha]$ is coprime to the *conductor* of $A_0[\alpha]$. See [Hua84] for the exact formulation of this condition. We remark that the journal version [Hua91a] (and [Hua91b; Rón92]) assumes the stronger condition that p is coprime to the discriminant of \tilde{f} .

induced from the morphism $\text{Spec } \mathcal{O}_F \rightarrow \text{Spec } A_0[\alpha]$. The latter morphism is known as the *normalization* of $\text{Spec } A_0[\alpha]$ (see [Har77, Exercise II.3.8]).

5.4 Producing a \mathcal{P} -scheme of double cosets \mathcal{C}

In this section, we present an algorithm that computes the idempotent decompositions of a collection of rings R_K corresponding to a \mathcal{P} -scheme of double cosets. It extends the algorithm in Section 3.4 and serves as (a preliminary version) of the main body of the generalized \mathcal{P} -scheme algorithm.

The pseudocode of the algorithm is given in Algorithm 11 below. Its input is a (K_0, \tilde{f}) -subfield system \mathcal{F} (see Definition 4.1). The algorithm outputs, for every $K \in \mathcal{F}$, an idempotent decomposition I_K of the ring R_K , together with some auxiliary data.

Algorithm 11 ComputeDoubleCosetPscheme

Input: (K_0, \tilde{f}) -subfield system \mathcal{F}

Output: for every $K \in \mathcal{F}$: the outputs of ComputeRings (see Lemma 5.7) on the input (K, p) , and an idempotent decomposition I_K of R_K

```

1: for  $K \in \mathcal{F}$  do
2:   call ComputeRings on  $(K, p)$ 
3:    $I_K \leftarrow \{1\}$ , where 1 denotes the unity of  $R_K$ 
4:   for  $(K, K') \in \mathcal{F}^2$  do
5:     call ComputeRelEmbeddings to compute all the embeddings from  $K$  to  $K'$ 
       over  $K_0$ 
6:     for embedding  $\phi : K \hookrightarrow K'$  over  $K_0$  do
7:       call ComputeRingHoms on  $p, K, K'$  and  $\phi$  to compute  $\bar{\phi}, \hat{\phi}$  and  $\tilde{\phi}$ 
8:   repeat
9:     call CompatibilityAndInvarianceTestV2
10:    call RegularityTestV2
11:    call StrongAntisymmetryTestV2
12:    call RamificationIndexTest
13:    call InertiaDegreeTest
14:  until  $I_K$  remains the same in the last iteration for all  $K \in \mathcal{F}$ 
15: return the outputs of ComputeRings on the input  $(K, p)$  and  $I_K$  for  $K \in \mathcal{F}$ 

```

We fix \mathcal{P} to be the subgroup system over $G = \text{Gal}(\tilde{f}/K_0)$ associated with \mathcal{F} , i.e.,

$$\mathcal{P} := \{H \subseteq G : L^H \cong_{K_0} K \text{ for some } K \in \mathcal{F}\}.$$

The first half (Lines 1–7) of the algorithm is the preprocessing stage: for each $K \in \mathcal{F}$, we run `ComputeRings` (see Lemma 5.7) on (K, p) which returns the following data:

- a p -maximal order \mathcal{O}'_K of K and the inclusion $\mathcal{O}'_K \hookrightarrow K$,
- $\bar{\mathcal{O}}_K$ and the quotient map $\mathcal{O}'_K \rightarrow \bar{\mathcal{O}}_K$,
- $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ and the quotient map $\bar{\mathcal{O}}_K \rightarrow \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$,
- R_K and the inclusion $R_K \hookrightarrow \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$.

For $(K, K') \in \mathcal{F}$, we also compute all the embeddings ϕ from K to K' and the corresponding ring homomorphisms $\bar{\phi} : \bar{\mathcal{O}}_K \rightarrow \bar{\mathcal{O}}_{K'}$, $\hat{\phi} : \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K) \rightarrow \bar{\mathcal{O}}_{K'}/\text{Rad}(\bar{\mathcal{O}}_{K'})$ and $\tilde{\phi} : R_K \rightarrow R_{K'}$. Moreover, for each $K \in \mathcal{F}$, we initialize the idempotent decomposition I_K of R_K to be the trivial one containing only the unity of R_K .

The second half (Lines 8–14) of the algorithm refines the idempotent decompositions I_K for $K \in \mathcal{F}$. To analyze it, we associate a \mathcal{P} -collection \mathcal{C} of double cosets with these idempotent decompositions. For each $H \in \mathcal{P}$, define a partition C_H of the coset space $H \backslash G / \mathcal{D}_{\Omega_0}$ as follows: Let K be the unique field in \mathcal{F} isomorphic to L^H over K_0 . Fix an isomorphism $\tau_H : K \rightarrow L^H$ over K_0 , which induces a ring isomorphism $\tilde{\tau}_H : R_K \rightarrow R_{L^H}$. Define $I_H := \tilde{\tau}_H(I_K)$, which is an idempotent decomposition of R_{L^H} . By Definition 5.4, it corresponds to a partition $P(I_H)$ of $H \backslash G / \mathcal{D}_{\Omega_0}$.¹² And we define

$$C_H := P(I_H).$$

Finally, define

$$\mathcal{C} := \{C_H : H \in \mathcal{P}\},$$

which is a \mathcal{P} -collection of double cosets (with respect to \mathcal{D}_{Ω_0}).

The subroutines in Lines 9–11 extend those in Section 3.5, 3.6, and 3.7 respectively:

Lemma 5.12. *There exists a subroutine `CompatibilityAndInvarianceTestV2` that updates I_K in time polynomial in $\log p$ and the size of \mathcal{F} so that the partitions $C_H \in \mathcal{C}$ are refined. Moreover, at least one partition C_H is properly refined if \mathcal{C} is not compatible or invariant.*

¹²Definition 5.4 is made with respect to a fixed prime ideal Ω_0 of \mathcal{O}_L lying over p . This ideal is chosen at the beginning of Section 5.3.

Lemma 5.13. *There exists a subroutine `RegularityTestV2` that updates I_K in time polynomial in $\log p$ and the size of \mathcal{F} so that the partitions $C_H \in \mathcal{C}$ are refined. Moreover, at least one partition C_H is properly refined if \mathcal{C} is compatible but not regular.*

Lemma 5.14. *Under GRH, there exists a subroutine `StrongAntisymmetryTestV2` that updates I_K in time polynomial in $\log p$ and the size of \mathcal{F} so that the partitions $C_H \in \mathcal{C}$ are refined. Moreover, at least one partition C_H is properly refined if \mathcal{C} is a \mathcal{P} -scheme of double cosets, but not strongly antisymmetric.*

The proofs of Lemma 5.12–5.14 (and the corresponding subroutines) are almost the same as those of Lemma 3.13–3.15 in Chapter 3. For this reason, we only list the changes that need to be made rather than describe the complete proofs and the subroutines.

Proof sketch of Lemma 5.12–5.14. We make the following changes to the proofs of Lemma 3.13–3.15 and the corresponding subroutines:

each quotient ring $\bar{\mathcal{O}}_K$ is replaced with the ring R_K , which is still isomorphic to a finite product of copies of \mathbb{F}_p . A maximal ideal \mathfrak{P} of $\bar{\mathcal{O}}_K$ is replaced with the maximal ideal $(\mathfrak{P}/\text{Rad}(\bar{\mathcal{O}}_K)) \cap R_K$ of R_K . The subroutines enumerate field embeddings over K_0 instead of arbitrary field embeddings. For each field embedding $\phi : K \rightarrow K'$ over K_0 , we use the ring homomorphism $\tilde{\phi} : R_K \rightarrow R_{K'}$ in place of $\bar{\phi} : \bar{\mathcal{O}}_K \rightarrow \bar{\mathcal{O}}_{K'}$. The ring isomorphisms $\bar{\tau}_H : \bar{\mathcal{O}}_K \rightarrow \bar{\mathcal{O}}_{L^H}$ are replaced with $\tilde{\tau}_H : R_K \rightarrow R_{L^H}$.

A right coset Hg is replaced with a double coset $Hg\mathcal{D}_{\Omega_0}$, and a right coset space $H \backslash G$ is replaced with $H \backslash G / \mathcal{D}_{\Omega_0}$. A projection $\pi_{H,H'}$ is replaced with $\pi_{H,H'}^{\mathcal{D}_{\Omega_0}}$, and a conjugation $c_{H,g}$ is replaced with $c_{H,g}^{\mathcal{D}_{\Omega_0}}$ (see Definition 5.5).

Finally, instead of applying Corollary 3.1, Lemma 3.5, and Lemma 3.6 from Chapter 3, we apply Lemma 5.3, Lemma 5.4, and Lemma 5.5, respectively. The details are left to the reader. \square

In addition, the subroutines at Line 12 and Line 13 properly refine the partitions in \mathcal{C} unless they all have locally constant ramification indices and inertia degrees:

Lemma 5.15. *There exists a subroutine `RamificationIndexTest` that updates I_K in time polynomial in $\log p$ and the size of \mathcal{F} so that the partitions $C_H \in \mathcal{C}$ are refined.*

Moreover, at least one partition C_H is properly refined unless all the partitions in \mathcal{C} have locally constant ramification indices (with respect to $(\mathcal{D}_{\Omega_0}, \mathcal{I}_{\Omega_0})$).

Lemma 5.16. *There exists a subroutine `InertiaDegreeTest` that updates I_K in time polynomial in $\log p$ and the size of \mathcal{F} so that the partitions $C_H \in \mathcal{C}$ are refined. Moreover, at least one partition C_H is properly refined unless all the partitions in \mathcal{C} have locally constant inertia degrees (with respect to $(\mathcal{D}_{\Omega_0}, \mathcal{I}_{\Omega_0})$).*

Lemma 5.15 and Lemma 5.16 are proved in Section 5.5 and Section 5.6, respectively.

Combining Lemma 5.12–5.16 yields the main result of this section:

Theorem 5.7 (Theorem 5.2 restated). *Under the assumption of GRH, the algorithm `ComputeDoubleCosetPscheme` runs in time polynomial in $\log p$ and the size of \mathcal{F} , and when it terminates, \mathcal{C} is a strongly antisymmetric \mathcal{P} -scheme of double cosets (with respect to \mathcal{D}_{Ω_0}). Moreover, all the partitions in \mathcal{C} have locally constant ramification indices and inertia degrees (with respect to $(\mathcal{D}_{\Omega_0}, \mathcal{I}_{\Omega_0})$).*

5.5 Testing local constantness of ramification indices

In this section, we describe the subroutine `RamificationIndexTest` that properly refines at least one partition in \mathcal{C} unless all the partition have locally constant ramification indices.

Algorithm 12 `RamificationIndexTest`

```

1: for  $K \in \mathcal{F}$  do
2:   for  $i \leftarrow 1$  to  $[K : K_0]$  do
3:      $J \leftarrow$  the image of  $\text{Ann}_{\bar{\mathcal{O}}_K}(\text{Rad}(\bar{\mathcal{O}}_K)^i) \subseteq \bar{\mathcal{O}}_K$  in  $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ 
4:     find  $\delta_0 \in J \cap R_K$  satisfying  $(1 - \delta_0)J = \{0\}$ 
5:     for  $\delta \in I_K$  satisfying  $\delta_0\delta \notin \{0, \delta\}$  do
6:        $I_K \leftarrow I_K - \{\delta\}$ 
7:        $I_K \leftarrow I_K \cup \{\delta_0\delta, (1 - \delta_0)\delta\}$ 

```

The pseudocode of the subroutine is given in Algorithm 12 above. We enumerate $K \in \mathcal{F}$ and $i = 1, 2, \dots, [K : K_0]$. For each K and i , we compute an ideal J of $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$, defined to be the image of $\text{Ann}_{\bar{\mathcal{O}}_K}(\text{Rad}(\bar{\mathcal{O}}_K)^i)$ under the quotient map $\bar{\mathcal{O}}_K \rightarrow \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$. We also compute an element $\delta_0 \in J \cap R_K \subseteq R_K$, satisfying $(1 - \delta_0)J = \{0\}$. As $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ and R_K are semisimple, and $\mathfrak{m} \mapsto \mathfrak{m} \cap R_K$ is a one-to-one correspondence between the maximal ideals of $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ and

those of R_K , we know δ_0 is the unique idempotent of $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ (resp. R_K) that generates J (resp. $J \cap R_K$). Then we use δ_0 to refine I_K .

Next we prove Lemma 5.15.

Proof of Lemma 5.15. The claim about the running time is straightforward. Suppose there exists $H \in \mathcal{P}$ such that C_H does not have locally constant ramification indices. Choose $B \in C_H$ and $g, g' \in G$ such that $Hg\mathcal{D}_{\Omega_0}, Hg'\mathcal{D}_{\Omega_0} \in B$ and $e(Hg\mathcal{D}_{\Omega_0}) < e(Hg'\mathcal{D}_{\Omega_0})$.

By Theorem 5.3 and Definition 5.2, the ideal $p\mathcal{O}_{L^H}$ splits into the product of prime ideals ${}^h\mathfrak{Q}_0 \cap \mathcal{O}_{L^H}$ by

$$p\mathcal{O}_{L^H} = \prod_{Hh\mathcal{D}_{\Omega_0} \in H \backslash G / \mathcal{D}_{\Omega_0}} ({}^h\mathfrak{Q}_0 \cap \mathcal{O}_{L^H})^{e(Hh\mathcal{D}_{\Omega_0})}.$$

For $Hh\mathcal{D}_{\Omega_0} \in H \backslash G / \mathcal{D}_{\Omega_0}$, define $\mathfrak{P}_{Hh\mathcal{D}_{\Omega_0}} := ({}^h\mathfrak{Q}_0 \cap \mathcal{O}_{L^H}) / p\mathcal{O}_{L^H}$, which is a maximal ideal of $\bar{\mathcal{O}}_{L^H}$. By the Chinese remainder theorem, we have

$$\bar{\mathcal{O}}_{L^H} \cong \prod_{x \in H \backslash G / \mathcal{D}_{\Omega_0}} \bar{\mathcal{O}}_{L^H} / \mathfrak{P}_x^{e(x)}.$$

And $\text{Rad}(\bar{\mathcal{O}}_{L^H}) = \prod_{x \in H \backslash G / \mathcal{D}_{\Omega_0}} \mathfrak{P}_x$. So for $i \in \mathbb{N}$, we have

$$\text{Ann}_{\bar{\mathcal{O}}_{L^H}}(\text{Rad}(\bar{\mathcal{O}}_{L^H})^i) = \prod_{x \in H \backslash G / \mathcal{D}_{\Omega_0}} \mathfrak{P}_x^{\max\{0, e(x)-i\}}. \quad (5.2)$$

Choose $i = e(Hg\mathcal{D}_{\Omega_0})$ and let J be the image of $\text{Ann}_{\bar{\mathcal{O}}_{L^H}}(\text{Rad}(\bar{\mathcal{O}}_{L^H})^i)$ in the quotient ring $\bar{\mathcal{O}}_{L^H} / \text{Rad}(\bar{\mathcal{O}}_{L^H})$. Let δ_0 be the unique idempotent of $\bar{\mathcal{O}}_{L^H} / \text{Rad}(\bar{\mathcal{O}}_{L^H})$ that generates J . It follows from (5.2) that

$$\delta_0 \equiv 1 \pmod{\mathfrak{P}_{Hg\mathcal{D}_{\Omega_0}} / \text{Rad}(\bar{\mathcal{O}}_{L^H})} \quad \text{and} \quad \delta_0 \equiv 0 \pmod{\mathfrak{P}_{Hg'\mathcal{D}_{\Omega_0}} / \text{Rad}(\bar{\mathcal{O}}_{L^H})}.$$

Therefore

$$g^{-1}(i_{L^H, L}(\delta_0)) \equiv 1 \pmod{\bar{\mathfrak{Q}}_0} \quad \text{and} \quad g'^{-1}(i_{L^H, L}(\delta_0)) \equiv 0 \pmod{\bar{\mathfrak{Q}}_0}, \quad (5.3)$$

where $i_{L^H, L} : R_{L^H} \hookrightarrow R_L$ is the inclusion induced from the natural inclusion $\mathcal{O}_{L^H} \hookrightarrow \mathcal{O}_L$.

On the other hand, by Lemma 5.4, the block $B \in C_H$ corresponds to an idempotent $\delta = \delta_B \in I_H$. And

$$h^{-1}(i_{L^H, L}(\delta)) \equiv 1 \pmod{\bar{\mathfrak{Q}}_0}$$

holds for all $h \in G$ satisfying $Hh\mathcal{D}_{\Omega_0} \in B$. In particular, it holds for $h = g$ and $h = g'$. It follows from (5.3) that $\delta_0\delta \notin \{0, \delta\}$.

Identifying L^H with a field $K \in \mathcal{F}$ using the isomorphism $\tau_H : K \rightarrow L^H$ over K_0 chosen in Section 5.4, we see that the subroutine is guaranteed to find an idempotent $\delta \in I_K$ satisfying $\delta_0\delta \notin \{0, \delta\}$ at Line 5. The lemma follows. \square

5.6 Testing local constantness of inertia degrees

In this section, we describe the subroutine `InertiaDegreeTest` that properly refines at least one partition in \mathcal{C} unless all the partition have locally constant inertia degrees.

Algorithm 13 `InertiaDegreeTest`

```

1: for  $K \in \mathcal{F}$  do
2:   for  $i \leftarrow 1$  to  $[K : K_0]$  do
3:      $J \leftarrow$  the ideal of  $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$  generated by  $\{x^{p^i} - x : x \in \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)\}$ 
4:     find  $\delta_0 \in J \cap R_K$  satisfying  $(1 - \delta_0)J = \{0\}$ 
5:     for  $\delta \in I_K$  satisfying  $\delta_0\delta \notin \{0, \delta\}$  do
6:        $I_K \leftarrow I_K - \{\delta\}$ 
7:        $I_K \leftarrow I_K \cup \{\delta_0\delta, (1 - \delta_0)\delta\}$ 

```

The pseudocode of the subroutine is given in Algorithm 13. We enumerate $K \in \mathcal{F}$ and $i = 1, 2, \dots, [K : K_0]$. For each K and i , we compute an ideal J of $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$, generated by the elements $x^{p^i} - x$, where x ranges over $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$. Note that J is just the \mathbb{F}_p -linear subspace of $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ spanned by $x^{p^i} - x$ where x ranges over an \mathbb{F}_p -basis of $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$. So it can be efficiently computed. We also compute an element $\delta_0 \in J \cap R_K \subseteq R_K$ satisfying $(1 - \delta_0)J = \{0\}$. As in Algorithm 12, it is the unique idempotent of $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ (resp. R_K) that generates J (resp. $J \cap R_K$). Then we use δ_0 to refine I_K .

Next we prove Lemma 5.16.

Proof of Lemma 5.16. The claim about the running time is straightforward. Suppose there exists $H \in \mathcal{P}$ such that C_H does not have locally constant inertia degrees. Choose $B \in C_H$ and $g, g' \in G$ such that $Hg\mathcal{D}_{\Omega_0}, Hg'\mathcal{D}_{\Omega_0} \in B$ and $f(Hg\mathcal{D}_{\Omega_0}) > f(Hg'\mathcal{D}_{\Omega_0})$.

For $Hh\mathcal{D}_{\Omega_0} \in H \setminus G/\mathcal{D}_{\Omega_0}$, define $\mathfrak{P}_{Hh\mathcal{D}_{\Omega_0}} := ({}^h\mathfrak{Q}_0 \cap \mathcal{O}_{L^H})/p\mathcal{O}_{L^H}$, which is a maximal ideal of $\bar{\mathcal{O}}_{L^H}$. By Theorem 5.3, Definition 5.2, and the Chinese remainder

theorem, we have

$$\bar{\mathcal{O}}_{LH}/\text{Rad}(\bar{\mathcal{O}}_{LH}) \cong \prod_{Hh\mathcal{D}_{\Omega_0} \in H \setminus G/\mathcal{D}_{\Omega_0}} \bar{\mathcal{O}}_{LH}/\mathfrak{P}_{Hh\mathcal{D}_{\Omega_0}}$$

and each factor $\bar{\mathcal{O}}_{LH}/\mathfrak{P}_{Hh\mathcal{D}_{\Omega_0}}$ is an extension field of \mathbb{F}_p of degree $f(Hh\mathcal{D}_{\Omega_0})$. Choose $i = f(Hg'\mathcal{D}_{\Omega_0})$ and let J be the ideal of $\bar{\mathcal{O}}_{LH}/\text{Rad}(\bar{\mathcal{O}}_{LH})$ generated by $x^{p^i} - x$ where x ranges over $\bar{\mathcal{O}}_{LH}/\text{Rad}(\bar{\mathcal{O}}_{LH})$. Let δ_0 be the unique idempotent of $\bar{\mathcal{O}}_{LH}/\text{Rad}(\bar{\mathcal{O}}_{LH})$ that generates J . Note that we have

$$x^{p^i} \not\equiv x \pmod{\mathfrak{P}_{Hg\mathcal{D}_{\Omega_0}}/\text{Rad}(\bar{\mathcal{O}}_{LH})} \quad \text{for some } x \in \bar{\mathcal{O}}_{LH}/\text{Rad}(\bar{\mathcal{O}}_{LH}),$$

and

$$x^{p^i} \equiv x \pmod{\mathfrak{P}_{Hg'\mathcal{D}_{\Omega_0}}/\text{Rad}(\bar{\mathcal{O}}_{LH})} \quad \text{for all } x \in \bar{\mathcal{O}}_{LH}/\text{Rad}(\bar{\mathcal{O}}_{LH}).$$

So J is contained in $\mathfrak{P}_{Hg'\mathcal{D}_{\Omega_0}}$ but not in $\mathfrak{P}_{Hg\mathcal{D}_{\Omega_0}}$. It follows that

$$\delta_0 \equiv 1 \pmod{\mathfrak{P}_{Hg\mathcal{D}_{\Omega_0}}/\text{Rad}(\bar{\mathcal{O}}_{LH})} \quad \text{and} \quad \delta_0 \equiv 0 \pmod{\mathfrak{P}_{Hg'\mathcal{D}_{\Omega_0}}/\text{Rad}(\bar{\mathcal{O}}_{LH})}.$$

Then (5.3) in the proof of Lemma 5.15 holds. The rest of the proof follows the proof of Lemma 5.15. \square

5.7 A \mathcal{P} -collection $\tilde{\mathcal{C}}$ induced from \mathcal{C} and auxiliary elements

The idempotent decompositions I_K produced in Section 5.4 define a \mathcal{P} -scheme of double cosets \mathcal{C} rather than an (ordinary) \mathcal{P} -scheme. Section 5.7–5.9 are devoted to turning it to a \mathcal{P} -scheme $\tilde{\mathcal{C}}$. In particular, this section focuses on the definition of $\tilde{\mathcal{C}}$ as a \mathcal{P} -collection.

We assume $p > \deg(f)$ in Section 5.7–5.9. As mentioned in Section 5.1, this assumption implies that the wild inertia group $\mathcal{W}_{\Omega_0} \subseteq G$ of Ω_0 over K_0 is trivial.

Suppose the partitions in \mathcal{C} all have locally constant ramification indices and inertia degrees (with respect to $(\mathcal{D}_{\Omega_0}, \mathcal{I}_{\Omega_0})$). Then for $K \in \mathcal{F}$ and $\delta \in I_K$, the (nonempty) set of maximal ideals \mathfrak{P} of \mathcal{O}_K satisfying

$$\delta \equiv 1 \pmod{\bar{\mathfrak{P}}} \quad \text{where} \quad \bar{\mathfrak{P}} := \frac{\mathfrak{P}/p\mathcal{O}_K}{\text{Rad}(\bar{\mathcal{O}}_K)} \cap R_K$$

all have the same ramification index $e(\mathfrak{P})$ and the same inertia degree $f(\mathfrak{P})$. We denote $e(\mathfrak{P})$ by e_δ and $f(\mathfrak{P})$ by f_δ . Note that e_δ and f_δ are coprime to p by Theorem 5.3 and the assumption $p > \deg(f)$.

Recall that for a finite extension K of K_0 and $i \in \mathbb{N}^+$, we denote by $A_{K,i}$ the ring $(\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)) \otimes_{\mathbb{F}_q} \mathbb{F}_{q^i}$. To define $\tilde{\mathcal{C}}$, we need an auxiliary collection of elements in rings $\bar{\mathcal{O}}_K$ or $A_{K,i}$. We call such a collection of elements an \mathcal{I} -*advice*:

Definition 5.7. Suppose $\mathcal{I} = \{I_K : K \in \mathcal{F}\}$ is a collection of idempotent decompositions of the rings R_K , $K \in \mathcal{F}$, that defines to a \mathcal{P} -collection of double cosets \mathcal{C} (with respect to \mathcal{D}_{Ω_0}), such that all the partitions in \mathcal{C} have locally constant ramification indices and inertia degrees (with respect to $(\mathcal{D}_{\Omega_0}, \mathcal{I}_{\Omega_0})$). An \mathcal{I} -advice $\{\mathcal{S}, \mathcal{T}\}$ consists of the following data:

- $\mathcal{S} = \{s_\delta : \delta \in I_K, e_\delta > 1\}$, where each $s_\delta \in \mathcal{S}$ is an element of $\bar{\mathcal{O}}_K$ such that $s_\delta \in \mathfrak{m} - \mathfrak{m}^2$ for all the maximal ideals \mathfrak{m} of $\bar{\mathcal{O}}_K$ satisfying $\delta \equiv 1 \pmod{\mathfrak{m}/\text{Rad}(\bar{\mathcal{O}}_K)}$.
- $\mathcal{T} = \{t_\delta : \delta \in I_K, f_\delta > 1\}$, where each $t_\delta \in \mathcal{T}$ is an element of A_{K, f_δ} such that $t_\delta \notin \mathfrak{m}$ for all the maximal ideals \mathfrak{m} of A_{K, f_δ} satisfying $\delta \equiv 1 \pmod{\mathfrak{m}}$, and $\sigma_{K, f_\delta}(t_\delta) = \xi \cdot t_\delta$, where $\xi \in \mathbb{F}_{q^{f_\delta}}$ is a primitive f_δ th root of unity.¹³

An \mathcal{I} -advice can be computed from \mathcal{I} by the following lemma. Its proof is deferred to Appendix C.

Lemma 5.17. Under GRH, there exists a subroutine `ComputeAdvice` that given $\mathcal{I} = \{I_K : K \in \mathcal{F}\}$ as in Definition 5.7, either properly refines some idempotent decomposition $I_K \in \mathcal{I}$, or computes e_δ, f_δ for $K \in \mathcal{F}$, $\delta \in I_K$ and an \mathcal{I} -advice.¹⁴ Moreover, the subroutine runs in time polynomial in $\log p$ and the size of \mathcal{F} .

We also need the following notations: recall that for $H \in \mathcal{P}$, we chose an isomorphism $\tau_H : K \rightarrow L^H$ over K_0 where K is the unique field in \mathcal{F} isomorphic to L^H over K_0 . The induced isomorphism $\bar{\mathcal{O}}_K \cong \bar{\mathcal{O}}_{L^H}$ identifies each $s_\delta \in \mathcal{S}$ (where \mathcal{S} is as in Definition 5.7) with an element in $\bar{\mathcal{O}}_{L^H}$, which we denote by $s_{\delta, H}$. Similarly, we identify each $t_\delta \in \mathcal{T}$ with an element in A_{L^H, f_δ} , denoted by $t_{\delta, H}$.

Next we define a \mathcal{P} -collection $\tilde{\mathcal{C}}$ using \mathcal{I} and an \mathcal{I} -advice:

Definition 5.8. Let $\mathcal{I} = \{I_K : K \in \mathcal{F}\}$ be as in Definition 5.7 and $\{\mathcal{S}, \mathcal{T}\}$ be an \mathcal{I} -advice. Let $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ be the \mathcal{P} -collection of double cosets with respect to \mathcal{D}_{Ω_0} associated with \mathcal{I} (see Section 5.4). For $H \in \mathcal{P}$, let K be the unique field in \mathcal{F} isomorphic to L^H over K_0 , and define the partition \tilde{C}_H of $H \setminus G$ so that $Hg, Hg' \in H \setminus G$ are in the same block of \tilde{C}_H iff the following conditions are satisfied:

¹³We regard $\delta \in R_K \subseteq \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ as an element of A_{K, f_δ} via $\delta \mapsto \delta \otimes 1$, and $\xi \in \mathbb{F}_{q^{f_\delta}}$ as an element of A_{K, f_δ} via $\xi \mapsto 1 \otimes \xi$.

¹⁴We need to compute the rings A_{K, f_δ} before computing the elements $t_\delta \in A_{K, f_\delta}$. These rings will be computed before the call of the subroutine `ComputeAdvice`. See Section 5.9.

1. $Hg\mathcal{D}_{\Omega_0}$ and $Hg'\mathcal{D}_{\Omega_0}$ are in the same block B of C_H .
2. Let δ be the unique idempotent in I_K such that $\tilde{\tau}_H(\delta) = \delta_B$ (see Definition 5.4), where $B \in C_H$ is as in the previous condition. If $e_\delta > 1$, the order of the unique element c in $\kappa_{\Omega_0}^\times$ satisfying

$$g^{-1}s_{\delta,H} + I = c \cdot (g'^{-1}s_{\delta,H} + I)$$

is coprime to e_δ , where $I = (\mathfrak{O}_{\Omega_0}/p\mathfrak{O}_L)^{e(\mathfrak{O}_{\Omega_0})/e_\delta+1}$.

3. Let $\delta \in I_K$ be as in the previous condition. Let \mathfrak{m}_0 be an arbitrary maximal ideal of A_{L,f_δ} containing $\frac{\mathfrak{O}_0/p\mathfrak{O}_L}{\text{Rad}(\mathfrak{O}_L)}$. If $f_\delta > 1$, the order of the unique element c in $(A_{L,f_\delta}/\mathfrak{m}_0)^\times$ satisfying

$$g^{-1}t_{\delta,H} + \mathfrak{m}_0 = c \cdot (g'^{-1}t_{\delta,H} + \mathfrak{m}_0)$$

is coprime to f_δ .

Define $\tilde{\mathcal{C}} = \{\tilde{C}_H : H \in \mathcal{P}\}$, which is a \mathcal{P} -collection. We say $\tilde{\mathcal{C}}$ is the \mathcal{P} -collection associated with \mathcal{I} and $\{\mathcal{S}, \mathcal{T}\}$.

We check that $\tilde{\mathcal{C}}$ is well defined:

Lemma 5.18. *The \mathcal{P} -collection $\tilde{\mathcal{C}}$ in Definition 5.8 is well defined.*

The proof of Lemma 5.18 is routine and can be found in Appendix C.

5.8 $(\mathcal{C}, \mathcal{D})$ -separated \mathcal{P} -collections

We continue the discussion in the previous section. Our goal is to compute $\mathcal{I} = \{I_K : K \in \mathcal{F}\}$ and an \mathcal{I} -advice $\{\mathcal{S}, \mathcal{T}\}$ such that the associated \mathcal{P} -collection $\tilde{\mathcal{C}}$ is a strongly antisymmetric \mathcal{P} -scheme. To achieve this goal, we introduce another property of \mathcal{P} -collections called $(\mathcal{C}, \mathcal{D})$ -separatedness:

Definition 5.9. *Let \mathcal{P} be a subgroup system over a finite group G , and let $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ be a \mathcal{P} -collection of double cosets with respect to a subgroup \mathcal{D} of G . We say a \mathcal{P} -collection $\tilde{\mathcal{C}} = \{\tilde{C}_H : H \in \mathcal{P}\}$ is $(\mathcal{C}, \mathcal{D})$ -separated if the following conditions are satisfied:*

1. All the partitions $\tilde{C}_H \in \tilde{\mathcal{C}}$ are invariant under the action of \mathcal{D} by inverse right translation, i.e. for all $B \in \tilde{C}_H$ and $g \in \mathcal{D}$, the set ${}^gB = \{Hhg^{-1} : Hh \in B\}$ is also in \tilde{C}_H .

2. For $H \in \mathcal{P}$, the map $\pi_H : H \backslash G \rightarrow H \backslash G / \mathcal{D}$ sending $Hg \in H \backslash G$ to $Hg\mathcal{D}$ maps each block of $\tilde{\mathcal{C}}_H$ bijectively to a block of C_H .

It is worth noting that if $\tilde{\mathcal{C}}$ is $(\mathcal{C}, \mathcal{D})$ -separated, then all the partitions in \mathcal{C} automatically have locally constant ramification indices and inertia degrees:

Lemma 5.19. *Suppose $\tilde{\mathcal{C}} = \{\tilde{\mathcal{C}}_H : H \in \mathcal{P}\}$ is a $(\mathcal{C}, \mathcal{D})$ -separated \mathcal{P} -collection where $\mathcal{P}, \mathcal{C}, \mathcal{D}$ are as in Definition 5.9. Let \mathcal{I} be a normal subgroup of \mathcal{D} . Then all the partitions in \mathcal{C} have locally constant ramification indices and inertia degrees with respect to $(\mathcal{D}, \mathcal{I})$.*

Proof. Fix $H \in \mathcal{P}$, $B \in C_H$, and $\tilde{B} \in \tilde{\mathcal{C}}_H$ such that $\pi_H(\tilde{B}) = B$, where π_H is as in Definition 5.9. Let \mathcal{D}' be a subgroup of \mathcal{D} . Consider arbitrary $Hg\mathcal{D}', Hg'\mathcal{D}' \in B$ and lift them to $Hg, Hg' \in \tilde{B}$ respectively. Choose $h_1, \dots, h_k \in \mathcal{D}'$ such that the \mathcal{D}' -orbit of Hg is $\{Hgh_1, \dots, Hgh_k\}$ and the cosets Hgh_i are all distinct. We claim $Hg'h_1, \dots, Hg'h_k$ are also distinct. Assume to the contrary that $Hg'h_{i_1} = Hg'h_{i_2}$ holds for distinct $i_1, i_2 \in [k]$. Then $Hg'h_{i_1}$ and $Hg'h_{i_2}$ are in the same block of $\tilde{\mathcal{C}}_H$. It follows by the first condition in Definition 5.9 that Hgh_{i_1} and Hgh_{i_2} are also in the same block. But $Hgh_{i_1} \neq Hgh_{i_2}$ and they are both mapped to $Hg\mathcal{D}$ by π_H , contradicting the second condition in Definition 5.9. This proves the claim. So the cardinality of the \mathcal{D}' -orbit of any $Hg \in H \backslash G$ only depends on the block in C_H containing $Hg\mathcal{D}$. In particular, this holds for $\mathcal{D}' = \mathcal{D}$ and $\mathcal{D}' = \mathcal{I}$. The lemma then follows from Definition 5.2. \square

The following lemma provides a criterion for a $(\mathcal{C}, \mathcal{D})$ -separated \mathcal{P} -collection to be a strongly antisymmetric \mathcal{P} -scheme.

Lemma 5.20. *Let \mathcal{P} be a subgroup system over a finite group G , and let $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ be a \mathcal{P} -scheme of double cosets with respect to $\mathcal{D} \subseteq G$. Suppose $\tilde{\mathcal{C}} = \{\tilde{\mathcal{C}}_H : H \in \mathcal{P}\}$ is a compatible, invariant, $(\mathcal{C}, \mathcal{D})$ -separated \mathcal{P} -collection. Then it is actually a \mathcal{P} -scheme. Moreover, if \mathcal{C} is antisymmetric (resp. strongly antisymmetric), so is $\tilde{\mathcal{C}}$.*

Proof. For the first claim, we just need to show $\tilde{\mathcal{C}}$ is regular. Consider $H, H' \in \mathcal{P}$ with $H \subseteq H'$. Let $\pi_H : H \backslash G \rightarrow H \backslash G / \mathcal{D}$ be the map sending $Hg \in H \backslash G$ to $Hg\mathcal{D}$,

and define $\pi_{H'}$ similarly. Then the following diagram commutes.

$$\begin{array}{ccc} H \backslash G & \xrightarrow{\pi_{H,H'}} & H' \backslash G \\ \pi_H \downarrow & & \downarrow \pi_{H'} \\ H \backslash G / \mathcal{D} & \xrightarrow{\pi_{H,H'}^{\mathcal{D}}} & H' \backslash G / \mathcal{D} \end{array}$$

For $B \in \tilde{\mathcal{C}}_H$ and $B' \in \tilde{\mathcal{C}}_{H'}$ containing $\pi_{H,H'}(B)$, we need to show the map $\pi_{H,H'}|_B : B \rightarrow B'$ has the constant degree, i.e., the cardinality of $\pi_{H,H'}^{-1}(y) \cap B$ is independent of $y \in B'$. As $\tilde{\mathcal{C}}$ is $(\mathcal{C}, \mathcal{D})$ -separated, the map π_H sends B bijectively to $\pi_H(B) \in C_H$, and similarly $\pi_{H'}$ sends B' bijectively to $\pi_{H'}(B') \in C_{H'}$. The claim then follows from regularity of \mathcal{C} .

Note that the conjugations also commute with the maps π_H , i.e., $\pi_{hHh^{-1}} \circ c_{H,h} = c_{H,h}^{\mathcal{D}} \circ \pi_H$ for $H \in \mathcal{P}$ and $h \in G$. Assume $\tilde{\mathcal{C}}$ is not strongly antisymmetric. Then there exists a nontrivial permutation τ of a block $B \in \tilde{\mathcal{C}}_H$ for some $H \in \mathcal{P}$ that arises as a composition of maps $\sigma_i : B_{i-1} \rightarrow B_i$, $i = 1 \dots, k$ where B_i is a block of $\tilde{\mathcal{C}}_{H_i}$, $H_i \in \mathcal{P}$, and σ_i is of the form $c_{H_{i-1},h}|_{B_{i-1}}$ (where $h \in G$), $\pi_{H_{i-1},H_i}|_{B_{i-1}}$, or $(\pi_{H_{i-1},H_i}|_{B_{i-1}})^{-1}$ (see Definition 2.7). As the maps $\pi_{H_i}|_{B_i} : B_i \rightarrow \pi_{H_i}(B_i)$ are bijective and commute with projections and conjugations, we see $\tau' := \sigma'_k \circ \dots \circ \sigma'_1$ is a nontrivial permutation of $\pi_H(B) \in C_H$, where each map $\sigma'_i := \pi_{H_i}|_{B_i} \circ \sigma_i \circ (\pi_{H_{i-1}}|_{B_{i-1}})^{-1}$ is of the form $c_{H_{i-1},h}^{\mathcal{D}}|_{B_{i-1}}$, $\pi_{H_{i-1},H_i}^{\mathcal{D}}|_{B_{i-1}}$, or $(\pi_{H_{i-1},H_i}^{\mathcal{D}}|_{B_{i-1}})^{-1}$. So \mathcal{C} is not strongly antisymmetric. The proof of antisymmetry is the same except that we only consider maps σ_i that are conjugations. \square

We need to compute $\mathcal{I} = \{I_K : K \in \mathcal{F}\}$ and an \mathcal{I} -advice $\{\mathcal{S}, \mathcal{T}\}$ such that the associated \mathcal{P} -collection $\tilde{\mathcal{C}}$ is $(\mathcal{C}, \mathcal{D}_{\Omega_0})$ -separated. The following lemma states that for \mathcal{P} -collections arising from Definition 5.8, the first condition of $(\mathcal{C}, \mathcal{D}_{\Omega_0})$ -separatedness is in fact automatic.

Lemma 5.21. *Let \mathcal{I} , $\{\mathcal{S}, \mathcal{T}\}$, \mathcal{C} and $\tilde{\mathcal{C}}$ be as in Definition 5.8. Then all the partitions in $\tilde{\mathcal{C}}$ are invariant under the action of \mathcal{D}_{Ω_0} by inverse right translation.*

To prove it, we need the following observation.

Lemma 5.22. *Let \mathfrak{m}_0 be a maximal ideal of A_{L,f_δ} containing $\frac{\Omega_0/p\mathcal{O}_L}{\text{Rad}(\mathcal{O}_L)}$. For all $x \in A_{L,f_\delta}$, $\omega \in \mathcal{I}_{\mathfrak{p}}$, and $\sigma \in \mathcal{D}_{\Omega_0}$ such that the image of σ in $\text{Gal}(\kappa_{\Omega_0}/\bar{\mathcal{O}}_{K_0})$ is the Frobenius automorphism $x \mapsto x^q$ over \mathbb{F}_q , it holds that $\omega x \equiv x \pmod{\mathfrak{m}_0}$ and $\sigma x \equiv \sigma_{L,f_\delta}(x) \pmod{\mathfrak{m}_0}$.*

Proof. By bilinearity, we may assume $x = a \otimes b$ where $a \in \bar{\mathcal{O}}_L/\text{Rad}(\bar{\mathcal{O}}_L)$ and $b \in \mathbb{F}_{q^{f_\delta}}$. As $\omega \in \mathcal{I}_{\mathfrak{p}}$, it holds that ${}^\omega a \equiv a \pmod{\frac{\mathfrak{Q}_0/p\mathcal{O}_L}{\text{Rad}(\bar{\mathcal{O}}_L)}}$ and hence ${}^\omega(a \otimes b) \equiv {}^\omega a \otimes b \equiv a \otimes b \pmod{\mathfrak{m}_0}$. Similarly, we have ${}^\sigma a \equiv a^q \pmod{\frac{\mathfrak{Q}_0/p\mathcal{O}_L}{\text{Rad}(\bar{\mathcal{O}}_L)}}$ by definition and hence ${}^\sigma(a \otimes b) \equiv {}^\sigma a \otimes b \equiv a^q \otimes b \equiv \sigma_{L,f_\delta}(a \otimes b) \pmod{\mathfrak{m}_0}$. \square

Now we are ready to prove Lemma 5.21.

Proof of Lemma 5.21. Consider $H \in \mathcal{P}$ and Hg, Hg' in the same block of \tilde{C}_H . Fix $h \in \mathcal{D}_{\mathfrak{Q}_0}$. We prove $Hgh^{-1}, Hg'h^{-1}$ are also in the same block by verifying the three conditions in Definition 5.8. Let B be the block of C_H containing both $Hg\mathcal{D}_{\mathfrak{Q}_0}$ and $Hg'\mathcal{D}_{\mathfrak{Q}_0}$. The first condition in Definition 5.8 obviously holds for Hgh^{-1} and $Hg'h^{-1}$ since $Hgh^{-1}\mathcal{D}_{\mathfrak{Q}_0} = Hg\mathcal{D}_{\mathfrak{Q}_0} \in B$ and $Hg'h^{-1}\mathcal{D}_{\mathfrak{Q}_0} = Hg'\mathcal{D}_{\mathfrak{Q}_0} \in B$.

Let K be the field in \mathcal{F} isomorphic to L^H over K_0 . Let δ be the idempotent in I_K satisfying $\tilde{\tau}_H(\delta) = \delta_B$ (see Definition 5.4). Suppose $e_\delta > 1$. By Definition 5.8, the order of the unique element c in $\kappa_{\mathfrak{Q}_0}^\times$ satisfying

$$g^{-1}s_{\delta,H} + I = c \cdot (g'^{-1}s_{\delta,H} + I)$$

is coprime to e_δ , where $I = (\mathfrak{Q}_0/p\mathcal{O}_L)^{e(\mathfrak{Q}_0)/e_\delta+1}$. We have ${}^h I = I$ since $h \in \mathcal{D}_{\mathfrak{Q}_0}$. Therefore

$${}^{hg^{-1}}s_{\delta,H} + I = {}^h c \cdot ({}^{hg'^{-1}}s_{\delta,H} + I),$$

where ${}^h c \in \kappa_{\mathfrak{Q}_0}^\times$ has the same order as c . So the second condition in Definition 5.8 is satisfied by Hgh^{-1} and $Hg'h^{-1}$.

Now suppose $f_\delta > 1$. Let \mathfrak{m}_0 be a maximal ideal of A_{L,f_δ} containing $\frac{\mathfrak{Q}_0/p\mathcal{O}_L}{\text{Rad}(\bar{\mathcal{O}}_L)}$. By Definition 5.8, the order of the unique element c in $(A_{L,f_\delta}/\mathfrak{m}_0)^\times$ satisfying

$$g^{-1}t_{\delta,H} + \mathfrak{m}_0 = c \cdot (g'^{-1}t_{\delta,H} + \mathfrak{m}_0)$$

is coprime to f_δ . Fix $\sigma \in \mathcal{D}_{\mathfrak{Q}_0}$ whose image in $\text{Gal}(\kappa_{\mathfrak{Q}_0}/\bar{\mathcal{O}}_{K_0})$ is the Frobenius automorphism $x \mapsto x^q$ over \mathbb{F}_q . Choose $\omega \in \mathcal{I}_{\mathfrak{Q}_0}$ and $i \in \mathbb{Z}$ such that $h = \omega\sigma^i$. By Lemma 5.22, we have

$$\begin{aligned} {}^{hg^{-1}}t_{\delta,H} &\equiv {}^{\omega\sigma^i}(g^{-1}t_{\delta,H}) \equiv \sigma^i(g^{-1}t_{\delta,H}) \equiv \sigma_{L,f_\delta}^i(g^{-1}t_{\delta,H}) \equiv g^{-1}(\sigma_{L,f_\delta}^i(t_{\delta,H})) \\ &\equiv g^{-1}(\xi^i \cdot t_{\delta,H}) \equiv \xi^i \cdot g^{-1}t_{\delta,H} \pmod{\mathfrak{m}_0}, \end{aligned}$$

where ξ is the primitive f_δ th root of unity satisfying $\sigma_{K,f_\delta}(t_\delta) = \xi \cdot t_\delta$ as in Definition 5.7. The same argument shows ${}^{hg'^{-1}}t_{\delta,H} \equiv \xi^i \cdot g'^{-1}t_{\delta,H} \pmod{\mathfrak{m}_0}$. It follows that

$${}^{hg^{-1}}t_{\delta,H} + \mathfrak{m}_0 = c \cdot ({}^{hg'^{-1}}t_{\delta,H} + \mathfrak{m}_0).$$

So the third condition in Definition 5.8 is also satisfied by Hgh^{-1} and $Hg'h^{-1}$. \square

We also show that \mathcal{P} -collections arising from Definition 5.8 always satisfy a weakening of the second condition of $(\mathcal{C}, \mathcal{D}_{\Omega_0})$ -separatedness, where bijectivity is replaced by injectivity:

Lemma 5.23. *Let \mathcal{I} , $\{\mathcal{S}, \mathcal{T}\}$, \mathcal{C} and $\tilde{\mathcal{C}}$ be as in Definition 5.8. Then for $H \in \mathcal{P}$, the map $\pi_H : H \backslash G \rightarrow H \backslash G / \mathcal{D}_{\Omega_0}$ sending $Hg \in H \backslash G$ to $Hg\mathcal{D}_{\Omega_0}$ maps each block of $\tilde{\mathcal{C}}_H$ injectively to a block of \mathcal{C}_H .*

Proof. Consider $H \in \mathcal{P}$ and $g \in G$ and $h \in \mathcal{D}_{\Omega_0}$ such that $Hg \neq Hgh^{-1}$. We want to prove that Hg and Hgh^{-1} are in different blocks of $\tilde{\mathcal{C}}_H$.

Let B be the block of \mathcal{C}_H containing $Hg\mathcal{D}_{\Omega_0} = Hgh^{-1}\mathcal{D}_{\Omega_0}$. Let K be the field in \mathcal{F} isomorphic to L^H over K_0 . Let δ be the idempotent in I_K satisfying $\tilde{\tau}_H(\delta) = \delta_B$ (see Definition 5.4). Fix $\sigma \in \mathcal{D}_{\Omega_0}$ whose image in $\text{Gal}(\kappa_{\Omega_0}/\bar{\mathcal{O}}_{K_0})$ is the Frobenius automorphism $x \mapsto x^q$ over \mathbb{F}_q .

As we assume $p > \deg(f)$, the wild inertia group $\mathcal{W}_{\Omega_0} \subseteq G$ of Ω_0 over K_0 is trivial. So \mathcal{I}_{Ω_0} is a cyclic group of order $e(\Omega_0)$. Fix a generator ω of \mathcal{I}_{Ω_0} . By Theorem 5.3 and Definition 5.2, we know e_δ is the smallest positive integer k satisfying $Hg\omega^{-k} = Hg$, and f_δ is the smallest positive integer k satisfying $Hg\sigma^{-k}\mathcal{I}_{\Omega_0} = Hg\mathcal{I}_{\Omega_0}$. So there exist unique $i \in \{0, \dots, f_\delta - 1\}$ and $j \in \{0, \dots, e_\delta - 1\}$ such that $Hgh^{-1} = Hg\sigma^{-i}\omega^{-j}$. As $Hg \neq Hgh^{-1}$, we have $(i, j) \neq (0, 0)$. By replacing h with $\omega^j\sigma^i$ if necessary, we may assume $h = \omega^j\sigma^i$.

First assume $i \neq 0$. Then $f_\delta > 1$. Let \mathfrak{m}_0 be a maximal ideal of A_{L, f_δ} containing $\frac{\Omega_0/p\mathcal{O}_L}{\text{Rad}(\mathcal{O}_L)}$. As shown in the proof of Lemma 5.21, we have

$$hg^{-1}t_{\delta, H} \equiv \xi^i \cdot g^{-1}t_{\delta, H} \pmod{\mathfrak{m}_0},$$

where ξ is a primitive f_δ th root of unity. The order of ξ^i is $f_\delta / \gcd(f_\delta, i) > 1$ and is a divisor of f_δ . So the third condition in Definition 5.8 is not satisfied by Hg and Hgh^{-1} . It follows that Hg and Hgh^{-1} are in different blocks of $\tilde{\mathcal{C}}_H$, as desired.

Now assume $i = 0$ and $j \neq 0$. Then $e_\delta > 1$. Let $\mathfrak{m}_e = \Omega_0/p\mathcal{O}_L$ and $k = e(\Omega_0)/e_\delta$. As shown in the proof of Lemma 5.18, we have $g^{-1}s_{\delta, H} \in \mathfrak{m}_e^k - \mathfrak{m}_e^{k+1}$. Choose $\pi_L \in \mathfrak{m}_e - \mathfrak{m}_e^2$. We have a group homomorphism $\mathcal{I}_{\Omega_0} \rightarrow \kappa_{\Omega_0}^\times$ sending $g \in \mathcal{I}_{\Omega_0}$ to the unique element $c_g \in \kappa_{\Omega_0}^\times$ satisfying $g\pi_L + \mathfrak{m}_e^2 = c_g(\pi_L + \mathfrak{m}_e^2)$. This map is

injective since its kernel is $\mathcal{W}_{\Omega_0} = \{e\}$. In particular, we know c_ω is a primitive $e(\Omega_0)$ th root of unity in $\kappa_{\Omega_0}^\times$. Choose $c \in \kappa_{\Omega_0}^\times$ such that

$$g^{-1} s_{\delta,H} + \mathfrak{m}_e^{k+1} = c(\pi_L^k + \mathfrak{m}_e^{k+1}),$$

which exists since $g^{-1} s_{\delta,H}$ and π_L^k are both in $\mathfrak{m}_e^k - \mathfrak{m}_e^{k+1}$. Then we have

$$\begin{aligned} hg^{-1} s_{\delta,H} + \mathfrak{m}_e^{k+1} &= \omega^j (g^{-1} s_{\delta,H} + \mathfrak{m}_e^{k+1}) = \omega^j (c(\pi_L^k + \mathfrak{m}_e^{k+1})) \\ &= c \cdot c_\omega^{jk} \cdot (\pi_L^k + \mathfrak{m}_e^{k+1}) = c_\omega^{jk} \cdot (g^{-1} s_{\delta,H} + \mathfrak{m}_e^{k+1}). \end{aligned}$$

The order of $c_\omega^{jk} \in \kappa_{\Omega_0}^\times$ is $e(\Omega_0) / \gcd(e(\Omega_0), jk) = e_\delta / \gcd(e_\delta, j) > 1$, which is a divisor of e_δ . So the second condition in Definition 5.8 is not satisfied by Hg and Hgh^{-1} . It follows that Hg and Hgh^{-1} are in different blocks of $\tilde{\mathcal{C}}_H$, as desired. \square

In the next section, we give subroutines that refine the idempotent decompositions I_K so that $\tilde{\mathcal{C}}$ is eventually a compatible, invariant, $(\mathcal{C}, \mathcal{D})$ -separated \mathcal{P} -collection, and hence a strongly antisymmetric \mathcal{P} -scheme.

5.9 Producing an ordinary \mathcal{P} -scheme

We modify the algorithm `ComputeDoubleCosetPscheme` in Section 5.4 so that a $(\mathcal{C}, \mathcal{D}_{\Omega_0})$ -separated strongly antisymmetric \mathcal{P} -scheme is produced.

The pseudocode of the modified algorithm is given in Algorithm 14. Again, the algorithm takes a (K_0, \tilde{f}) -subfield system \mathcal{F} as the input, and outputs for every $K \in \mathcal{F}$ an idempotent decomposition I_K of the ring R_K , together with some auxiliary data.

The first half (Lines 1–10) of the algorithm is the preprocessing stage: we compute the same data as in the algorithm `ComputeDoubleCosetPscheme`. In addition, for $K \in \mathcal{F}$, we compute the inclusion $\mathbb{F}_q \hookrightarrow \bar{\mathcal{O}}_K / \text{Rad}(\bar{\mathcal{O}}_K)$ at Line 4, endowing $\bar{\mathcal{O}}_K / \text{Rad}(\bar{\mathcal{O}}_K)$ the structure of an \mathbb{F}_q -algebra.¹⁵ And for $1 \leq i \leq [K : K_0]$, we compute the ring $A_{K,i} = \bar{\mathcal{O}}_K / \text{Rad}(\bar{\mathcal{O}}_K) \otimes_{\mathbb{F}_q} \mathbb{F}_{q^i}$ together with the inclusions $\bar{\mathcal{O}}_K / \text{Rad}(\bar{\mathcal{O}}_K) \hookrightarrow A_{K,i}$, $\mathbb{F}_{q^i} \hookrightarrow A_{K,i}$ defined by $a \mapsto a \otimes 1$ and $b \mapsto 1 \otimes b$ respectively.

The second half (Lines 11–24) of the algorithm refines the idempotent decompositions I_K for $K \in \mathcal{F}$. The loop in Lines 13–19 is the same as in the algorithm

¹⁵To achieve this, we compute the image \bar{Y} of $Y + (\tilde{h}(Y)) \in \mathcal{O}_{K_0} \subseteq \mathcal{O}_K$ in $\bar{\mathcal{O}}_K$ by Lemma 3.9. Then compute the map $\mathbb{F}_p[Y]/(h(Y)) \rightarrow \bar{\mathcal{O}}_K$ sending $Y + (h(Y))$ to \bar{Y} , and compose it with the isomorphism $\psi_0^{-1} : \mathbb{F}_q \rightarrow \mathbb{F}_p[Y]/(h(Y))$ and the quotient map $\bar{\mathcal{O}}_K \rightarrow \bar{\mathcal{O}}_K / \text{Rad}(\bar{\mathcal{O}}_K)$.

Algorithm 14 ComputeOrdinaryPscheme

Input: (K_0, \tilde{f}) -subfield system \mathcal{F}

Output: for every $K \in \mathcal{F}$: the outputs of ComputeRings (see Lemma 5.7) on the input (K, p) , and an idempotent decomposition I_K of R_K

```

1: for  $K \in \mathcal{F}$  do
2:   call ComputeRings on  $(K, p)$ 
3:    $I_K \leftarrow \{1\}$ , where 1 denotes the unity of  $R_K$ 
4:   compute the inclusion  $\mathbb{F}_q \hookrightarrow \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ 
5:   for  $i \leftarrow 1$  to  $[K : K_0]$  do
6:     compute  $A_{K,i}$  and the inclusions  $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K) \hookrightarrow A_{K,i}, \mathbb{F}_{q^i} \hookrightarrow A_{K,i}$ 
7:   for  $(K, K') \in \mathcal{F}^2$  do
8:     call ComputeRelEmbeddings to compute all the embeddings from  $K$  to  $K'$ 
       over  $K_0$ 
9:     for embedding  $\phi : K \hookrightarrow K'$  over  $K_0$  do
10:      call ComputeRingHoms on  $p, K, K'$  and  $\phi$  compute  $\bar{\phi}, \hat{\phi}$  and  $\tilde{\phi}$ 
11:   repeat
12:     repeat
13:       repeat
14:         call CompatibilityAndInvarianceTestV2
15:         call RegularityTestV2
16:         call StrongAntisymmetryTestV2
17:         call RamificationIndexTest
18:         call InertiaDegreeTest
19:       until  $I_K$  remains the same in the last iteration for all  $K \in \mathcal{F}$ 
20:       call ComputeAdvice on  $\mathcal{I} = \{I_K : K \in \mathcal{F}\}$ 
21:     until  $I_K$  remains the same in the last iteration for all  $K \in \mathcal{F}$ 
22:     call SurjectivityTest
23:     call RingHomTest
24:   until  $I_K$  remains the same in the last iteration for all  $K \in \mathcal{F}$ 
25: return the outputs of ComputeRings on the input  $(K, p)$  and  $I_K$  for  $K \in \mathcal{F}$ 

```

`ComputeDoubleCosetPscheme`. It produces idempotent decompositions I_K that define a strongly antisymmetric \mathcal{P} -scheme of double cosets $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ with respect to \mathcal{D}_{Ω_0} , in which all the partitions have locally constant ramification indices and inertia degrees (with respect to $(\mathcal{D}_{\Omega_0}, \mathcal{I}_{\Omega_0})$). After that, we call the subroutine `ComputeAdvice` in Lemma 5.17 on $\mathcal{I} = \{I_K : K \in \mathcal{F}\}$ at Line 20. It either properly refines some I_K or returns an \mathcal{I} -advice. In the former case, we start over from Line 13.

So assume an \mathcal{I} -advice $\{\mathcal{S}, \mathcal{T}\}$ is returned at Line 20. Let $\tilde{\mathcal{C}} = \{\tilde{C}_H : H \in \mathcal{P}\}$ be the \mathcal{P} -collection associated with \mathcal{I} and $\{\mathcal{S}, \mathcal{T}\}$. Next we need two new subroutines, `SurjectivityTest` and `RingHomTest`:

Lemma 5.24. *Under GRH, there exists a subroutine `SurjectivityTest` that updates I_K in time polynomial in $\log p$ and the size of \mathcal{F} so that the partitions $C_H \in \mathcal{C}$ are refined. Moreover, at least one partition C_H is properly refined unless for all $H \in \mathcal{P}$, the map $\pi_H : H \backslash G \rightarrow H \backslash G / \mathcal{D}_{\Omega_0}$ sending $Hg \in H \backslash G$ to $Hg\mathcal{D}_{\Omega_0}$ maps each block of \tilde{C}_H surjectively to a block of C_H .*

Lemma 5.25. *Under GRH, there exists a subroutine `RingHomTest` that updates I_K in time polynomial in $\log p$ and the size of \mathcal{F} so that the partitions $C_H \in \mathcal{C}$ are refined. Moreover, at least one partition C_H is properly refined unless $\tilde{\mathcal{C}}$ is compatible and invariant.*

The proofs of the above two lemmas are the most technical part of this chapter. We defer them to Appendix C.

We run these two subroutines and repeat, until no idempotent decomposition I_K is properly refined in the last iteration. By Lemma 5.21, Lemma 5.23, Lemma 5.24, and Lemma 5.25, the resulting \mathcal{P} -collection $\tilde{\mathcal{C}}$ is a compatible, invariant, $(\mathcal{C}, \mathcal{D}_{\Omega_0})$ -separated \mathcal{P} -collection. Also note that \mathcal{C} is a strongly antisymmetric \mathcal{P} -scheme of double cosets with respect to \mathcal{D}_{Ω_0} . It follows from Lemma 5.20 that $\tilde{\mathcal{C}}$ is a strongly antisymmetric $(\mathcal{C}, \mathcal{D}_{\Omega_0})$ -separated \mathcal{P} -scheme. We conclude

Theorem 5.8. *Under GRH, the algorithm `ComputeOrdinaryPscheme` runs in time polynomial in $\log p$ and the size of \mathcal{F} , and when it terminates, \mathcal{C} is a strongly antisymmetric \mathcal{P} -scheme of double cosets (with respect to \mathcal{D}_{Ω_0}), and $\tilde{\mathcal{C}}$ is a $(\mathcal{C}, \mathcal{D}_{\Omega_0})$ -separated strongly antisymmetric \mathcal{P} -scheme.*

5.10 Putting it together

We combine the results in previous sections to obtain the generalized \mathcal{P} -scheme algorithm. For simplicity, we first focus on computing the complete factorization of the input polynomial f . The problem of computing a proper factorization of f is discussed later in this section.

The algorithm takes a polynomial $f(X) \in \mathbb{F}_q[X]$ and an irreducible lifted polynomial $\tilde{f}(X) \in A_0[X]$ as the input, and outputs the complete factorization of f . Its pseudocode is given in Algorithm 15 below.

Algorithm 15 GeneralizedPschemeAlgorithm

Input: $f(X) \in \mathbb{F}_q[X]$ and its irreducible lifted polynomial $\tilde{f}(X) \in A_0[X]$

Output: factorization of f

- 1: **if** $p \leq \deg(f)$ **then** run Berlekamp's algorithm in [Ber70] to compute the complete factorization of f , output it and halt
 - 2: call `ComputeRelNumberFields` to compute a (K_0, \tilde{f}) -subfield system \mathcal{F} such that (1) $F = K_0[X]/(\tilde{f}(X)) \in \mathcal{F}$, and (2) for some $H \in \mathcal{P}$ satisfying $L^H \cong_{K_0} F$, all strongly antisymmetric \mathcal{P} -schemes are discrete on H , where \mathcal{P} is the subgroup system over $G = \text{Gal}(\tilde{f}/K_0)$ associated with \mathcal{F}
 - 3: call `ComputeOrdinaryPscheme` on \mathcal{F} to obtain I_K for $K \in \mathcal{F}$
 - 4: call `ExtractFactorsV2` to extract a factorization of f from I_F , and output it
-

Line 1 checks whether $p > \deg(f)$ holds. If $p \leq \deg(f)$, we just run Berlekamp's algorithm in [Ber70] to compute the complete factorization of f in time polynomial in p and $\deg(f)$, output it, and halt. This step justifies the assumption $p > \deg(f)$ made in Section 5.7–5.9.

The subroutine `ComputeRelNumberFields` at Line 2 is the generic part of the algorithm. It is supposed to compute a (K_0, \tilde{f}) -subfield system \mathcal{F} such that $F \in \mathcal{F}$, and the associated subgroup system \mathcal{P} over G satisfies a certain combinatorial property (see Theorem 5.9 below). The algorithm `ComputeOrdinaryPscheme` (see Section 5.4) at Line 3 takes the input \mathcal{F} and outputs data that includes the idempotent decomposition I_F of R_F . Finally, we call the subroutine `ExtractFactorsV2` (see Section 5.3) at Line 4 to extract a factorization of f from I_F .

The following theorem is the main result of this chapter.

Theorem 5.9. *Suppose there exists a deterministic algorithm that given a polynomial $g(X) \in K_0[X]$ irreducible over K_0 , constructs a (K_0, g) -subfield system \mathcal{F} in time*

$T(g)$ such that

- $K_0[X]/(g(X))$ is in \mathcal{F} , and
- for some $H \in \mathcal{P}$ satisfying $(L(g))^H \cong_{K_0} K_0[X]/(g(X))$, all strongly antisymmetric \mathcal{P} -schemes are discrete on H , where \mathcal{P} is the subgroup system over $\text{Gal}(g/K_0)$ associated with \mathcal{F} , and $L(g)$ is the splitting field of g over K_0 .

Then under GRH, there exists a deterministic algorithm that given a polynomial $f(X) \in \mathbb{F}_q[X]$ and an irreducible lifted polynomial $\tilde{f}(X) \in A_0[X]$ of f , outputs the complete factorization of f over \mathbb{F}_q in time polynomial in $T(\tilde{f})$ and the size of the input.

Proof. Consider the algorithm `GeneralizedPschemeAlgorithm` above and implement the subroutine `ComputeRelNumberFields` using the hypothetical algorithm in the theorem. The case $p \leq \deg(f)$ is solved by Berlekamp's algorithm in [Ber70]. So assume $p > \deg(f)$. Choose $g = \tilde{f}$. By Theorem 5.8, the \mathcal{P} -collection $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ defined by $C_H = P(\tilde{\tau}_H(I_K))$ is a strongly antisymmetric \mathcal{P} -scheme of double cosets with respect to \mathcal{D}_{Ω_0} , and the \mathcal{P} -collection $\tilde{\mathcal{C}} = \{\tilde{C}_H : H \in \mathcal{P}\}$ associated with the collection of idempotent decompositions $\mathcal{I} = \{I_K : K \in \mathcal{F}\}$ and the \mathcal{I} -advice produced by the algorithm `ComputeOrdinaryPscheme` is a $(\mathcal{C}, \mathcal{D}_{\Omega_0})$ -separated strongly antisymmetric \mathcal{P} -scheme. By the second condition in the theorem, we have $\tilde{C}_H = \infty_{H \setminus G}$ (and hence $C_H = \infty_{H \setminus G / \mathcal{D}_{\Omega_0}}$) for some $H \in \mathcal{P}$ satisfying $L^H \cong_{K_0} F$. So the idempotent decomposition I_F is complete. By Theorem 5.6, the algorithm outputs the complete factorization of f over \mathbb{F}_q .

The subroutine `ComputeRelNumberFields` runs in time $T(\tilde{f})$. In particular, the size of \mathcal{F} is bounded by $T(\tilde{f})$. The claim about the running time then follows from Theorem 5.8 and Theorem 5.6. \square

By Theorem 5.9 and Lemma 4.10, we have the following partial generalization of Corollary 3.2.

Corollary 5.2. *Under GRH, there exists a deterministic algorithm that, given a polynomial $f(X) \in \mathbb{F}_q[X]$ of degree $n \in \mathbb{N}^+$ and an irreducible¹⁶ lifted polynomial $\tilde{f}(X) \in A_0[X]$ of f , computes the complete factorization of f over \mathbb{F}_q in time*

¹⁶The assumption that \tilde{f} is irreducible is not necessary, and can be avoided by adapting Lemma 4.10. We omit the details.

polynomial in $n^{d(G)}$ and the size of the input, where G is the permutation group $\text{Gal}(f/K_0)$ acting on the set of roots of \tilde{f} .

The unifying framework via the generalized \mathcal{P} -scheme algorithm. In the following, we use Theorem 5.9 and Corollary 5.2 to derive generalizations of the main results in [Hua91a; Hua91b; Rón88; Rón92; Evd92; Evd94; IKS09] in a uniform way.

Given a polynomial $f(X) \in \mathbb{F}_q[X]$ and a (possibly reducible) lifted polynomial $\tilde{f}(X) \in A_0[X]$ of f . We reduce to the case that the lifted polynomial is irreducible as follows: using Lemma 5.1, we compute an integer D satisfying $D \equiv 1 \pmod{p}$ and a factorization of $D \cdot \tilde{f}$ into irreducible factors $\tilde{f}_1, \dots, \tilde{f}_k \in A_0[X]$ over K_0 . Then we have

$$f(X) = \prod_{i=1}^k \tilde{\psi}_0(\tilde{f}_i)(X)$$

and the problem of factoring $f(X)$ is reduced to the problem of factoring each $\tilde{\psi}_0(\tilde{f}_i) \in \mathbb{F}_q[X]$ with the aid of its irreducible lifted polynomial $\tilde{f}_i(X)$ (see the discussion after Lemma 5.1). Moreover, for $i \in [k]$, the Galois group $\text{Gal}(\tilde{f}_i(X)/K_0)$ is a quotient group of $\text{Gal}(\tilde{f}/K_0)$, and hence $|\text{Gal}(\tilde{f}_i(X)/K_0)| \leq |\text{Gal}(\tilde{f}(X)/K_0)|$.

So assume \tilde{f} is irreducible over K_0 . We choose $\mathcal{F} = \{F, L\}$ where $F = K_0[X]/(\tilde{f}(X))$ and L is the splitting field of \tilde{f} over K_0 . Compute \mathcal{F} in time polynomial in $[L : K_0] = |\text{Gal}(\tilde{f}(X)/K_0)|$ and the size of \tilde{f} using Lemma 4.9. By Lemma 2.4, all antisymmetric \mathcal{P} -schemes are discrete on H for all $H \in \mathcal{P}$ since the trivial subgroup $\{e\}$ is in \mathcal{P} . Therefore by Theorem 5.9 and the reduction above, we have the following generalization of Theorem 3.10.

Theorem 5.10. *Under GRH, there exists a deterministic algorithm that, given a polynomial $f(X) \in \mathbb{F}_q[X]$ and a lifted polynomial $\tilde{f}(X) \in A_0[X]$ of f , computes the complete factorization of f over \mathbb{F}_q in time polynomial in $|\text{Gal}(\tilde{f}/K_0)|$ and the size of the input.*

Note $|\text{Gal}(\tilde{f}/\mathbb{Q})| \leq \deg(f)$ when $\text{Gal}(\tilde{f}/\mathbb{Q})$ is abelian. So we have the following generalization of Corollary 3.3.

Corollary 5.3. *Under GRH, there exists a deterministic algorithm that, given a polynomial $f(X) \in \mathbb{F}_q[X]$ and a lifted polynomial $\tilde{f}(X) \in A_0[X]$ of f such that $\text{Gal}(\tilde{f}/K_0)$ is abelian, computes the complete factorization of f over \mathbb{F}_q in polynomial time.*

Suppose only the polynomial f is known. Let $n = \deg(f)$. We can efficiently compute a lifted polynomial $\tilde{f}(X) \in A_0[X]$ of f whose size is polynomial in n and $\log q$.¹⁷ Reduce to the case that \tilde{f} is irreducible over K_0 as above. As $\text{Gal}(\tilde{f}(X)/K_0)$ is a subgroup of $\text{Sym}(n)$, we have the following generalization of Theorem 3.11.

Theorem 5.11. *Under GRH, there exists a deterministic algorithm that, given a polynomial $f(X) \in \mathbb{F}_q[X]$ of degree $n \in \mathbb{N}^+$, computes the complete factorization of f in time polynomial in $n!$ and $\log q$.*

Now suppose we lift f to \tilde{f} , reduce to the case that \tilde{f} is irreducible over K_0 , but compute \mathcal{F} using Lemma 4.10 instead. By Corollary 5.2 and Lemma 2.6, we have the following generalization of Theorem 3.12.

Theorem 5.12. *Under GRH, there exists a deterministic algorithm that, given a polynomial $f(X) \in \mathbb{F}_q[X]$ of degree $n \in \mathbb{N}^+$, computes the complete factorization of f over \mathbb{F}_q in time polynomial in $n^{\log n}$ and $\log q$.*

We also have the following theorem that generalizes Theorem 4.3 and the main result of [Evd92]. The proof is the same as that of Theorem 4.3, except that Theorem 5.9 is used instead of Theorem 3.9, and the base field \mathbb{Q} is replaced by K_0 .

Theorem 5.13. *Under GRH, there exists a deterministic polynomial-time algorithm that, given a polynomial $f(X) \in \mathbb{F}_q[X]$ and a lifted polynomial $\tilde{f}(X) \in A_0[X]$ of f whose Galois group $\text{Gal}(\tilde{f}/K_0)$ is solvable, computes the complete factorization of f over \mathbb{F}_q .*

Computing a proper factorization of f . Unlike the special case considered in Chapter 3, replacing discreteness by inhomogeneity in the second condition of Theorem 5.9 does not automatically yield an algorithm computing a proper factorization of f . The reason is that even if a $(\mathcal{C}, \mathcal{D}_{\Omega_0})$ -separated \mathcal{P} -scheme $\tilde{\mathcal{C}}$ is inhomogeneous on a subgroup $H \in \mathcal{P}$, the corresponding \mathcal{P} -scheme of double cosets \mathcal{C} may still be homogeneous on H . In fact, this is always the case when $H \backslash G / \mathcal{D}_{\Omega_0}$ is a singleton, or equivalently, when \mathcal{D}_{Ω_0} acts transitively on $H \backslash G$ by inverse right translation.

Still, by adapting the condition, we obtain some results on computing a proper factorization of f :

¹⁷We also need to choose $\tilde{h}(Y) \in \mathbb{Z}[Y]$, $h(Y) = \tilde{h}(Y) \bmod p \in \mathbb{F}_p[Y]$ and $\psi_0 : \mathbb{F}_p[Y]/(h(Y)) \rightarrow \mathbb{F}_q$ first, so that $A_0 = \mathbb{Z}[Y]/(\tilde{h}(Y))$, $K_0 = \mathbb{Q}[Y]/(\tilde{h}(Y))$ and the notion of lifted polynomials are defined. The isomorphism ψ_0 can be efficiently computed by [Len91].

- We formulate a new condition on \mathcal{P} -schemes and use it to obtain algorithms computing one irreducible factor of f . See Theorem 5.14 and Corollary 5.4.
- We formulate conditions on \mathcal{P} that involve not only ordinary \mathcal{P} -schemes but also \mathcal{P} -schemes of double cosets, and these conditions can be used for computing the complete factorization as well as a proper factorization. See Theorem 5.14.
- Finally, we prove a generalization of Lemma 2.18 for \mathcal{P} -schemes of double cosets, and use it to prove a generalization of Theorem 3.13. See Theorem 5.15.

First, we introduce the following definition.

Definition 5.10. *For a \mathcal{P} -scheme (resp. \mathcal{P} -scheme of double cosets) $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ and $H \in \mathcal{P}$, we say \mathcal{C} has a singleton on H if the partition C_H has a block that is a singleton.*

The following theorem is a variant of Theorem 5.9 with weakened conditions on the subgroup system \mathcal{P} .

Theorem 5.14. *Suppose there exists a deterministic algorithm that, given a polynomial $g(X) \in K_0[X]$ irreducible over K_0 , constructs a (K_0, g) -subfield system \mathcal{F} in time $T(g)$ such that*

- $K_0[X]/(g(X))$ is in \mathcal{F} , and
- for some $H \in \mathcal{P}$ satisfying $(L(g))^H \cong_{K_0} K_0[X]/(g(X))$, all strongly antisymmetric \mathcal{P} -schemes of double cosets \mathcal{C} with respect to \mathcal{D}_{Ω_0} that admit a $(\mathcal{C}, \mathcal{D}_{\Omega_0})$ -separated strongly antisymmetric \mathcal{P} -scheme are discrete (resp. are inhomogeneous, have a singleton) on H ,¹⁸ where \mathcal{P} is the subgroup system over $\text{Gal}(g/K_0)$ associated with \mathcal{F} , and $L(g)$ is the splitting field of g over K_0 .

Then under GRH, there exists a deterministic algorithm that, given a polynomial $f(X) \in \mathbb{F}_q[X]$ and an irreducible lifted polynomial $\tilde{f}(X) \in A_0[X]$ of f , outputs the complete factorization (resp. a proper factorization, an irreducible factor) of f over \mathbb{F}_q in time polynomial in $T(\tilde{f})$ and the size of the input.

¹⁸Here \mathcal{D}_{Ω_0} is the decomposition group of a fixed prime ideal Ω_0 of $\mathcal{O}_{L(g)}$ lying over p . Different choices of Ω_0 lead to conjugate subgroups and hence do not matter.

Proof. The proof is the almost same as that of Theorem 5.9. The second condition in the theorem are used to show $C_H = \infty_{\mathcal{D}_{\Omega_0}}$ (resp. $C_H \neq 0_{\mathcal{D}_{\Omega_0}}$, C_H has a singleton) for some $H \in \mathcal{P}$ satisfying $L^H \cong_{K_0} K_0[X]/(\tilde{f}(X))$, and hence the corresponding idempotent decomposition is complete (resp. is proper, has a singleton). Then apply Theorem 5.6. The details are left to the reader. \square

Observe that if a \mathcal{P} -scheme of double cosets \mathcal{C} has a singleton on H , then a $(\mathcal{C}, \mathcal{D}_{\Omega_0})$ -separated \mathcal{P} -scheme also has a singleton on H . So we have the following corollary, which is an analogue of Theorem 5.9.

Corollary 5.4. *Suppose there exists a deterministic algorithm that, given a polynomial $g(X) \in K_0[X]$ irreducible over K_0 , constructs a (K_0, g) -subfield system \mathcal{F} in time $T(g)$ such that*

- $K_0[X]/(g(X))$ is in \mathcal{F} , and
- for some $H \in \mathcal{P}$ satisfying $(L(g))^H \cong_{K_0} K_0[X]/(g(X))$, all strongly anti-symmetric \mathcal{P} -schemes have a singleton on H , where \mathcal{P} is the subgroup system over $\text{Gal}(g/K_0)$ associated with \mathcal{F} and $L(g)$ is the splitting field of g over K_0 .

Then under GRH, there exists a deterministic algorithm that, given a polynomial $f(X) \in \mathbb{F}_q[X]$ and an irreducible lifted polynomial $\tilde{f}(X) \in A_0[X]$ of f , outputs an irreducible factor of f over \mathbb{F}_q in time polynomial in $T(\tilde{f})$ and the size of the input.

Finally, we give a generalization of Theorem 3.13:

Theorem 5.15. *Under GRH, there exists a deterministic algorithm that, given a polynomial $f(X) \in \mathbb{F}_q[X]$ of degree $n \in \mathbb{N}^+$ that has $k > 1$ irreducible factors over \mathbb{F}_q , computes a proper factorization of f in time polynomial in n^ℓ and $\log q$, where ℓ is the least prime factor of k .*

To prove Theorem 5.15, we need the following generalization of Lemma 2.18, whose proof is deferred to Appendix C.

Lemma 5.26. *Let G be a finite group acting transitively on a set S . Let \mathcal{D} be a subgroup of G and let k be the number of \mathcal{D} -orbits in S . Suppose $k > 1$. Let $\ell \in \mathbb{N}^+$ be the least prime factor of k . Let $\mathcal{P} = \mathcal{P}_m$ be the system of stabilizers of depth m for some $m \geq \ell$ (with respect to the action of G on S). Then for any $x \in S$ and any \mathcal{P} -scheme of double cosets \mathcal{C} with respect to \mathcal{D} that is homogeneous on G_x , there exists no antisymmetric $(\mathcal{C}, \mathcal{D})$ -separated \mathcal{P} -scheme.*

Proof of Theorem 5.15. We may assume the irreducible factors of f over \mathbb{F}_q are all distinct and have the same degree d , since otherwise a proper factorization of f can be found by square-free factorization [Yun76; Knu98] or distinct-degree factorization [CZ81]. Compute d as the smallest positive integer for which the automorphism $x \mapsto x^{q^d}$ fixes $\mathbb{F}_q[X]/(f(X))$. Then compute $k = n/d$ and ℓ .

As in the proof of Theorem 5.11, we choose a lifted polynomial $\tilde{f} \in A_0[X]$ of f whose size is polynomial in n and $\log q$, and reduce to the case that \tilde{f} is irreducible over K_0 . Use Lemma 4.10 to compute \mathcal{F} so that the associated subgroup system \mathcal{P} is the system of stabilizers of depth ℓ with respect to the action of $\text{Gal}(\tilde{f}/K_0)$ on the set of roots of \tilde{f} in L . This step takes time polynomial in $c(\mathcal{P})$ and the size of \tilde{f} , which is polynomial in n^ℓ and $\log q$. The theorem then follows from Theorem 5.14 and Lemma 5.26. \square

Remark. We may also derive Theorem 5.15 from Theorem 3.13 by reducing to the case that f satisfies Condition 3.1: by square-free factorization, we may assume f is square-free. Compute the subring R of $\mathbb{F}_q[X]/(f(X))$ fixed by the Frobenius automorphism $x \mapsto x^p$ over \mathbb{F}_p . Then find an element $z \in R$ such that the minimal polynomial g of z over \mathbb{F}_p is a degree- k polynomial satisfying Condition 3.1. Such an element z exists if $p \geq k$. Then reduce to the problem of computing a proper factorization of g over \mathbb{F}_p . We leave the details to the reader.

Chapter 6

CONSTRUCTING NEW \mathcal{P} -SCHEMES FROM OLD ONES

In the previous chapters, we developed a framework for polynomial factoring whose correctness relies on combinatorial properties of \mathcal{P} -schemes. Motivated by it, we continue our study on \mathcal{P} -schemes in this chapter and also in subsequent chapters. Techniques introduced in this chapter have a common theme, namely constructing new \mathcal{P} -schemes from old ones. Such techniques include

- Inverse right translation on the set of \mathcal{P} -schemes.
- Restriction of \mathcal{P} -schemes to a subgroup, and its analogue for m -schemes.
- Passing to quotient groups.
- Induction of \mathcal{P} -schemes.
- Extension of \mathcal{P} -schemes to the closure of \mathcal{P} .
- Restriction of m -schemes to a subset, and its generalization for \mathcal{P} -schemes.
- Constructing primitive m -schemes from a general one.
- Direct products and wreath products.

The first three of them are introduced in Section 6.1. We use them to prove Lemma 4.12, as promised in Chapter 4.

In Section 6.2, we discuss the *induction* of \mathcal{P} -schemes. For $G' \subseteq G$ and a subgroup system \mathcal{P} over G , this operation produces a \mathcal{P} -scheme from a \mathcal{P}' -scheme, where \mathcal{P}' is a certain subgroup system over G' . We apply this operation in Section 6.3 to establish reductions among a family of conjectures concerning \mathcal{P} -schemes, whose resolution would imply that polynomial factoring over finite fields can be solved in deterministic polynomial time under GRH if an irreducible lifted polynomial with a special Galois group is given. See below for a more detailed discussion on these conjectures.

The rest of the above list is discussed in Section 6.4–6.7. In particular, we discuss *primitivity* of homogeneous m -schemes in Section 6.6. By exploiting the connection

between homogeneous primitive orbit m -schemes and primitive permutation groups, we prove that for $m \geq 3$, every antisymmetric homogeneous orbit m -scheme on a finite set S where $|S| > 1$ has a matching. Previously this was known for $m \geq 4$, as proved in [IKS09].

Schemes conjectures. The work [IKS09] proposed a conjecture on m -schemes called the *schemes conjecture*.

Conjecture (schemes conjecture). *There exists a constant $m \in \mathbb{N}^+$ such that every antisymmetric homogeneous m -scheme on a finite set S where $|S| > 1$ has a matching.*

Assuming this conjecture (and GRH), polynomial factorization over finite fields can be solved in deterministic polynomial time, as shown in [IKS09]. We reprove this result in Section 6.3 using a \mathcal{P} -scheme algorithm.

For each family \mathcal{G} of finite permutation groups, we also formulate an analogous conjecture, called the *schemes conjecture for \mathcal{G}* , in terms of the notation $d(G)$ introduced in Definition 2.8.

Conjecture (schemes conjecture for \mathcal{G}). *There exists a constant $m \in \mathbb{N}^+$ such that $d(G) \leq m$ for all $G \in \mathcal{G}$.*

We show that assuming this conjecture (and GRH), a polynomial f over finite fields can be factorized in deterministic polynomial time if we are also given an irreducible lifted polynomial \tilde{f} of f whose Galois group is in \mathcal{G} (as a permutation group on the set of roots of \tilde{f}).

Using induction of \mathcal{P} -schemes, we establish reductions among these conjectures for various families \mathcal{G} , so that the schemes conjecture for \mathcal{G} reduces to that for \mathcal{G}' if the permutation groups in \mathcal{G} are “less complex” than those in \mathcal{G}' . In particular, all these conjectures reduce to the one for the family of symmetric groups, and the latter turns out to be equivalent to (a slight relaxation of) the original schemes conjecture. In summary, the scheme conjectures for various families of finite permutation groups form a hierarchy of relaxations of the original schemes conjecture.

Therefore, in order to prove the original schemes conjecture, it is necessary to prove our analogous conjectures for families of less complex permutation groups. On the other hand, one may hope that progress on the latter would shed some light on the

original conjecture. We will follow this approach in subsequent chapters and prove some nontrivial results.

6.1 Basic operations on \mathcal{P} -schemes

In this section, we introduce some basic operations on \mathcal{P} -schemes, including inverse right translation, restriction, and passing to quotient groups. We then use them to prove Lemma 4.12.

Inverse right translation of \mathcal{P} -schemes. Let \mathcal{P} be a subgroup system over a finite group G . For each $H \in \mathcal{P}$, the group G acts on $H \backslash G$ by inverse right translation ${}^g H h = H h g^{-1}$. This action induces an action of G on the set of partitions of $H \backslash G$, defined by ${}^g P = \{{}^g B : B \in P\}$ for a partition P of $H \backslash G$. Then G also acts on the set of \mathcal{P} -collections by inverse right translation:

Definition 6.1. *The action of G on the set of \mathcal{P} -collections by inverse right translation is defined as follows: for a \mathcal{P} -collection $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ and $g \in G$, define ${}^g \mathcal{C} = \{{}^g C_H : H \in \mathcal{P}\}$.*

Lemma 6.1. *For a \mathcal{P} -scheme \mathcal{C} and $g \in G$, the \mathcal{P} -collection ${}^g \mathcal{C}$ is also a \mathcal{P} -scheme. Moreover, if \mathcal{C} is antisymmetric (resp. strongly antisymmetric), so is ${}^g \mathcal{C}$.*

Proof. This follows in a straightforward manner from G -equivariance of projections and conjugations (see Lemma 2.2). \square

So G also acts on the set of \mathcal{P} -schemes by inverse right translation, which preserves antisymmetry and strong antisymmetry.

Restriction to a subgroup. We define the *restriction* of a subgroup system \mathcal{P} over G and that of \mathcal{P} -collections to a subgroup of G .

Definition 6.2 (restriction). *Let \mathcal{P} be a subgroup system over a finite group G . For a subgroup G' of G , define*

$$\mathcal{P}|_{G'} := \{H \in \mathcal{P} : H \subseteq G'\},$$

which is a subgroup system over G' , called the restriction of \mathcal{P} to G' .

Let $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ be a \mathcal{P} -collection. For $H \in \mathcal{P}|_{G'}$, regard $H \backslash G'$ as a subset of $H \backslash G$ in the obvious way. Then the partition C_H of $H \backslash G$ restricts to a partition

of $H \backslash G'$, denoted by $C_H|_{G'}$. Define

$$\mathcal{C}|_{G'} := \{C_H|_{G'} : H \in \mathcal{P}|_{G'}\}$$

which is a $\mathcal{P}|_{G'}$ -collection, called the restriction of \mathcal{C} to G' .

Next we show that when \mathcal{C} is \mathcal{P} -scheme, its restriction $\mathcal{C}|_{G'}$ to a subgroup G' is a $\mathcal{P}|_{G'}$ -scheme. Moreover, antisymmetry and strong antisymmetry are preserved by restriction.

Lemma 6.2. *Let \mathcal{P} be a subgroup system over a finite group G . For a subgroup G' of G and a \mathcal{P} -scheme \mathcal{C} , the restriction $\mathcal{C}|_{G'}$ is a $\mathcal{P}|_{G'}$ -scheme. Moreover, if \mathcal{C} is antisymmetric (resp. strongly antisymmetric), so is $\mathcal{C}|_{G'}$.*

Proof. We have projections $\pi_{H,H'}$ and conjugations $c_{H,g}$ defined between coset spaces $H \backslash G$ for various subgroups $H \subseteq G$. And we also have projections and conjugations between coset spaces $H \backslash G'$ for $H \subseteq G'$. We use $\pi'_{H,H'}$ and $c'_{H,g}$ for the latter maps to distinguish them from the former.

For each $H \in \mathcal{P}|_{G'}$, we have a projection $\pi_{H,G'} : H \backslash G \rightarrow G' \backslash G$. This allows us to partition $H \backslash G$ into “fibers” of $\pi_{H,G'}$, i.e., preimages of elements in $G' \backslash G$:

$$H \backslash G = \coprod_{y \in G' \backslash G} \pi_{H,G'}^{-1}(y).$$

We say $x \in H \backslash G$ is in the y -fiber if $\pi_{H,G'}(x) = y$, and y is called the index of x . Note that the subset $H \backslash G' \subseteq H \backslash G$ is precisely the y -fiber with $y = G'e \in G' \backslash G$.

Consider $H, H' \in \mathcal{P}|_{G'}$ and a map $\tau : H \backslash G \rightarrow H' \backslash G$ that is either a projection $\pi_{H,H'}$, or a conjugation $c_{H,g}$ for some $g \in G'$ satisfying $H' = gHg^{-1}$. We claim $\pi_{H,G'} = \pi_{H',G'} \circ \tau$, i.e., the map τ preserves the indices of elements. This can be checked directly: for $Hh \in H \backslash G$, we have $\pi_{H,G'}(Hh) = G'h$. If $\tau = \pi_{H,H'}$, we have $\pi_{H',G'} \circ \tau(Hh) = \pi_{H',G'}(H'h) = G'h$. And if $\tau = c_{H,g}$ with $g \in G'$, we have $\pi_{H',G'} \circ \tau(Hh) = \pi_{H',G'}(H'gh) = G'gh = G'h$. So the claim holds.

This means the map τ is also fibered over $G' \backslash G$ such that its “ y -fiber” $\tau_y := \tau|_{\pi_{H,G'}^{-1}(y)}$ maps the y -fiber of $H \backslash G$ to the y -fiber of $H' \backslash G$. Setting $y = G'e$ gives us the map $\tau_y : H \backslash G' \rightarrow H' \backslash G'$ that is either the projection $\pi'_{H,H'}$, or the conjugation $c'_{H,g}$.

From this observation it is easy to see that compatibility, invariance, or regularity of $\mathcal{C}|_{G'}$ follows from the corresponding property of \mathcal{C} : fix $y = G'e$. Assume to the

contrary that $\mathcal{C}|_{G'}$ does not satisfy compatibility. Then some projection $\tau_y = \pi'_{H,H'}$ maps two elements in the same block of $C_H|_{G'}$ into different blocks of $C_{H'}|_{G'}$. But then $\tau = \pi_{H,H'}$ also maps these two elements that are in the same block of C_H into different blocks of $C_{H'}$, contradicting compatibility of \mathcal{C} . Invariance is proved in the same way except that we consider conjugations instead of projections. For regularity, note that for each projection $\pi'_{H,H'} : H \setminus G' \rightarrow H' \setminus G'$ and blocks $B \in C_H|_{G'}$, $B' \in C_{H'}|_{G'}$, we have $B = \tilde{B} \cap (H \setminus G')$, $B' = \tilde{B}' \cap (H' \setminus G')$ where $\tilde{B} \in C_H$, $\tilde{B}' \in C_{H'}$. And for $z \in B'$ we have $\pi'^{-1}_{H,H'}(z) \cap B = \pi^{-1}_{H,H'}(z) \cap \tilde{B} \cap (H \setminus G') = \pi^{-1}_{H,H'}(z) \cap \tilde{B}$. Regularity of $\mathcal{C}|_{G'}$ then follows from regularity of \mathcal{C} .

Now assume $\mathcal{C}|_{G'}$ is not antisymmetric. Then for some $H \in \mathcal{P}|_{G'}$ and $g \in N_{G'}(H)$, the map $c'_{H,g}$ restricts to a nontrivial permutation of some block $B \in C_H|_{G'}$. Then we have $g \in N_G(H)$ and $c_{H,g}$ restricts to a nontrivial permutation of \tilde{B} , where \tilde{B} is the block of C_H satisfying $\tilde{B} \cap (H \setminus G') = B$. So \mathcal{C} is not antisymmetric.

Finally, assume $\mathcal{C}|_{G'}$ is not strongly antisymmetric. Then there exists a nontrivial permutation τ of a block $B \in C_H|_{G'}$ for some subgroup $H \in \mathcal{P}|_{G'}$ such that τ is a composition of maps $\sigma_i : B_{i-1} \rightarrow B_i$, $i = 1 \dots, k$, where each B_i is a block of $C_{H_i}|_{G'}$, $H_i \in \mathcal{P}|_{G'}$, and σ_i is of the form $c'_{H_{i-1},g}|_{B_{i-1}}$ (where $g \in G'$), $\pi'_{H_{i-1},H_i}|_{B_{i-1}}$, or $(\pi'_{H_i,H_{i-1}}|_{B_i})^{-1}$ (see Definition 2.7). Each block B_i is of the form $\tilde{B}_i \cap (H_i \setminus G')$ for some $\tilde{B}_i \in C_{H_i}$. In the case that σ_i is of the form $(\pi'_{H_i,H_{i-1}}|_{B_i})^{-1}$, we know $|\pi'^{-1}_{H_i,H_{i-1}}(z) \cap B_i| = |\pi^{-1}_{H_i,H_{i-1}}(z) \cap \tilde{B}_i| = 1$ for all $z \in B_{i-1}$. So $\pi_{H_i,H_{i-1}}|_{B_i}$ is well defined. Then $\tau = \tilde{\tau}|_B$ for the nontrivial permutation $\tau = \sigma_k \cdots \circ \sigma_1$ of the block $\tilde{B} = \tilde{B}_0 \in C_H$, where each map $\tilde{\sigma}_i$ is of the form $c_{H_{i-1},g}|_{\tilde{B}_{i-1}}$, $\pi_{H_{i-1},H_i}|_{\tilde{B}_{i-1}}$, or $(\pi_{H_i,H_{i-1}}|_{\tilde{B}_i})^{-1}$. So \mathcal{C} is not strongly antisymmetric. \square

Next we describe the analogue of Definition 6.2 for m -schemes.

Definition 6.3. Let $\Pi = \{P_1, \dots, P_m\}$ be an m -scheme on a finite set S . For $(x_1, \dots, x_k) \in S^{(k)}$ where $k < m$, define the $(m - k)$ -collection

$$\Pi|_{x_1, \dots, x_k} := \{P'_1, \dots, P'_{m-k}\}$$

on the set $S - \{x_1, \dots, x_k\}$, where P'_i is the partition of $S^{(i)}$ such that two elements $(y_1, \dots, y_i), (y'_1, \dots, y'_i)$ are in the same block of $S^{(i)}$ iff $(x_1, \dots, x_k, y_1, \dots, y_i)$ and $(x_1, \dots, x_k, y'_1, \dots, y'_i)$ are in the same block of $S^{(i+k)}$.

We also have the analogue of Lemma 6.2 for m -schemes.

Lemma 6.3. *The $(m-k)$ -collection $\Pi|_{x_1, \dots, x_k}$ in Definition 6.2 is an $(m-k)$ -scheme. Moreover, if Π is antisymmetric (resp. strongly antisymmetric), so is $\Pi|_{x_1, \dots, x_k}$. And if Π does not have a matching, neither does $\Pi|_{x_1, \dots, x_k}$.*

The proof is straightforward by definition. Indeed, if we view Π as a \mathcal{P} -scheme via Definition 2.12 and Definition 2.13, where \mathcal{P} is the system of stabilizers of depth m with respect to the natural action of $G = \text{Sym}(S)$ on S . Then $\Pi|_{x_1, \dots, x_k}$ is simply the restriction of this \mathcal{P} -scheme to the subgroup G_{x_1, \dots, x_k} . We leave the details to the reader.

Passing to quotient groups. Let G be a finite group and let N be a normal in G . Write \bar{G} for G/N and ϕ for the quotient map $G \rightarrow \bar{G}$.

For a subgroup $H \subseteq \bar{G}$, the group G acts on $H \backslash \bar{G}$ by inverse right translation (through its quotient group \bar{G}). The stabilizer of $He \in H \backslash \bar{G}$ is $\phi^{-1}(H)$. So by Lemma 2.1, we have an equivalence between the action of G on $H \backslash \bar{G}$ and that on $\phi^{-1}(H) \backslash G$, given by the bijection $\lambda_{He} : H \backslash \bar{G} \rightarrow \phi^{-1}(H) \backslash G$ sending $H\phi(g)$ to $\phi^{-1}(H)g$ for $g \in G$.

Let \mathcal{P} be a subgroup system over \bar{G} . Define $\tilde{\mathcal{P}} = \{\phi^{-1}(H) : H \in \mathcal{P}\}$, which is a subgroup system over G . By identifying $H \backslash \bar{G}$ with $\phi^{-1}(H) \backslash G$ via λ_{He} for $H \in \mathcal{P}$, we see that a \mathcal{P} -scheme over \bar{G} is equivalent to a $\tilde{\mathcal{P}}$ -scheme over G . This is made formal by the following lemma.

Lemma 6.4. *Let \mathcal{P} and $\tilde{\mathcal{P}}$ be as above. For a \mathcal{P} -collection $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$, define the $\tilde{\mathcal{P}}$ -collection $\mathcal{C}' = \{C'_{\phi^{-1}(H)} : H \in \mathcal{P}\}$ by choosing*

$$C'_{\phi^{-1}(H)} = \{\lambda_{He}(B) : B \in C_H\}.$$

Then $\mathcal{C} \mapsto \mathcal{C}'$ is a one-to-one correspondence between \mathcal{P} -schemes over \bar{G} and $\tilde{\mathcal{P}}$ -schemes over G . Moreover, \mathcal{C} is antisymmetric (resp. strongly antisymmetric) iff \mathcal{C}' is antisymmetric (resp. strongly antisymmetric). And \mathcal{C} is homogeneous (resp. discrete) on a subgroup $H \in \mathcal{P}$ iff \mathcal{C}' is homogeneous (resp. discrete) on $\phi^{-1}(H)$.

Proof. We check that the maps λ_{He} commute with conjugations and projections: write $\pi_{H, H'}$ and $c_{H, g}$ for conjugations and projections between coset spaces of \bar{G} and write $\pi'_{H, H'}$ and $c'_{H, g}$ for those between coset spaces of G . Then we always have

$$\lambda_{H'e} \circ \pi_{H, H'} = \pi'_{\phi^{-1}(H), \phi^{-1}(H')} \circ \lambda_{He}$$

for $H, H' \in \mathcal{P}$, $H \subseteq H'$, and

$$\lambda_{H'e} \circ c_{H, \phi(g)} = c'_{\phi^{-1}(H), g} \circ \lambda_{He}.$$

for $H, H' \in \mathcal{P}$, $g \in G$, $H' = \phi(g)H\phi(g)^{-1}$. Also note that the maps λ_{He} are bijections. The lemma then follows easily by definition. \square

We conclude this section by proving Lemma 4.12 using the results developed above. First we prove the following lemma.

Lemma 6.5. *Let $k \in \mathbb{N}^+$ and $G_k \subseteq G_{k-1} \subseteq \cdots \subseteq G_1 \subseteq G_0$ be a chain of finite groups. Let \mathcal{P} be a subgroup system over G_0 . We have:*

1. *If for all $i \in [k]$, all strongly antisymmetric $\mathcal{P}|_{G_{i-1}}$ -schemes are discrete on G_i , then all strongly antisymmetric \mathcal{P} -schemes are discrete on G_k .*
2. *If for some $i \in [k]$, all strongly antisymmetric $\mathcal{P}|_{G_{i-1}}$ -schemes are inhomogeneous on G_i , then all strongly antisymmetric \mathcal{P} -schemes are inhomogeneous on G_k .*

The same holds if strong antisymmetry is replaced by antisymmetry.

Proof. Assume that there exists a strongly antisymmetric \mathcal{P} -scheme $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ that is not discrete on G_k . Then there exist two different elements $x, x' \in G_k \setminus G$ lying in the same block of C_{G_k} . Pick the greatest integer $i \in [k]$ satisfying $\pi_{G_k, G_{i-1}}(x) = \pi_{G_k, G_{i-1}}(x')$. Such i exists as $\pi_{G_k, G_0}(x) = \pi_{G_k, G_0}(x')$. Let $y = \pi_{G_k, G_i}(x)$ and $y' = \pi_{G_k, G_i}(x')$. Then (1) $y \neq y'$ by maximality of i and the fact that $x \neq x'$, (2) y, y' are in the same block of C_{G_i} by compatibility of \mathcal{C} and the fact that x, x' are in the same block of C_{G_k} , and (3) $\pi_{G_i, G_{i-1}}(y) = \pi_{G_i, G_{i-1}}(y')$ since $\pi_{G_i, G_{i-1}}(y) = \pi_{G_k, G_{i-1}}(x)$ and $\pi_{G_i, G_{i-1}}(y') = \pi_{G_k, G_{i-1}}(x')$.

Suppose $\pi_{G_i, G_{i-1}}(y) = \pi_{G_i, G_{i-1}}(y') = G_{i-1}g$. By replacing \mathcal{C} with ${}^g\mathcal{C}$ (with respect to the action of G_k on the set of \mathcal{P} -schemes by inverse right translation) and applying Lemma 6.1, we may assume $G_{i-1}g = G_{i-1}e$. Then we can write $y = G_i h$ and $y' = G_i h'$ for some $h, h' \in G_{i-1}$. By Lemma 6.2, the restriction $\mathcal{C}|_{G_{i-1}} = \{C_H|_{G_{i-1}} : H \in \mathcal{P}|_{G_{i-1}}\}$ is a strongly antisymmetric $\mathcal{P}|_{G_{i-1}}$ -scheme. As y, y' are in the same block of C_{G_i} , they are also in the same block of $C_{G_i}|_{G_{i-1}}$. As $y \neq y'$, we know $\mathcal{C}|_{G_{i-1}}$ is not discrete on G_i . This proves the first claim of the lemma.

For the second claim, assume to the contrary that it does not hold. Choose $i \in [k]$ such all strongly antisymmetric $\mathcal{P}|_{G_{i-1}}$ -schemes are inhomogeneous on G_i . Let $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ be a strongly antisymmetric \mathcal{P} -scheme that is homogeneous on G_k . By compatibility, we know \mathcal{C} is homogeneous on G_i . Then $\mathcal{C}|_{G_{i-1}}$ is also homogeneous on G_i . It is also strongly antisymmetric by Lemma 6.2, which contradicts the assumption.

The proof for antisymmetry is the same. \square

Now we are ready to prove Lemma 4.12. For convenience, we restate the lemma.

Lemma 6.6. *Let $k \in \mathbb{N}^+$ and $G_k \subseteq G_{k-1} \subseteq \cdots \subseteq G_1 \subseteq G_0$ be a chain of finite groups. For $i \in [k]$, let N_i be a subgroup of G_i that is normal in G_{i-1} , $\pi_i : G_{i-1} \rightarrow G_{i-1}/N_i$ be the corresponding quotient map, and \mathcal{P}_i be a subgroup system over G_{i-1}/N_i that contains G_i/N_i . Define*

$$\mathcal{P} = \{g\pi_i^{-1}(H)g^{-1} : 1 \leq i \leq k, H \in \mathcal{P}_i, g \in G_0\},$$

which is a subgroup system over G_0 and contains $\pi_i^{-1}(G_i/N_i) = G_i$ for all $i \in [k]$. Then we have

1. *If for all $i \in [k]$, all strongly antisymmetric \mathcal{P}_i -schemes are discrete on G_i/N_i , then all strongly antisymmetric \mathcal{P} -schemes are discrete on G_k .*
2. *If for some $i \in [k]$, all strongly antisymmetric \mathcal{P}_i -schemes are inhomogeneous on G_i/N_i , then all strongly antisymmetric \mathcal{P} -schemes are inhomogeneous on G_k .*

The same holds if strong antisymmetry is replaced by antisymmetry.

Proof. Fix $i \in [k]$. By Lemma 6.4 and the definition of \mathcal{P} , if all strongly antisymmetric \mathcal{P}_i -schemes are discrete (resp. inhomogeneous) on G_i/N_i , then all strongly antisymmetric $\mathcal{P}|_{G_{i-1}}$ -schemes are discrete (resp. inhomogeneous) on G_i . The same holds if strong antisymmetry is replaced by antisymmetry. The lemma now follows from Lemma 6.5. \square

6.2 Induction of \mathcal{P} -schemes

Let G be a finite group and let G' be a subgroup of G . Let \mathcal{P} be a subgroup system over G and let

$$\mathcal{P}' = \{G' \cap gHg^{-1} : H \in \mathcal{P}, g \in G\},$$

which is a subgroup system over G' . In this section, we show that every \mathcal{P}' -scheme induces a \mathcal{P} -scheme in a way that preserves antisymmetry and strong antisymmetry. To achieve it, we need the following lemma.

Lemma 6.7. *Given $g_1, \dots, g_k \in G$ such that $\{g_1^{-1}, \dots, g_k^{-1}\}$ is a complete set of representatives of $H \backslash G / G'$, there exists a bijection*

$$\phi : \prod_{i=1}^k (G' \cap g_i H g_i^{-1}) \backslash G' \rightarrow H \backslash G$$

defined as follows: For $g \in G$, define the map

$$\phi_{H,g} : (G' \cap g H g^{-1}) \backslash G' \rightarrow H \backslash G$$

sending $(G' \cap g H g^{-1})h$ to $H g^{-1}h$ for $h \in G'$. The maps $\phi_{H,g}$ are well defined. For $i \in [k]$, the restriction of ϕ to $(G' \cap g_i H g_i^{-1}) \backslash G'$ is ϕ_{H,g_i} .

Proof. Consider the action of G' on $H \backslash G$ by inverse right translation. For $i \in [k]$, let $O_i = \{H g_i^{-1} g^{-1} : g \in G'\}$ be the G' -orbits of $H g_i^{-1} \in H \backslash G$. Then $\{O_1, \dots, O_k\}$ is the partition of $H \backslash G$ into the G' -orbits, i.e., $H \backslash G = \coprod_{i=1}^k O_i$. Fix $i \in [k]$. The stabilizer of $H g_i^{-1}$ is $G' \cap g_i H g_i^{-1}$. So by Lemma 2.1, we have an equivalence of actions of G'

$$\lambda_{H g_i^{-1}} : O_i \rightarrow (G' \cap g_i H g_i^{-1}) \backslash G'$$

sending ${}^h(H g_i^{-1}) = H g_i^{-1} h^{-1}$ to $(G' \cap g_i H g_i^{-1}) h^{-1}$ for $h \in G'$. The inverse of this map is exactly ϕ_{H,g_i} . As we are allowed to choose g_i to be any $g \in G$, all the maps $\phi_{H,g}$ are well defined. \square

For each $H \in \mathcal{P}$, the subgroups $G' \cap g H g^{-1}$ are in \mathcal{P}' for all $g \in G$. By Lemma 6.7, we can combine partitions of $G' \cap g_i H g_i^{-1} \backslash G'$, $i = 1, \dots, k$, into a partition of $H \backslash G$. This leads to the following definition.

Definition 6.4 (induction). *Let G, G', \mathcal{P} and \mathcal{P}' be as above. Let $\mathcal{C}' = \{C'_H : H \in \mathcal{P}'\}$ be a \mathcal{P}' -scheme. For $H \in \mathcal{P}$, choose $g_1, \dots, g_k \in G$ such that $\{g_1^{-1}, \dots, g_k^{-1}\}$ is a complete set of representatives of $H \backslash G / G'$. Define the partition C_H of $H \backslash G$ by*

$$C_H = \left\{ \phi_{H,g_i}(B) : i \in [k], B \in C'_{G' \cap g_i H g_i^{-1}} \right\},$$

where the maps ϕ_{H,g_i} are as in Lemma 6.7, i.e., each ϕ_{H,g_i} sends $(G' \cap g_i H g_i^{-1})h$ to $H g_i^{-1}h$ for $h \in G'$. Define the \mathcal{P} -collection $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$, called the induction of \mathcal{C}' to \mathcal{P} .

The \mathcal{P} -collection \mathcal{C} constructed as above is indeed a \mathcal{P} -scheme:

Theorem 6.1. *The \mathcal{P} -collection \mathcal{C} in Definition 6.4 is a well defined \mathcal{P} -scheme, which does not depend on the choices of the elements g_i . Moreover, if \mathcal{C}' is antisymmetric (resp. strongly antisymmetric), so is \mathcal{C} .*

Proof. Fix $H \in \mathcal{P}$. It follows from Lemma 6.7 that C_H is indeed a partition of $H \backslash G$. We need to show that C_H is independent of the choices of the elements g_1, \dots, g_k . So consider $g'_1, \dots, g'_k \in G$ such that $\{g'_1, \dots, g'_k\}$ is a complete set of representatives of $H \backslash G/G'$ as well. We want to show

$$C_H = \left\{ \phi_{H, g'_i}(B) : i \in [k], B \in C'_{G' \cap g'_i H g'^{-1}_i} \right\}.$$

As the right hand side is also a partition of $H \backslash G$, it suffices to show that $\phi_{H, g'_i}(B) \in C_H$ for $i \in [k]$ and $B \in C'_{G' \cap g'_i H g'^{-1}_i}$. Fix i and B . Choose $j \in [k]$ such that $H g_j^{-1} G' = H g'_i G'$. And choose $g \in G'$ such that $H g_j^{-1} = H g'_i g^{-1}$. We have the conjugation

$$c_{G' \cap g'_i H g'^{-1}_i, g} : (G' \cap g'_i H g'^{-1}_i) \backslash G' \rightarrow (G' \cap g_j H g_j^{-1}) \backslash G'$$

sending $(G' \cap g'_i H g'^{-1}_i)h$ to $(G' \cap g_j H g_j^{-1})gh$ for $h \in G'$. By invariance of \mathcal{C}' , the set $c_{G' \cap g'_i H g'^{-1}_i, g}(B)$ is a block of $C'_{G' \cap g_j H g_j^{-1}}$. So $\phi_{H, g_j} \circ c_{G' \cap g'_i H g'^{-1}_i, g}(B)$ is a block of C_H . We claim

$$\phi_{H, g_j} \circ c_{G' \cap g'_i H g'^{-1}_i, g} = \phi_{H, g'_i},$$

which holds since

$$\begin{aligned} \phi_{H, g_j} \circ c_{G' \cap g'_i H g'^{-1}_i, g}((G' \cap g'_i H g'^{-1}_i)h) &= \phi_{H, g_j}((G' \cap g_j H g_j^{-1})gh) = H g_j^{-1}gh \\ &= H g'^{-1}_i h = \phi_{H, g'_i}((G' \cap g'_i H g'^{-1}_i)h) \end{aligned}$$

for $h \in G'$. It follows that $\phi_{H, g'_i}(B) \in C_H$, as desired. So C_H does not depend on the choices of g_1, \dots, g_k .

Next we prove that \mathcal{C} is a \mathcal{P} -scheme. To prove compatibility, consider $H, H' \in \mathcal{P}$ satisfying $H \subseteq H'$. For $g \in G$, the following diagram commutes:

$$\begin{array}{ccc} (G' \cap g H g^{-1}) \backslash G' & \xrightarrow{\pi_{G' \cap g H g^{-1}, G' \cap g H' g^{-1}}} & (G' \cap g H' g^{-1}) \backslash G' \\ \phi_{H, g} \downarrow & & \downarrow \phi_{H', g} \\ H \backslash G & \xrightarrow{\pi_{H, H'}} & H' \backslash G. \end{array}$$

For $B \in C_H$, we want to show that $\pi_{H,H'}(B)$ is contained in a block of $C_{H'}$. Choose $g \in G$ and $\tilde{B} \in C'_{G' \cap gHg^{-1}}$ such that $B = \phi_{H,g}(\tilde{B})$. Such g and \tilde{B} exist by the definition of C_H and the fact that this definition does not depend on the choices of g_1, \dots, g_k (in particular we may choose $g_1 = g$). Then

$$\pi_{H,H'}(B) = \pi_{H,H'} \circ \phi_{H,g}(\tilde{B}) = \phi_{H',g}(y) \circ \pi_{G' \cap gHg^{-1}, G' \cap gH'g^{-1}}(\tilde{B}).$$

Here $\pi_{G' \cap gHg^{-1}, G' \cap gH'g^{-1}}(\tilde{B})$ is contained in a block of $C'_{G' \cap gH'g^{-1}}$ by compatibility of C' , and hence $\pi_{H,H'}(B)$ is contained in a block of $C_{H'}$. It follows that \mathcal{C} is compatible.

For regularity, consider H, H' as above and $B \in C_H$. Choose $B' \in C_{H'}$ containing $\pi_{H,H'}(B)$. We claim that $\pi_{H,H'}|_B : B \rightarrow B'$ has constant degree, i.e., the number of preimages $|(\pi_{H,H'}|_B)^{-1}(y)|$ is independent of the choices of $y \in B'$. Choose $g \in G$ and $\tilde{B} \in C'_{G' \cap gHg^{-1}}$ such that $B = \phi_{H,g}(\tilde{B})$. Let $\tilde{B}' = \pi_{G' \cap gHg^{-1}, G' \cap gH'g^{-1}}(\tilde{B})$. Then $B' = \phi_{H',g}(\tilde{B}')$. By regularity of C' , the map $\pi_{G' \cap gHg^{-1}, G' \cap gH'g^{-1}}|_{\tilde{B}} : \tilde{B} \rightarrow \tilde{B}'$ has constant degree. The claim follows by noting that $\phi_{H,g}|_{\tilde{B}} : \tilde{B} \rightarrow B$ and $\phi_{H',g}|_{\tilde{B}'} : \tilde{B}' \rightarrow B'$ are bijective. So \mathcal{C} is regular.

For invariance, consider $H, H' \in \mathcal{P}$ and $h \in G$ satisfying $H' = hHh^{-1}$. For $g \in G$, we have $G' \cap gHg^{-1} = G' \cap gh^{-1}H'(gh^{-1})^{-1}$, and the following diagram commutes

$$\begin{array}{ccc} (G' \cap gHg^{-1}) \backslash G' & \xrightarrow{\text{id}} & (G' \cap gh^{-1}H'(gh^{-1})^{-1}) \backslash G' \\ \phi_{H,g} \downarrow & & \downarrow \phi_{H',gh^{-1}} \\ H \backslash G & \xrightarrow{c_{H,h}} & H' \backslash G, \end{array}$$

where id denotes the identity map. It follows that $c_{H,h}$ maps blocks of C_H to blocks of $C_{H'}$. So \mathcal{C} is invariant.

Now assume \mathcal{C} is not strongly antisymmetric and we prove that C' is not either. By definition, there exists a nontrivial permutation $\tau = \sigma_k \circ \dots \circ \sigma_1$ of a block $B \in C_H$ for some $H \in \mathcal{P}$ such that each $\sigma_i : B_{i-1} \rightarrow B_i$ is a map of the form $\pi_{H_{i-1}, H_i}|_{B_{i-1}}$, $(\pi_{H_i, H_{i-1}}|_{B_i})^{-1}$, or $c_{H_{i-1}, h}|_{B_{i-1}}$, and $B_i \in C_{H_i}$, $H_i \in \mathcal{P}$, $B = B_0 = B_k$, $H = H_0 = H_k$ (see Definition 2.7).

By the two diagrams above, we can choose $g_i \in G$ and $\tilde{B}_i \in C'_{G' \cap g_i H g_i^{-1}}$ for $0 \leq i \leq k$, and choose $\tilde{\sigma}_i : \tilde{B}_{i-1} \rightarrow \tilde{B}_i$ of the form $\pi_{G' \cap g_{i-1} H_{i-1} g_{i-1}^{-1}, G' \cap g_i H_i g_i^{-1}}|_{\tilde{B}_{i-1}}$, $(\pi_{G' \cap g_i H_i g_i^{-1}, G' \cap g_{i-1} H_{i-1} g_{i-1}^{-1}}|_{\tilde{B}_i})^{-1}$, or the identity map on \tilde{B}_i for $i \in [k]$, such that $\phi_{H_i, g_i}(\tilde{B}_i) = B_i$ and $\sigma_i \circ \phi_{H_{i-1}, g_{i-1}}|_{\tilde{B}_{i-1}} = \phi_{H_i, g_i}|_{\tilde{B}_i} \circ \tilde{\sigma}_i$ for $i \in [k]$.¹ Define

¹For the case that $\sigma_i = (\pi_{H_i, H_{i-1}}|_{B_i})^{-1}$, we choose $\tilde{\sigma}_i = (\pi_{G' \cap g_i H_i g_i^{-1}, G' \cap g_{i-1} H_{i-1} g_{i-1}^{-1}}|_{\tilde{B}_i})^{-1}$, which is well defined since $\phi_{H_{i-1}, g_{i-1}}$ and ϕ_{H_i, g_i} are bijective.

$\tilde{\tau} := \tilde{\sigma}_k \circ \cdots \circ \tilde{\sigma}_1$ which is a map from \tilde{B}_0 to \tilde{B}_k . Then the following diagram commutes.

$$\begin{array}{ccc} \tilde{B}_0 & \xrightarrow{\tilde{\tau}} & \tilde{B}_k \\ \phi_{H,g_0}|_{\tilde{B}_0} \downarrow & & \downarrow \phi_{H,g_k}|_{\tilde{B}_k} \\ B & \xrightarrow{\tau} & B \end{array}$$

We have $Hg_0^{-1}G' = Hg_k^{-1}G'$, since otherwise the image of ϕ_{H,g_0} and that of ϕ_{H,g_k} would be disjoint (see Lemma 6.7). So $Hg_0^{-1} = Hg_k^{-1}g^{-1}$ for some $g \in G'$. The first part of the proof shows that $\phi_{H,g_0} \circ c_{G' \cap g_k Hg_k^{-1}, g} = \phi_{H,g_k}$. By composing $\tilde{\tau}$ with $c_{G' \cap g_k Hg_k^{-1}, g}$, we may assume $g_k = g_0$ and $\tilde{B}_k = \tilde{B}_0$. Then as τ is a nontrivial permutation of B and $\phi_{H,g_0}|_{\tilde{B}_0} : \tilde{B}_0 \rightarrow B$ is bijective, we know $\tilde{\tau}$ is a nontrivial permutation of \tilde{B}_0 . So \mathcal{C} is not strongly antisymmetric.

The proof for antisymmetry is the same except that we only consider maps τ that are conjugations. \square

Corollary 6.1. *Let $G, G', \mathcal{P}, \mathcal{P}'$ be as above and let H be a subgroup in \mathcal{P} .*

1. *Suppose all antisymmetric \mathcal{P} -schemes are discrete on H . Then all antisymmetric \mathcal{P}' -schemes are discrete on $G' \cap gHg^{-1}$ for all $g \in G$.*
2. *Suppose all antisymmetric \mathcal{P} -schemes are inhomogeneous on H , and G' acts transitively on $H \backslash G$ by inverse right translation. Then all antisymmetric \mathcal{P}' -schemes are inhomogeneous on $G' \cap gHg^{-1}$ for all $g \in G$.*

The same claims hold if antisymmetry is replaced with strong antisymmetry.

Proof. We prove the claims by contrapositive. For the first claim, suppose $\mathcal{C}' = \{C'_{H'} : H' \in \mathcal{P}'\}$ is an antisymmetric \mathcal{P}' -scheme that is not discrete on $G' \cap gHg^{-1}$ for some $g \in G$. Choose $B \in C'_{G' \cap gHg^{-1}}$ that is not a singleton. By Theorem 6.4, the induced \mathcal{P} -scheme $\mathcal{C} = \{C_{H'} : H' \in \mathcal{P}\}$ is antisymmetric. Moreover, we know \mathcal{C} is not discrete on H since the block $\phi_{H,g}(B) \in C_H$ is not a singleton.

For the second claim, suppose $\mathcal{C}' = \{C'_{H'} : H' \in \mathcal{P}'\}$ is an antisymmetric \mathcal{P}' -scheme that is homogeneous on $G' \cap gHg^{-1}$ for some $g \in G$. By Theorem 6.4, the induced \mathcal{P} -scheme $\mathcal{C} = \{C_{H'} : H' \in \mathcal{P}\}$ is antisymmetric. As G' acts transitively on $H \backslash G$, the double coset space $H \backslash G / G'$ has only one double coset $Hg^{-1}G'$, which implies that $\phi_{H,g} : G' \cap gHg^{-1} \backslash G \rightarrow H \backslash G$ is surjective. As \mathcal{C}' is homogeneous on $G' \cap gHg^{-1}$, we know \mathcal{C} is homogeneous on H .

The proof for strong antisymmetry is the same. \square

Now let S be a finite G -set and let G' be a subgroup of G . Fix $m \in \mathbb{N}^+$ and let $\mathcal{P} = \{G_T : 1 \leq |T| \leq m\}$ be the system of stabilizers of depth m with respect to the action of G on S . Note that for $T \subseteq S$, we have $G' \cap gG_Tg^{-1} = G' \cap G_{gT} = G'_{gT}$. So $\mathcal{P}' = \{G' \cap gHg^{-1} : H \in \mathcal{P}\}$ is exactly the system of stabilizers of depth m with respect to the action of G' on S restricted from that of G . Therefore we have:

Corollary 6.2. *Let G be a finite group acting on a finite set S , G' a subgroup of G , and $m \in \mathbb{N}^+$. Let \mathcal{P} (resp. \mathcal{P}') be the system of stabilizers of depth m over G (resp. G') with respect to the action of G (resp. G') on S .*

1. *Suppose all antisymmetric \mathcal{P} -schemes are discrete on G_x for all $x \in S$. Then all antisymmetric \mathcal{P}' -schemes are discrete on G'_x for all $x \in S$.*
2. *Suppose all antisymmetric \mathcal{P} -schemes are inhomogeneous on G_{x_0} for some $x_0 \in S$, and G' acts transitively on S . Then all antisymmetric \mathcal{P}' -schemes are inhomogeneous on G'_x for all $x \in S$.*

The same claims hold if antisymmetry is replaced with strong antisymmetry.

In particular, we see $d(G)$ and $d'(G)$ (cf. Definition 2.8) are monotone with respect to inclusion of permutation groups:

Corollary 6.3. *Let G be a finite permutation group on a finite set S , and let G' be a subgroup of G on S . Then $d(G') \leq d(G)$ and $d'(G') \leq d'(G)$.*

We also mention the following variant of Corollary 6.2, which allows $G' \subseteq G$ to act on a proper subset of S .

Corollary 6.4. *Let G be a finite group acting on a finite set S , G' a subgroup of G , and $m \in \mathbb{N}^+$. Let T a subset of S such that the action of G' on S fixes T setwisely and $S - T$ pointwisely. Let \mathcal{P} (resp. \mathcal{P}') be the system of stabilizers of depth m over G (resp. G') with respect to the action of G (resp. G') on S (resp. T). Suppose all antisymmetric \mathcal{P} -schemes are discrete on G_x for all $x \in S$. Then all antisymmetric \mathcal{P}' -schemes are discrete on G'_x for all $x \in T$. The same claims hold if antisymmetry is replaced with strong antisymmetry.*

Proof. If $S = T$, the claim holds by Corollary 6.2. So assume $S \neq T$. Let \mathcal{P}'' be the system of stabilizers of depth m over G' with respect to the action of G' on S . Then $\mathcal{P}'' = \mathcal{P}' \cup \{G\}$. A \mathcal{P}' -scheme \mathcal{C} always extends to a \mathcal{P}'' -scheme

$\mathcal{C}' := \mathcal{C} \cup \{C_G\}$, where C_G is the only partition of the singleton $G \setminus G$, and such an extension clearly preserves antisymmetry and strong antisymmetry. The claim then follows from Corollary 6.2. \square

6.3 Schemes conjectures

We investigate the following conjecture proposed in [IKS09], known as the *schemes conjecture*.

Conjecture 6.1 (schemes conjecture). *There exists a constant $m \in \mathbb{N}^+$ such that every antisymmetric homogeneous m -scheme on a finite set S where $|S| > 1$ has a matching.*

It was shown in [IKS09] that this conjecture is true for orbit m -schemes with $m = 4$. We improve this result in Section 6.6 by showing that one can even choose $m = 3$. For general m -schemes, antisymmetric homogeneous m -schemes with no matching do exist for $m = 1, 2, 3$ (see Section 2.5) but no counterexamples are known for $m \geq 4$.

The following theorem was proved in [IKS09].

Theorem 6.2. *Assuming GRH and the schemes conjecture, there exists a deterministic polynomial-time algorithm that computes the complete factorization of a given polynomial $f(X) \in \mathbb{F}_q[X]$ over a finite field \mathbb{F}_q .*

We reprove this theorem using the machinery of \mathcal{P} -schemes. First note that by Lemma 2.10, an m -scheme with a matching is not strongly antisymmetric. So we can replace the schemes conjecture by the following variant, which is implied by the original one.

Conjecture 6.2. *There exists a constant $m \in \mathbb{N}^+$ such that every strongly antisymmetric m -scheme on a finite set S where $|S| > 1$ is inhomogeneous.*

We also need the following simple lemma whose proof is deferred to Section 6.5. It shows that inhomogeneity in Conjecture 6.2 can be replaced by discreteness.

Lemma 6.8. *Suppose there exists a strongly antisymmetric m -scheme on a finite set S that is not discrete, where $m \in \mathbb{N}^+$ and $|S| > 1$. Then for some finite set T satisfying $1 < |T| \leq |S|$, there exists a strongly antisymmetric homogeneous m -scheme on T .*

Now we complete the proof of Theorem 6.2.

Proof of Theorem 6.2. First assume that $\mathbb{F}_q = \mathbb{F}_p$ is a prime field and f is square-free and completely reducible over \mathbb{F}_p . Fix the constant $m \in \mathbb{N}^+$ as guaranteed by the schemes conjecture, and let $n = \deg(f)$. The algorithm first lifts f to $\tilde{f}(X) \in \mathbb{Z}[X]$ of degree n such that all coefficients of \tilde{f} are between zero and p . We can assume \tilde{f} is irreducible over \mathbb{Q} using the factoring algorithm for rational polynomials [LLL82]. Let S be the set of roots of \tilde{f} in its splitting field. The Galois group $\text{Gal}(\tilde{f}/\mathbb{Q})$ of \tilde{f} is then a permutation group on S .

Run the \mathcal{P} -scheme algorithm in Chapter 3 that we used to prove Corollary 3.2. By Corollary 3.2, it suffices to prove $d(\text{Gal}(\tilde{f}/\mathbb{Q})) \leq m$. Assume to the contrary that $d(\text{Gal}(\tilde{f}/\mathbb{Q})) > m$. By Corollary 6.3, we have $d(\text{Sym}(S)) > m$, where $\text{Sym}(S)$ acts naturally on S . Then by Lemma 2.7, there exists a strongly antisymmetric non-discrete m -scheme on S . By Lemma 6.8, for some finite set T satisfying $|T| > 1$, there exists a strongly antisymmetric homogeneous m -scheme on T . But this is a contradiction to Conjecture 6.2 and hence to the schemes conjecture.

For general f and \mathbb{F}_q , we either reduce to the previous case using Berlekamp's reduction [Ber70] and square-free factorization [Yun76; Knu98], or run the generalized \mathcal{P} -scheme algorithm in Chapter 5 and apply Corollary 5.2 instead. \square

Schemes conjectures for a family of permutation groups. In the proof of Theorem 6.2, we reduce to the case of the full symmetric group $\text{Sym}(S)$ and then apply the schemes conjecture. On the other hand, if the Galois group G is “less complex” than $\text{Sym}(S)$, we expect that the schemes conjecture can be replaced with a more moderate assumption. Formalizing this intuition leads to a hierarchy of conjectures, which we explain now.

Let \mathcal{G} be a family of finite permutation groups. We formulate a conjecture for \mathcal{G} as follows.

Conjecture 6.3 (schemes conjecture for \mathcal{G}). *There exists a constant $m \in \mathbb{N}^+$ such that $d(G) \leq m$ for all $G \in \mathcal{G}$.*

By Corollary 3.2 and Corollary 5.2, assuming this conjecture (and GRH) guarantees a polynomial-time factoring algorithm for the case that the Galois group G is in \mathcal{G} as a permutation group:

Theorem 6.3. *Assuming GRH and the schemes conjecture for \mathcal{G} , there exists a deterministic polynomial-time algorithm that given a polynomial $f(X) \in \mathbb{F}_q[X]$ and an irreducible² lifted polynomial \tilde{f} of f , computes the complete factorization of f over \mathbb{F}_q , provided that the Galois group of \tilde{f} , as a permutation group on the set of roots of \tilde{f} , is permutation isomorphic to some group in \mathcal{G} .*

There exist reductions among these schemes conjectures defined for various families \mathcal{G} . To formulate them, we need the following notation: for two families \mathcal{G} and \mathcal{G}' , write $\mathcal{G} \preceq \mathcal{G}'$ if any permutation group $G \in \mathcal{G}$ is permutation isomorphic to a subgroup of some permutation group $G' \in \mathcal{G}'$ (where action of this subgroup is restricted from that of G'). Then we have

Theorem 6.4. *The schemes conjecture for \mathcal{G} is implied by that for \mathcal{G}' if $\mathcal{G} \preceq \mathcal{G}'$.*

Proof. This follows directly from Corollary 6.3. □

In particular, all these conjectures are subsumed by that for the family of symmetric groups $\{\text{Sym}(n) : n \in \mathbb{N}^+\}$, where each symmetric group $\text{Sym}(n)$ acts naturally on $[n]$. The latter is equivalent to Conjecture 6.2 by the connection between m -schemes and \mathcal{P} -schemes (see Theorem 2.1).

Therefore, the conjectures for different families of finite permutation groups form a hierarchy, partially ordered by the relation \preceq , and Conjecture 6.2 is the most difficult one. One possible approach to the schemes conjecture is first relaxing it to those for simpler permutation groups which may be easier to prove. We will prove results in the same spirit in subsequent chapters.

Finally, we note that the schemes conjecture hold for the family of primitive solvable permutation groups, or more generally for primitive permutation groups G not involving $\text{Alt}(d)$ (i.e., $\text{Alt}(d)$ is not isomorphic to a subquotient of G), where d is a constant.

Theorem 6.5. *The schemes conjecture for \mathcal{G} is true if \mathcal{G} is the family of primitive solvable permutation groups, or the family of primitive permutation groups G not involving $\text{Alt}(d)$, where $d \in \mathbb{N}^+$ is a constant.*

²The assumption that \tilde{f} is irreducible is not necessary, and can be avoided by adapting Lemma 4.10. We omit the details.

Proof. Let G be a primitive permutation group. Seress [Ser96] proved $b(G) \leq 4$ when G is solvable. More generally, it was shown in [GSS98] that there exists a function $g(\cdot)$ such that $b(G) \leq g(d)$ if G does not involve $\text{Alt}(d)$. The theorem then follows from Lemma 2.5. \square

Remark. The scheme conjectures in this section are formulated in terms of discreteness of \mathcal{P} -schemes and are used for complete factorization. One can also formulate conjectures in terms of inhomogeneity and use them for proper factorization. We leave the details to the reader. To establish reductions between these conjectures (in terms of inhomogeneity rather than discreteness), one needs to restrict to families of *transitive* permutation groups as transitivity is required in Corollary 6.2.

6.4 Extension to the closure of a subgroup system

Suppose $\mathcal{P}, \mathcal{P}'$ are subgroup systems over a finite group G and $\mathcal{P} \subseteq \mathcal{P}'$. We can construct a \mathcal{P} -scheme from a \mathcal{P}' -scheme by simply discarding the partitions of $H \setminus G$ for $H \in \mathcal{P}' - \mathcal{P}$. Conversely, we want to know if a \mathcal{P} -scheme can be extended to a \mathcal{P}' -scheme. In this section, we show that this is possible in some cases by formulating the notion of the *closure* \mathcal{P}_{cl} of a subgroup system \mathcal{P} and proving that \mathcal{P} -scheme can always be extended to a \mathcal{P}_{cl} -scheme. As an application, we prove Lemma 4.16 and Lemma 4.17 as promised before.

Definition 6.5 (closure). *Let \mathcal{P} be a subgroup system over a group G . Denote by \mathcal{P}_{cl} the set of subgroups H of G satisfying the following conditions:*

1. \mathcal{P} contains a subgroup $H' \subseteq H$, and the set of such subgroups has a unique maximal element (with respect to inclusion), denoted by $u_{\mathcal{P}}(H)$, or simply $u(H)$ when there is no confusion.
2. $u(H)$ is a normal subgroup of H .

Then \mathcal{P}_{cl} is a subgroup system³ over G containing \mathcal{P} , called the *closure* of \mathcal{P} .

The usage of the term *closure* is justified by the obvious fact $\mathcal{P} \subseteq \mathcal{P}_{\text{cl}}$ and the next lemma.

Lemma 6.9. $(\mathcal{P}_{\text{cl}})_{\text{cl}} = \mathcal{P}_{\text{cl}}$.

³It is easy to see that \mathcal{P}_{cl} is closed under conjugation in G , so it is indeed a subgroup system over G .

Proof. Consider $H \in (\mathcal{P}_{\text{cl}})_{\text{cl}}$. Write $H' = u_{\mathcal{P}_{\text{cl}}}(H)$ and $H'' = u_{\mathcal{P}}(H')$. We show that $H \in \mathcal{P}_{\text{cl}}$ and $u_{\mathcal{P}}(H) = H''$.

We first verify that H'' is normal in H . By definition, we know H' is normal in H . Then for any $g \in H$, we have

$$H'' = u_{\mathcal{P}}(H') = u_{\mathcal{P}}(gH'g^{-1}) = gu_{\mathcal{P}}(H')g^{-1} = gH''g^{-1}.$$

So H'' is normal in H .

Next we show that H'' is the unique maximal element in \mathcal{P} subject to $H'' \subseteq H$. Assume to the contrary that there exists an element $U \subsetneq H''$ in $\mathcal{P} \subseteq \mathcal{P}_{\text{cl}}$ that is a subgroup of H . As H' is the unique maximal element in \mathcal{P}_{cl} subject to $H' \subseteq H$, we have $U \subseteq H'$. Furthermore, as H'' is the unique maximal element in \mathcal{P} subject to $H'' \subseteq H'$, we have $U \subseteq H''$, contradicting the assumption $U \subsetneq H''$.

By definition, we have $H \in \mathcal{P}_{\text{cl}}$ and $u_{\mathcal{P}}(H) = H''$. \square

We show that a \mathcal{P} -scheme can always be extended to a \mathcal{P}_{cl} -scheme where antisymmetry and strong antisymmetry are preserved.

Lemma 6.10. *Let \mathcal{P} be a subgroup system over a group G and let $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ be a \mathcal{P} -scheme. There exists a unique \mathcal{P}_{cl} -scheme $\mathcal{C}' = \{C'_H : H \in \mathcal{P}_{\text{cl}}\}$ extending \mathcal{C} (i.e., $C'_H = C_H$ for $H \in \mathcal{P}$), given by*

$$C'_H = \{\pi_{u(H),H}(B) : B \in C_{u(H)}\}.$$

Moreover, if \mathcal{C} is antisymmetric (resp. strongly antisymmetric), so is \mathcal{C}' . And \mathcal{C}' is not discrete on $H \in \mathcal{P}_{\text{cl}}$ if \mathcal{C} is antisymmetric and not discrete on $u(H)$.

Proof. We have $u(H) \in \mathcal{P} \subseteq \mathcal{P}_{\text{cl}}$ for $H \in \mathcal{P}_{\text{cl}}$. It follows from Lemma 2.3 that \mathcal{C}' as defined above is the only possible one extending \mathcal{C} .

Then we check that \mathcal{C}' is indeed well defined, i.e., for $H \in \mathcal{P}_{\text{cl}}$, the set $C'_H = \{\pi_{u(H),H}(B) : B \in C_{u(H)}\}$ is indeed a partition of $H \setminus G$. For two blocks $B_1, B_2 \in C_{u(H)}$, we prove that $\pi_{u(H),H}(B_1)$ and $\pi_{u(H),H}(B_2)$ are either identical or disjoint. Suppose there exist $u(H)g_1 \in B_1$ and $u(H)g_2 \in B_2$ satisfying $\pi_{u(H),H}(u(H)g_1) = \pi_{u(H),H}(u(H)g_2)$, i.e., $Hg_1 = Hg_2$. Then $g_2g_1^{-1} \in H \subseteq N_G(u(H))$. Note that $c_{u(H),g_2g_1^{-1}}(u(H)g_1) = u(H)g_2$. So by invariance of \mathcal{C} , we have $c_{u(H),g_2g_1^{-1}}(B_1) = B_2$. Then by Lemma 2.2, we have

$$\pi_{u(H),H}(B_2) = \pi_{u(H),H} \circ c_{u(H),g_2g_1^{-1}}(B_1) = c_{H,g_2g_1^{-1}} \circ \pi_{u(H),H}(B_1) = \pi_{u(H),H}(B_1)$$

as desired. So \mathcal{C}' is well defined. Moreover, we have $u(H) = H$ for $H \in \mathcal{P}$. It follows that \mathcal{C}' does extend \mathcal{C} .

Next we show that \mathcal{C}' is a \mathcal{P}_{cl} -scheme. For $H, H' \in \mathcal{P}_{\text{cl}}$ with $H \subseteq H'$, we have $u(H) \subseteq H'$ and hence $u(H) \subseteq u(H')$ by the unique maximality of $u(H')$. By transitivity of projections (see Lemma 2.2), the following diagram commutes:

$$\begin{array}{ccc} u(H) \backslash G & \xrightarrow{\pi_{u(H), u(H')}} & u(H') \backslash G \\ \pi_{u(H), H} \downarrow & & \downarrow \pi_{u(H'), H'} \\ H \backslash G & \xrightarrow{\pi_{H, H'}} & H' \backslash G \end{array}$$

To show compatibility, consider $y, y' \in H \backslash G$ lying in the same block $B \in C'_H$. Choose $\tilde{B} \in C_{u(H)}$ satisfying $\pi_{u(H), H}(\tilde{B}) = B$ and choose $x, x' \in \tilde{B}$ satisfying $\pi_{u(H), H}(x) = y$, $\pi_{u(H), H}(x') = y'$. By compatibility of \mathcal{C} , the elements $\pi_{u(H), u(H')}(x)$ and $\pi_{u(H), u(H')}(x')$ lie in the same block of $C_{u(H')}$. Then $\pi_{u(H'), H'} \circ \pi_{u(H), u(H')}$ maps x and x' into the same block of $C'_{H'}$ by the definition of \mathcal{C}' . By commutativity of the diagram above and the facts $\pi_{u(H), H}(x) = y$, $\pi_{u(H), H}(x') = y'$, we see that $\pi_{H, H'}(y)$ and $\pi_{H, H'}(y')$ lie in the same block of $C'_{H'}$. So \mathcal{C}' is compatible.

For regularity, let B be a block of C'_H . Then $\pi_{H, H'}(B)$ is contained in a unique block B' of $C'_{H'}$ by compatibility of \mathcal{C}' . Lift B to a block $\tilde{B} \in C_{u(H)}$ along $\pi_{u(H), H}$, and let $\tilde{B}' = \pi_{u(H), u(H')}(\tilde{B}) \in C_{u(H')}$. By regularity of \mathcal{C} , the map $\pi_{u(H), u(H')}|_{\tilde{B}} : \tilde{B} \rightarrow \tilde{B}'$ has constant degree, i.e., the number of preimages $|(\pi_{u(H), u(H')}|_{\tilde{B}})^{-1}(y)|$ is independent of the choices of $y \in \tilde{B}'$. We show that $\pi_{u(H), H}|_{\tilde{B}}$ (and similarly $\pi_{u(H'), H'}|_{\tilde{B}'}$) also has constant degree. Consider $y, y' \in B$. As $\pi_{u(H), H}(\tilde{B}) = B$, there exists $x, x' \in \tilde{B}$ satisfying $\pi_{u(H), H}(x) = y$ and $\pi_{u(H), H}(x') = y'$. Note that all the elements in $(\pi_{u(H), H}|_{\tilde{B}})^{-1}(y)$ (resp. $(\pi_{u(H), H}|_{\tilde{B}})^{-1}(y')$) are of the form $c_{u(H), g}(x)$ (resp. $c_{u(H), g}(x')$) for some $g \in H$ since $H \subseteq N_G(u(H))$. And we have $c_{u(H), g}(x) \in \tilde{B}$ iff $c_{u(H), g}(x') \in \tilde{B}$ for $g \in H$ by invariance of \mathcal{C} . It follows that $|(\pi_{u(H), H}|_{\tilde{B}})^{-1}(y)| = |(\pi_{u(H), H}|_{\tilde{B}})^{-1}(y')|$. So $\pi_{u(H), H}|_{\tilde{B}}$ (and similarly $\pi_{u(H'), H'}|_{\tilde{B}'}$) has constant degree. Then $\pi_{H, H'}|_B$ also has constant degree by the commutativity of the diagram above. So \mathcal{C}' is regular.

For invariance, note that for $H, H' \in \mathcal{P}_{\text{cl}}$ with $H' = gHg^{-1}$, we have $u(H') =$

$gu(H)g^{-1}$. And the following diagram commutes by Lemma 2.2:

$$\begin{array}{ccc} u(H)\backslash G & \xrightarrow{c_{u(H),g}} & u(H')\backslash G \\ \pi_{u(H),H} \downarrow & & \downarrow \pi_{u(H'),H'} \\ H\backslash G & \xrightarrow{c_{H,g}} & H'\backslash G \end{array}$$

For a block B of C'_H , lift it to a block \tilde{B} of $C_{u(H)}$. Then $c_{H,g}(B) = \pi_{u(H'),H'} \circ c_{u(H),g}(\tilde{B})$ by the commutativity of the diagram above. Note that $c_{u(H),g}(\tilde{B})$ is a block of $C_{u(H')}$ by invariance of \mathcal{C} . So $c_{H,g}(B)$ is a block of C'_H by definition. Therefore \mathcal{C}' is invariant.

Now assume \mathcal{C}' is not strongly antisymmetric and we prove that \mathcal{C} is not either. By definition, there exists a nontrivial permutation $\tau = \sigma_k \circ \cdots \circ \sigma_1$ of a block $B \in C'_H$ for some $H \in \mathcal{P}_{\text{cl}}$ such that each $\sigma_i : B_{i-1} \rightarrow B_i$ is a map of the form $c_{H_{i-1},g}|_{B_{i-1}}$, $\pi_{H_{i-1},H_i}|_{B_{i-1}}$, or $(\pi_{H_i,H_{i-1}}|_{B_i})^{-1}$, and $B_i \in C'_{H_i}$, $H_i \in \mathcal{P}_{\text{cl}}$, $B = B_0 = B_k$, $H = H_0 = H_k$ (see Definition 2.7). By the two diagrams above, we can lift each B_i to $\tilde{B}_i \in C_{u(H_i)}$ for $0 \leq i \leq k$ and lift each σ_i to a map $\tilde{\sigma}_i : \tilde{B}_{i-1} \rightarrow \tilde{B}_i$ of the form $c_{u(H_{i-1}),g}|_{\tilde{B}_{i-1}}$, $\pi_{u(H_{i-1}),u(H_i)}|_{\tilde{B}_{i-1}}$, or $(\pi_{u(H_i),u(H_{i-1})}|_{\tilde{B}_i})^{-1}$ respectively, i.e., $\pi_{u(H_i),H_i}(\tilde{B}_i) = B_i$ and $\sigma_i \circ \pi_{u(H_{i-1}),H_{i-1}}|_{\tilde{B}_{i-1}} = \pi_{u(H_i),H_i}|_{\tilde{B}_i} \circ \tilde{\sigma}_i$.⁴ Then $\tilde{\tau} := \tilde{\sigma}_k \circ \cdots \circ \tilde{\sigma}_1$ is a map from \tilde{B}_0 to \tilde{B}_k lifting τ . Note that $\pi_{u(H),H}(\tilde{B}_0) = \pi_{u(H),H}(\tilde{B}_k) = B$. So $c_{u(H),g}(\tilde{B}_k) = \tilde{B}_0$ for some $g \in H$. By composing $\tilde{\tau}$ with $c_{u(H),g}$ (and noting that $c_{H,g}$ is the identity map), we may assume $\tilde{B}_k = \tilde{B}_0$. So $\tilde{\tau}$ is a permutation of \tilde{B}_0 . Moreover $\tilde{\tau}$ is nontrivial since it lifts τ . So \mathcal{C} is not strongly antisymmetric. The proof for antisymmetry is the same except that we only consider maps τ that are conjugations.

Finally, to prove the last claim, assume \mathcal{C} is antisymmetric and \mathcal{C}' is discrete on $H \in \mathcal{P}_{\text{cl}}$. We prove that \mathcal{C} is discrete on $u(H)$. Consider distinct elements $x, x' \in u(H)\backslash G$ and let $y = \pi_{u(H),H}(x)$, $y' = \pi_{u(H),H}(x')$. If $y \neq y'$, they are in different blocks of C'_H and hence x, x' are in different blocks of $C_{u(H)}$ by the definition of C'_H . So assume $y = y'$. Then $x = u(H)g$, $x' = u(H)g'$ for some $g, g' \in G$ satisfying $Hg = Hg'$, i.e., $g'g^{-1} \in H \subseteq N_G(u(H))$. As $x' = c_{u(H),g'g^{-1}}(x)$, the elements x and x' are in different blocks of $C_{u(H)}$ by antisymmetry of \mathcal{C} . So \mathcal{C} is discrete on $u(H)$, as desired. \square

⁴For the case that $\sigma_i = (\pi_{H_i,H_{i-1}}|_{B_i})^{-1}$, we lift $\pi_{H_i,H_{i-1}}|_{B_i}$ to $\pi_{u(H_i),u(H_{i-1})}|_{\tilde{B}_i}$. As \mathcal{C} is antisymmetric, both $\pi_{u(H_{i-1}),H_{i-1}}|_{\tilde{B}_{i-1}}$ and $\pi_{u(H_i),H_i}|_{\tilde{B}_i}$ are bijective. So $\pi_{u(H_i),u(H_{i-1})}|_{\tilde{B}_i}$ is also bijective and its inverse is well defined.

Recall that for a subgroup system \mathcal{P} over a finite group G , we let $\mathcal{P}_+ = \{H : H' \subseteq H \subseteq N_G(H'), H' \in \mathcal{P}\}$ which is also a subgroup system over G (see Section 4.4). Clearly $\mathcal{P}_{\text{cl}} \subseteq \mathcal{P}_+$. We show that equality holds if \mathcal{P} is *join-closed*.

Lemma 6.11. *Let \mathcal{P} be a subgroup system that is join-closed, i.e., $\langle H, H' \rangle \in \mathcal{P}$ for all $H, H' \in \mathcal{P}$. Then $\mathcal{P}_{\text{cl}} = \mathcal{P}_+$.*

Proof. Consider $H \in \mathcal{P}_+$. We prove $H \in \mathcal{P}_{\text{cl}}$ by verifying the conditions in Definition 6.5.

Choose a maximal element $H' \in \mathcal{P}$ subject to $H' \subseteq H$. Such an element exists by the definition of \mathcal{P}_+ . We first show that H' is unique. Assume to the contrary that there exists another maximal element $H'' \subseteq H$ in \mathcal{P} different from H' . Then $\langle H', H'' \rangle \supsetneq H'$ is also a subgroup of H and lies in \mathcal{P} by join-closedness, contradicting maximality of H' . So H' is unique.

Next we prove H' is normal in H . Assume to the contrary that there exists $g \in H$ such that $gH'g^{-1} \neq H'$. As $gH'g^{-1} \subseteq gHg^{-1} = H$ and $gH'g^{-1} \in \mathcal{P}$, the join $\langle H', gH'g^{-1} \rangle \supsetneq H'$ is also a subgroup of H and lies in \mathcal{P} by join-closedness, again contradicting maximality of H' . \square

As an application, we consider a system of stabilizers with respect to the natural action of a symmetric group or an alternating group.

Lemma 6.12. *Let S be a finite G -set where G is $\text{Sym}(S)$ or $\text{Alt}(S)$ acting naturally on S . Let $\mathcal{P} = \mathcal{P}_m$ be the corresponding system of stabilizers of depth m , where $m < |S|/2$. Then $\mathcal{P}' := \mathcal{P} \cup \{G\}$ is join-closed.*

Proof. Note $\mathcal{P}' = \{G_T : 0 \leq T \leq m\}$. Let T and T' be subsets of S of cardinality at most m . We show that $\langle G_T, G_{T'} \rangle \in \mathcal{P}'$. Obviously we have $\langle G_T, G_{T'} \rangle \subseteq G_{T \cap T'}$.

First assume $G = \text{Sym}(S)$. We have $G_T \cong \text{Sym}(S - T)$, $G_{T'} \cong \text{Sym}(S - T')$ and $G_{T \cap T'} \cong \text{Sym}(S - (T \cap T'))$ by restricting to the subsets $S - T$, $S - T'$ and $S - (T \cap T')$ respectively. The group $\text{Sym}(S - (T \cap T'))$ is generated by transpositions $(x y)$ with $x, y \in S - (T \cap T')$. We claim that every such $(x y)$ is contained in $\langle G_T, G_{T'} \rangle$. This is obvious if x and y are both in $S - T$ or $S - T'$. So we assume $x \in T - T'$ and $y \in T' - T$. As $m < |S|/2$, the set $S - (T \cup T')$ is not empty. Pick $z \in S - (T \cup T')$. Then $(x y) = (y z)(x z)(y z)^{-1} \in \langle G_T, G_{T'} \rangle$ since $y, z \in S - T$ and $x, z \in S - T'$. So $\langle G_T, G_{T'} \rangle = G_{T \cap T'} \in \mathcal{P}'$.

Next assume $G = \text{Alt}(S)$. If $|S| \leq 4$, one can directly verify that $\langle G_T, G_{T'} \rangle$ equals G , G_T or $G_{T'}$. So assume $|S| \geq 5$. Note that $G_{T \cap T'} \cong \text{Alt}(S - (T \cap T'))$ is generated by 3-cycles $(x y z)$ with $x, y, z \in S - (T \cap T')$. We claim that every such $(x y z)$ is contained in $\langle G_T, G_{T'} \rangle$. This is obvious if x, y, z are all in $S - T$ or $S - T'$. So we assume $x, y \in T - T'$ and $z \in T' - T$ (the other cases are symmetric). Pick $w \in S - (T \cup T')$ and let $(w z u)$ be a 3-cycle for some $u \in S - T - \{z, w\}$. Then $(x y z) = (w z u)(x y w)(w z u)^{-1} \in \langle G_T, G_{T'} \rangle$ since $w, z, u \in S - T$ and $x, y, w \in S - T'$. So again $\langle G_T, G_{T'} \rangle = G_{T \cap T'} \in \mathcal{P}'$. \square

Corollary 6.5. *Let S be a finite G -set where G is $\text{Sym}(S)$ or $\text{Alt}(S)$ acting naturally on S . Let $\mathcal{P} = \mathcal{P}_m$ be the corresponding system of stabilizers of depth m , where $m < |S|/2$. Then $\mathcal{P}_{\text{cl}} = \mathcal{P}_+$.*

Proof. Let $\mathcal{P}' = \mathcal{P} \cup \{G\}$. Then by Lemma 6.12, we have $\mathcal{P}_+ \subseteq \mathcal{P}'_+ = \mathcal{P}'_{\text{cl}} = \mathcal{P}_{\text{cl}} \cup \{G\}$. If $G \in \mathcal{P}$, we have $\mathcal{P}_{\text{cl}} \cup \{G\} = \mathcal{P}_{\text{cl}}$ and hence $\mathcal{P}_+ \subseteq \mathcal{P}_{\text{cl}}$. On the other hand, if $G \notin \mathcal{P}$, none of the groups in \mathcal{P} is normal in G , and hence $G \notin \mathcal{P}_+$. So we still have $\mathcal{P}_+ \subseteq \mathcal{P}_{\text{cl}}$. \square

Remark. The condition $m < |S|/2$ is necessary: suppose $|S| \geq 6$ is even and let $m = |S|/2$. Partition S into S_1 and S_2 of the same cardinality m . When $G = \text{Sym}(S)$ (resp. $G = \text{Alt}(S)$), the subgroup $\langle G_{S_1}, G_{S_2} \rangle$ is the product of two copies of the symmetric group (resp. alternating group) of degree m . It is a proper subgroup of G but stabilizes no element of S . Therefore $\langle G_{S_1}, G_{S_2} \rangle \notin \mathcal{P}_m \cup \{G\}$. Indeed, we have $\langle G_{S_1}, G_{S_2} \rangle \in (\mathcal{P}_m)_+ - (\mathcal{P}_m)_{\text{cl}}$ since $G_{S_1} \subseteq \langle G_{S_1}, G_{S_2} \rangle \subseteq N_G(G_{S_1})$ whereas both G_{S_1} and G_{S_2} are maximal among subgroups of $\langle G_{S_1}, G_{S_2} \rangle$ in \mathcal{P}_m .

Lemma 4.16 now follows from Lemma 6.10 and Corollary 6.5.

We also consider the case $G = \text{GL}(V)$ with the natural action on a vector space V .

Lemma 6.13. *Let V be a finite dimensional vector space over a finite field F . Let $\mathcal{P} = \mathcal{P}_m$ be the system of stabilizers of depth m with respect to the natural action of $G := \text{GL}(V)$ on $S := V - \{0\}$, where $m < \dim_F V$. Then $\mathcal{P}_{\text{cl}} = \mathcal{P}_+$.*

Proof. Consider $H \in \mathcal{P}_+$ and we prove that $H \in \mathcal{P}_{\text{cl}}$. Choose $H' \in \mathcal{P}$ such that $H' \subseteq H \subseteq N_G(H')$. It suffices to show that H' is the unique maximal element in \mathcal{P} subject to $H' \subseteq H$. Assume to the contrary that there exists another maximal element $H'' \subseteq H$ in \mathcal{P} . As $m < \dim_F V$, we have $H' = G_{V'}$ and $H'' = G_{V''}$ for

some proper linear subspaces V', V'' of V . As $H'' \not\subseteq H'$, we have $V' \not\subseteq V''$. Also note that $V - (V' \cup V'') \neq \emptyset$ since

$$|V' \cup V''| = |V'| + |V''| - |V' \cap V''| < 2|V|/|F| \leq |V|.$$

Pick $v \in V' - V''$ and $v' \in V - (V' \cup V'')$. Choose $g \in H'' = G_{V''}$ sending v to v' which is possible since $v, v' \notin V''$. As $g \in H'' \subseteq H \subseteq N_G(H') = N_G(G_{V'})$, we have ${}^gV' = V'$. But ${}^gv = v' \notin V'$, and we get a contradiction. \square

Lemma 4.16 now follows from Lemma 6.10 and Lemma 6.13.

6.5 Restricting to a subset

Suppose $\Pi = \{P_1, \dots, P_m\}$ is an m -scheme on a finite set S and T is a subset of S . Then we can restrict Π to T and obtain an m -collection on T , denote by $\Pi|_T$.⁵ In this section, we investigate this operation and use it to prove Lemma 6.8 in Section 6.2. We also discuss its generalization for \mathcal{P} -schemes, where \mathcal{P} is a system of stabilizers.

Definition 6.6. Let $\Pi = \{P_1, \dots, P_m\}$ be an m -collection on a finite set S , where $m \in \mathbb{N}^+$. For a subset T of S , define the m -collection $\Pi|_T := \{P'_1, \dots, P'_m\}$ on T , where $P'_k := P_k|_{T^{(k)}}$ is the restriction of P_k to $T^{(k)} \subseteq S^{(k)}$ for $k \in [m]$.

Lemma 6.14. Suppose $\Pi = \{P_1, \dots, P_m\}$ is an m -scheme on S and $T \subseteq S$ is a disjoint union of blocks in P_1 . Then $\Pi|_T$ is also an m -scheme. Moreover, if Π is antisymmetric (resp. strongly antisymmetric), so is $\Pi|_T$. And if Π does not have a matching, neither does $\Pi|_T$.

Proof. By compatibility of Π , for $k \in [m]$ and $B \in P_k$, either $B \subseteq T^{(k)}$ or $B \cap T^{(k)} = \emptyset$, and hence $T^{(k)}$ is a disjoint union of blocks of P_k . Then the various properties of $\Pi|_T$ (compatibility, regularity, etc.) follow from those of Π in a straightforward manner. \square

In particular, suppose $\Pi = \{P_1, \dots, P_m\}$ is a strongly antisymmetric m -scheme on S that is not discrete. Let T be a block of P_1 such that $|T| > 1$. Then $\Pi|_T$ is a strongly antisymmetric homogeneous m -scheme on T . Lemma 6.8 now follows.

⁵It should not be confused with the notation $\Pi|_{x_1, \dots, x_k}$ in Definition 6.3, which is an $(m - k)$ -scheme on $S - \{x_1, \dots, x_k\}$.

Next we discuss the analogue of Lemma 6.14 for \mathcal{P} -schemes. Let G be a finite group acting on a finite set S . Let $\mathcal{P} = \mathcal{P}_m$ be the corresponding system of stabilizers of depth m over G for some $m \in \mathbb{N}^+$. By Lemma 2.1, for $x \in S$, we have an equivalence of group actions

$$\lambda_x : Gx \rightarrow G_x \backslash G$$

between the action of G on the G -orbit Gx and that on $G_x \backslash G$ by inverse right translation. It sends $g \cdot x$ to $G_x g^{-1}$ for $g \in G$.

Definition 6.7. *Let m , G and \mathcal{P} be as above. Let $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ be a \mathcal{P} -scheme. Let T be a subset of S such that for $z \in T$, the set $\lambda_z(Gz \cap T)$ is a disjoint union of blocks in C_{G_z} . Moreover, define G' to be the setwise stabilizer $G_{\{T\}}$ and suppose it satisfies the following conditions:*

1. *For $U, U' \subseteq T$ satisfying $1 \leq |U|, |U'| \leq m$ and $G'_U \subseteq G'_{U'}$, we have $G_U \subseteq G_{U'}$.*
2. *For $k \in [m]$ and $x \in T^{(k)}$, we have $G'x = Gx \cap T^{(k)}$.*

Let \mathcal{P}' be the system of stabilizers of depth m over G' with respect to the action of G' on T (restricted from the action of G on S). We define a \mathcal{P}' -collection $\mathcal{C}' = \{C'_H : H \in \mathcal{P}'\}$ as follows:

For $H \in \mathcal{P}'$, choose a nonempty subset $U \subseteq T$ of cardinality at most m such that $H = G'_U$. Identify $G'_U \backslash G'$ with a subset of $G_U \backslash G$ via the injective map $i_U : G'_U \backslash G' \hookrightarrow G_U \backslash G$ sending $G'_U g$ to $G_U g$ for $g \in G'$.⁶ Then define C'_H to be the restriction of C_{G_U} to $G'_U \backslash G'$.

The assumption that $\lambda_z(Gz \cap T)$ is a disjoint union of blocks in C_{G_z} for all $z \in T$ is the analogue of the assumption in Lemma 6.14 that T is a disjoint union of blocks in P_1 . If G acts transitively on S , we have $Gz = S$, in which case this assumption is equivalent to that $\lambda_z(T)$ is a disjoint union of blocks in C_{G_z} for some $z \in T$. Note that we also need two additional conditions on G' . They are satisfied in the following important cases.

⁶This is indeed a well defined injective map by Lemma 2.1. Let G' act on $G_U \backslash G$ by inverse right translation and let O be the G' -orbit of $G_U e$. The stabilizers of $G_U e$ is G'_U . So we have a bijection $O \rightarrow G'_U \backslash G'$ whose inverse (composed with $O \hookrightarrow G_U \backslash G$) is i_U .

Example 6.1. Suppose G is the full symmetric group $\text{Sym}(S)$ acting naturally on S . The image of the permutation representation $G' \rightarrow \text{Sym}(T)$ is $\text{Sym}(T)$. In this case the two conditions in Definition 6.7 are satisfied for any subset T of S whose cardinality greater than $m + 1$.⁷ Indeed, if we view the \mathcal{P} -scheme \mathcal{C} as an m -scheme by Theorem 2.1, the construction of \mathcal{C}' from \mathcal{C} is precisely the restriction of an m -scheme to the subset T (see Definition 6.6).

Example 6.2. Suppose $S = V - \{0\}$ where V is a finite dimensional vector space over a finite field F . Let G be the general linear group $\text{GL}(V)$ acting naturally on S . Let $T = V' - \{0\} \subseteq S$ where V' is a linear subspace of V . The image of the permutation representation $G' \rightarrow \text{Sym}(T)$ is isomorphic to $\text{GL}(V')$. It is easy to verify that in this case the two conditions in Definition 6.7 are also satisfied.

We prove the following generalization of Lemma 6.14.

Lemma 6.15. *The \mathcal{P}' -collection \mathcal{C}' is a well defined \mathcal{P}' -scheme. Moreover, if \mathcal{C} is antisymmetric (resp. strongly antisymmetric), so is \mathcal{C}' .*

Proof. In Definition 6.7 we define each C'_H by picking $U \subseteq T$ of cardinality at most m satisfying $H = G'_U$. Here the group G_U and the map i_U do not depend on the choice of U by the first condition in Definition 6.7. So \mathcal{C}' is well defined.

For $H, H' \in \mathcal{P}'$ satisfying $H \subseteq H'$, we pick nonempty subsets $U, U' \subseteq T$ of cardinality at most m such that $H = G'_U$ and $H' = G'_{U'}$. Then $G_U \subseteq G_{U'}$ by the first condition in Definition 6.7. And the following diagram commutes.

$$\begin{array}{ccc} G'_U \backslash G' & \xrightarrow{\pi_{G'_U, G'_{U'}}} & G'_{U'} \backslash G' \\ i_U \downarrow & & \downarrow i_{U'} \\ G_U \backslash G & \xrightarrow{\pi_{G_U, G_{U'}}} & G_{U'} \backslash G \end{array}$$

For $H, H' \in \mathcal{P}'$ and $g \in G'$ satisfying $H' = gHg^{-1}$, we pick a nonempty subset $U \subseteq T$ of cardinality at most m such that $H = G'_U$, and let $U' = {}^gU \subseteq T$. Then $H' = G'_{U'} = gG'_U g^{-1}$ and $G_{U'} = gG_U g^{-1}$. And the following diagram commutes.

$$\begin{array}{ccc} G'_U \backslash G' & \xrightarrow{c_{G'_U, g}} & G'_{U'} \backslash G' \\ i_U \downarrow & & \downarrow i_{U'} \\ G_U \backslash G & \xrightarrow{c_{G_U, g}} & G_{U'} \backslash G \end{array}$$

⁷The first condition does not hold for $|T| \leq m + 1$: if $U, U' \subseteq T$ are different subsets of cardinality $|T| - 1$, we have $G'_U = G'_{U'} = G'_T$, but $G_U \neq G_{U'}$ unless $T = S$.

Let U be a nonempty subset of T of cardinality at most m . We claim i_U maps each block of $C'_{G'_U}$ to a block of C_{G_U} . The rest of the proof focuses on this claim. Combining it with the two diagrams above, we can derive the various properties of \mathcal{C}' (compatibility, regularity, invariance, antisymmetry and strong antisymmetry) from the corresponding properties of \mathcal{C} in a straightforward manner.

Let B be a block of $C'_{G'_U}$ and B' be the block of C_{G_U} containing $i_U(B)$. Assume to the contrary that $i_U(B) \neq B'$. Choose $G_U g^{-1}, G_U g'^{-1} \in G_U \setminus G$, represented by $g^{-1}, g'^{-1} \in G$ respectively, such that $G_U g^{-1} \in i_U(B)$ and $G_U g'^{-1} \in B' - i_U(B)$. We may assume $g \in G'$ and hence ${}^g z \in {}^g T = T$ for all $z \in T$. Also note $B = i_U^{-1}(B')$ by construction. So from $G_U g'^{-1} \in B' - i_U(B)$ we know $G_U g'^{-1} \notin i_U(G'_U \setminus G')$.

Assume there exists $z \in U$ such that ${}^{g'} z \notin T$. As $G_U g^{-1}$ and $G_U g'^{-1}$ are in the same block B' of C_{G_U} , by compatibility of \mathcal{C} we know $\pi_{G_U, G_z}(G_U g^{-1}) = G_z g^{-1}$ and $\pi_{G_U, G_z}(G_U g'^{-1}) = G_z g'^{-1}$ are in the same block of C_{G_z} . On the other hand, we have $G_z g^{-1} = \lambda_z({}^g z) \in \lambda_z(Gz \cap T)$ and $G_z g'^{-1} = \lambda_z({}^{g'} z) \notin \lambda_z(Gz \cap T)$ since ${}^g z \in Gz \cap T$, ${}^{g'} z \notin T$ and $\lambda_z : Gz \rightarrow G_z \setminus G$ is a bijection. But this contradicts the assumption that $\lambda_z(Gz \cap T)$ is a disjoint union of blocks of C_{G_z} .

Now assume ${}^{g'} z \in T$ for all $z \in U$. Suppose $U = \{x_1, \dots, x_k\}$, where x_i are distinct and ordered in an arbitrary way. Let $x = (x_1, \dots, x_k) \in T^{(k)}$. Then ${}^{g'} x$ is in $Gx \cap T^{(k)}$ and hence in $G'x$ by the second condition in Definition 6.7. So ${}^{g'} x = {}^{g''} x$ for some $g'' \in G'$. Then $g'^{-1} g'' \in G_x = G_U$. So $G_U g'^{-1} = G_U g''^{-1} = i_U(G'_U g''^{-1})$, contradicting the fact $G_U g'^{-1} \notin i_U(G'_U \setminus G')$ above. This proves the claim that i_U maps each block of $C'_{G'_U}$ to a block of C_{G_U} . \square

6.6 Primitivity of homogeneous m -schemes

The notion of *primitivity* is important for permutation groups as well as association schemes. In this section, we extend it to homogeneous m -schemes. As an application, we show that every antisymmetric homogeneous orbit m -scheme on a finite set S has a matching if $|S| > 1$ and $m \geq 3$.

Definition 6.8 (primitivity). *Let $\Pi = \{P_1, \dots, P_m\}$ be a homogeneous m -scheme on a finite set S . For $B \in P_2$, denote by G_B the simple graph⁸ on the vertex set S such that there exists an edge between two distinct vertices u, v iff (u, v) or (v, u) is in B . We say Π is primitive if G_B is connected for all $B \in P_2$. Otherwise Π is imprimitive.*

⁸A simple graph is an undirected graph without loops or multiple edges.

The reader familiar with *primitivity of association schemes* (see, e.g., [CGS78]) may recognize that when $m \geq 3$, Definition 6.8 simply defines $\Pi = \{P_1, \dots, P_m\}$ to be primitive iff $P(\Pi')$ is primitive, where Π' denotes the homogeneous 3-scheme $\{P_1, P_2, P_3\}$ and $P(\Pi')$ is the corresponding association scheme (see Definition 2.16).

Remark. Our definition of primitivity coincides with the notion of *primitivity at level 2* introduced in the full version of [IKS09]. The same paper also generalizes the notion of primitivity to higher levels. We will not discuss their generalization in this thesis, but refer the interested reader to [IKS09] for further details.

Restricting to a connected component. We note that restricting a homogeneous m -scheme to a connected component yields another homogeneous m -scheme:

Lemma 6.16. *Let $\Pi = \{P_1, \dots, P_m\}$ be a homogeneous m -scheme on a finite set S where $m \geq 3$. For each $B \in P_2$ and a connected component $T \subseteq S$ of G_B , the m -collection $\Pi|_T$ (see Definition 6.6) is a homogeneous m -scheme on T . Moreover, if Π is antisymmetric (resp. strongly antisymmetric), then so is $\Pi|_T$. And if Π has no matching, then neither does $\Pi|_T$.*

Proof. Let $T \subseteq S$ be as in the lemma. It is well known that there exist blocks $B_1, \dots, B_k \in P_2$ such that the union of these blocks and $1_S = \{(x, x) : x \in S\}$ yields an equivalence relation \sim on S , and T is one of its equivalence classes (see, e.g., [CGS78]).

For $k \in [m]$, define the equivalence relation \sim_k on $S^{(k)}$ such that $(x_1, \dots, x_k) \sim_k (y_1, \dots, y_k)$ iff $x_i \sim y_i$ for all $i \in [k]$. These equivalence relations are respected by the maps π_i^k and c_g^k . The various properties of $\Pi|_T$ then follow from the corresponding properties of Π in a straightforward manner. \square

Primitivity of homogeneous orbit m -schemes. The next lemma states that primitivity of homogeneous orbit m -schemes is equivalent to primitivity of the associated permutation group.

Lemma 6.17. *A homogeneous orbit m -scheme on a finite set S associated with $K \subseteq \text{Sym}(S)$ is primitive iff K is a primitive permutation group on S .*

Proof. Let $\Pi = \{P_1, \dots, P_m\}$ be a homogeneous orbit m -scheme associated with a group $K \subseteq \text{Sym}(S)$. Then K acts transitively on S . The graphs G_B for $B \in P_2$

are known as the non-diagonal (undirected) *orbital graphs*. The lemma then follows from Definition 6.8 and the well known fact that a transitive permutation group is primitive iff every non-diagonal orbital graph is connected [Hig67]. \square

In general, we can obtain a primitive orbit m -scheme from a possibly imprimitive one by restricting to a minimal set that is a connected component:

Lemma 6.18. *Let $\Pi = \{P_1, \dots, P_m\}$ be a homogeneous orbit m -scheme on S associated with $K \subseteq \text{Sym}(S)$, where $|S| > 1$. Let T be a minimal subset of S such that T is a connected component of G_B for some $B \in P_2$. Let K' be the image of the permutation representation $K|_{\{T\}} \rightarrow \text{Sym}(T)$. Then $\Pi|_T$ is a primitive homogeneous orbit m -scheme on T , and is the orbit m -scheme associated with K' .*

Proof. As already noted, for $B \in P_2$ and any connected component T' of G_B , there exist blocks $B_1, \dots, B_k \in P_2$ such that the union of these blocks and $1_S = \{(x, x) : x \in S\}$ yields an equivalence relation on S where T' is an equivalence class [CGS78]. Primitivity of Π then follows from minimality of T .

Choose $B \in P_2$ such that T is a connected component of G_B . Note that for $g \in K$ and $(u, v) \in S^{(2)}$, the edge (u, v) is in G_B iff $({}^g u, {}^g v)$ is in G_B . So for $g \in K$, the set ${}^g T$ is a connected component of G_B . It follows that T is a set of imprimitivity of K , i.e., ${}^g T \cap T = \emptyset$ or ${}^g T = T$ for all $g \in K$.

Consider $k \in [m]$ and $x, y \in T^{(k)}$ in the same block of $P_k|_{T^{(k)}} \in \Pi|_T$. There exists $g \in K$ sending x to y . As T is a set of imprimitivity of K , we have ${}^g T = T$ and hence $g \in K'$. So $\Pi|_T$ is the orbit m -scheme on T associated with K' . \square

Antisymmetric homogeneous orbit m -schemes for $m \geq 3$. As an application, we prove that for $m \geq 3$, an antisymmetric homogeneous orbit m -scheme Π on a finite set S where $|S| > 1$ always has a matching. In particular, it is not strongly antisymmetric by Lemma 2.10. The same claim for $m \geq 4$ was proved in [IKS09]. Note that strongly antisymmetric homogeneous orbit m -schemes on sets S where $|S| > 1$ do exist for $m = 1$ and $m = 2$ (see Section 2.5).

We need the following result from finite group theory.

Lemma 6.19. *Let G be a primitive solvable permutation group on a finite set S . The set S can be identified with a finite dimensional vector space V over a finite field F such that G acts on it as a subgroup of the general affine group*

$$\text{AGL}(V) = \{\phi_{g,u} : g \in \text{GL}(V), u \in V\},$$

where $\phi_{g,u}$ sends $x \in V$ to ${}^g x + u$. Moreover, the group G contains the translation $\phi_{e,u} : x \mapsto x + u$ for all $u \in V$.

See [Sup76, Section I.4] for its proof. We have

Theorem 6.6. *Let $\Pi = \{P_1, \dots, P_m\}$ be an antisymmetric homogeneous orbit m -scheme on a finite set S associated with a group $K \subseteq \text{Sym}(S)$, where $m \geq 3$ and $|S| > 1$. Then Π has a matching.*

Proof. We may assume $m = 3$. Assume to the contrary that Π has no matching. Let T be a minimal subset of S such that T is a connected component of G_B for some $B \in P_2$. Let K' be the image of the permutation representation $K_{\{T\}} \rightarrow \text{Sym}(T)$. By Lemma 6.16 and Lemma 6.18, the m -scheme $\Pi|_T$ is the orbit m -scheme on T associated with K' which is antisymmetric, homogeneous, primitive and has no matching. By replacing Π with $\Pi|_T$, S with T , and K with K' , we may assume Π is primitive. Then K is a primitive permutation group on S by Lemma 6.17. Also note that $|K|$ is odd by Lemma 2.16. It follows by the *Odd Order Theorem* [FT63] that K is solvable. We conclude that K is a primitive solvable permutation group on S of odd order.

By Lemma 6.19, we can identify S with a finite dimensional vector space V over a finite field F , and K with a subgroup of $\text{AGL}(V)$ acting on V that contains all the translations $\phi_{e,u}, u \in V$. Moreover, we have $\text{char}(F) \neq 2$ since $|K|$ is odd.

Choose $v \in V - \{0\}$. Let $x = (0, v, 2v) \in S^{(3)}$, $y = \pi_3^3(x) = (0, v) \in S^{(2)}$ and $z = \pi_1^3(x) = (v, 2v) \in S^{(2)}$. Let $B = Kx \in P_3$, $B' = \pi_3^3(B) = Ky \in P_2$ and $B'' = \pi_1^3(B) = Kz \in P_2$. We claim that B together with the maps $\pi_3^3|_B : B \rightarrow B'$, $\pi_1^3|_B : B \rightarrow B''$ is a matching of Π , which contradicts the assumption. To see this, note that the translation $\phi_{e,v} : x \mapsto x + v$ is in K and sends y to z . So $B' = B''$. We also need to prove $|B| = |B'|$. By the orbit-stabilizer stabilizer theorem, it suffices to show $K_x = K_y$, which holds since $2v$ lies on the affine line spanned by 0 and v , and K acts affine linearly on V . The claim follows. \square

Remark. The first half of our proof basically follows [IKS09] which reduces to the case that K is primitive solvable. In [IKS09], the proof is completed by a result of Seress [Ser96] that bounds the minimal base size of primitive solvable permutation groups of odd order. This result allows them to prove the theorem for $m \geq 4$. We substitute it with the more elementary fact in Lemma 6.19, and use the above argument to prove the theorem for $m \geq 3$.

6.7 Direct products and wreath products

We describe two more techniques of constructing new \mathcal{P} -schemes (resp. m -schemes) from old ones, namely the direct product and the wreath product. They extend the direct product and the wreath product of association schemes (see, e.g., [SS98]). As an application, we show that either the schemes conjecture (Conjecture 6.1) is true, or there exist infinitely many counterexamples.

Direct products. Suppose \mathcal{P} and \mathcal{P}' are subgroup systems over finite groups G and G' respectively. Define

$$\mathcal{P} \times \mathcal{P}' := \{H \times H' : H \in \mathcal{P}, H' \in \mathcal{P}'\}$$

which is a subgroup system over $G \times G'$. For $H \in \mathcal{P}$ and $H' \in \mathcal{P}'$, we have a bijection

$$\phi_{H,H'} : H \backslash G \times H' \backslash G' \rightarrow (H \times H') \backslash (G \times G')$$

sending $(Hg, H'g')$ to $(H \times H')(g, g')$ for $g \in G$ and $g' \in G'$. Then we define the direct product of a \mathcal{P} -collection and a \mathcal{P}' -collection as follows.

Definition 6.9. For a \mathcal{P} -collection $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ and a \mathcal{P}' -collection $\mathcal{C}' = \{C'_{H'} : H' \in \mathcal{P}'\}$, define the $(\mathcal{P} \times \mathcal{P}')$ -collection $\mathcal{C} \times \mathcal{C}' = \{C''_{H \times H'} : H \times H' \in \mathcal{P} \times \mathcal{P}'\}$ by

$$C''_{H \times H'} = \{\phi_{H,H'}(B \times B') : B \in C_H, B' \in C'_{H'}\},$$

called the direct product of \mathcal{C} and \mathcal{C}' .

We have

Lemma 6.20. *The direct product $\mathcal{C} \times \mathcal{C}'$ is a $(\mathcal{P} \times \mathcal{P}')$ -scheme if \mathcal{C} is a \mathcal{P} -scheme and \mathcal{C}' is a \mathcal{P}' -scheme. Moreover, if \mathcal{C} and \mathcal{C}' are antisymmetric (resp. strongly antisymmetric), so is $\mathcal{C} \times \mathcal{C}'$.*

Proof. Write $\pi_{H,H'}$ (resp. $\pi'_{H,H'}$, $\pi''_{H,H'}$) for a projection between coset spaces of subgroups in G (resp. G' , $G \times G'$). Similarly write $c_{H,g}$ (resp. $c'_{H,g}$, $c''_{H,g}$) for a conjugation between coset spaces of subgroups in G (resp. G' , $G \times G'$). For $H = H_1 \times H_2$, $H' = H'_1 \times H'_2 \in \mathcal{P} \times \mathcal{P}'$ satisfying $H \subseteq H'$, we have $H_1 \subseteq H'_1$, $H_2 \subseteq H'_2$ and

$$\pi''_{H_1 \times H_2, H'_1 \times H'_2} \circ \phi_{H_1, H_2}(x, y) = \phi_{H'_1, H'_2}(\pi_{H_1, H'_1}(x), \pi'_{H_2, H'_2}(y))$$

for all $x \in H_1 \backslash G$ and $y \in H_2 \backslash G'$. Similarly, for $H_1 \times H_2 \in \mathcal{P} \times \mathcal{P}'$ and $(g, g') \in G \times G'$, we have

$$c''_{H_1 \times H_2, (g, g')} \circ \phi_{H_1, H_2}(x, y) = \phi_{gH_1g^{-1}, g'H_2g'^{-1}}(c_{H_1, g}(x), c'_{H_2, g'}(y))$$

for all $x \in H_1 \backslash G$ and $y \in H_2 \backslash G'$. The various properties of $\mathcal{C} \times \mathcal{C}'$ (compatibility, regularity, invariance, antisymmetry, and strong antisymmetry) then follow from those of \mathcal{C} and \mathcal{C}' in a straightforward manner. \square

Similarly, we define the direct product of m -schemes:

Definition 6.10. Let $\Pi = \{P_1, \dots, P_m\}$ and $\Pi' = \{P'_1, \dots, P'_m\}$ be m -schemes on finite sets S and S' respectively, where $m \in \mathbb{N}^+$. Define the m -collection $\Pi \times \Pi' = \{P''_1, \dots, P''_m\}$ on $S \times S'$ in the following way: for $k \in [m]$, two elements $z = ((x_1, y_1), \dots, (x_k, y_k))$, $z' = ((x'_1, y'_1), \dots, (x'_k, y'_k)) \in (S \times S')^{(k)}$ are in the same block of P''_k iff the following conditions are satisfied:

1. For $i, j \in [k]$, it holds that $x_i = x_j$ iff $x'_i = x'_j$, and $y_i = y_j$ iff $y'_i = y'_j$.
2. Omit a minimal subset T of coordinates in $[k]$ such that all x_i are distinct, and so are all x'_i . Let $k' = k - |T|$. Suppose the remaining x -coordinates of z and z' are $x_{i_1}, \dots, x_{i_{k'}}$ and $x'_{i_1}, \dots, x'_{i_{k'}}$ respectively. Then $(x_{i_1}, \dots, x_{i_{k'}})$ and $(x'_{i_1}, \dots, x'_{i_{k'}})$ are in the same block of $P_{k'}$.⁹
3. The previous condition holds with x -coordinates replaced by y -coordinates and $P_{k'}$ replaced by $P'_{k'}$.

We have the following analogue of Lemma 6.20 whose proof is left to the reader.

Lemma 6.21. The m -collection $\Pi \times \Pi'$ is an m -scheme on $S \times S'$. Moreover, if Π and Π' are antisymmetric (resp. strongly antisymmetric), so is $\Pi \times \Pi'$. And if Π and Π' have no matching, neither does $\Pi \times \Pi'$.

Remark. The connection between Definition 6.9 and Definition 6.10 is as follows. Given $m \in \mathbb{N}^+$, let \mathcal{P} (resp. \mathcal{P}' , \mathcal{P}'') be the system of stabilizers of depth m over $G = \text{Sym}(S)$ (resp. $G' = \text{Sym}(S')$, $G'' = \text{Sym}(S \times S')$) with respect to the natural action of G on S (resp. G' on S' , G'' on $S \times S'$). Let $\tilde{\mathcal{P}}$ be the system of

⁹The order of these coordinates does not matter by invariance of Π . Under the previous condition, the choice of T does not matter either.

stabilizers of depth m with respect to the *product action* of $G \times G'$ on $S \times S'$.¹⁰ Then $\tilde{\mathcal{P}} \subseteq \mathcal{P} \times \mathcal{P}'$.¹¹ So we obtain a $\tilde{\mathcal{P}}$ -scheme $\tilde{\mathcal{C}}$ from $\mathcal{C} \times \mathcal{C}'$. Using induction of $\tilde{\mathcal{P}}$ -schemes, we obtain a \mathcal{P}'' -scheme \mathcal{C}'' (see Definition 6.4). Using the connection between m -schemes and \mathcal{P} -schemes (see Theorem 2.1), we see that the construction of \mathcal{C}'' from \mathcal{C} and \mathcal{C}' corresponds to a construction of an m -scheme on $S \times S'$ from those on S and S' . This is exactly Definition 6.10.

It is obvious that the direct product also preserves homogeneity and discreteness. By taking iterated direct products, we can construct infinitely many antisymmetric homogeneous m -schemes with no matching if there exists a single one. As an application, we know that either the schemes conjecture (Conjecture 6.1) is true, or there exist infinitely many counterexamples.¹²

Corollary 6.6. *For any $m \in \mathbb{N}^+$, there exist either infinitely many antisymmetric homogeneous m -schemes with no matching or none.*

Wreath products. There exists another operation of \mathcal{P} -schemes and m -schemes called the *wreath product*. While this operation is interesting on its own, we do not need it anywhere else in this thesis, except that it provides an alternative proof of Corollary 6.6. For this reason, we only give the definitions as well as the statements, and leave the proofs to the reader.

We first define the wreath product of groups.

Definition 6.11. *Let G and G' be groups and let Ω be a G' -set. Let G^Ω be the group consisting of all the functions $f : \Omega \rightarrow G$. Its group operation is defined by $(ff')(x) = f(x)f'(x)$. Define the wreath product $G \wr G'$ as the group consisting of all the pairs $(f, g) \in G^\Omega \times G'$, with its group operation defined by*

$$(f, g)(f', g') = (f \cdot {}^g f', gg')$$

for $(f, g), (f', g') \in G \wr G'$, where ${}^g f' : \Omega \rightarrow G$ sends $x \in \Omega$ to $f'(g^{-1}x)$. In other words, the group $G \wr G'$ is the semidirect product $G \rtimes_\varphi G'$ where $\varphi : G' \rightarrow \text{Aut}(G^\Omega)$

¹⁰The product action is defined by ${}^{(g, g')}(x, x') = ({}^g x, {}^{g'} x')$ for $(g, g') \in G \times G'$ and $(x, x') \in S \times S'$.

¹¹To see this, note that for a subset $U \subseteq S \times S'$ whose projections to S and S' are U_1 and U_2 , respectively, we have $(G \times G')_U = G_{U_1} \times G'_{U_2}$.

¹²This claim also holds for the variant of the schemes conjecture (Conjecture 6.2) for the same reason.

sends $g \in G'$ to the automorphism $f \mapsto {}^g f$ of G^Ω . For convenience, we identify G^Ω and G' with subgroups of $G \wr G'$ and write $(f, g) \in G \wr G'$ as fg .

Use the following notations: let G and G' be finite groups and let Ω be a finite G' -set. For a family $\mathcal{H} = \{H_x : x \in \Omega\}$ of subgroups of G indexed by Ω and a subgroup H' of G' satisfying the following condition:

$$H_x = G \text{ for all } x \in \Omega \text{ not fixed by } G', \quad (6.1)$$

write $\mathcal{H} \wr H'$ for the subset

$$\{fg : f(x) \in H_x \text{ for all } x \in \Omega, g \in H'\}$$

of $G \wr G'$, which is a subgroup of $G \wr G'$ by (6.1). Suppose \mathcal{P} and \mathcal{P}' are subgroup systems over finite groups G and G' respectively. Define $\mathcal{P} \wr \mathcal{P}'$ to be the poset of subgroups of $G \wr G'$ consisting of the subgroups $\mathcal{H} \wr H'$ for all $\mathcal{H} = \{H_x \in \mathcal{P} : x \in \Omega\}$ and $H' \in \mathcal{P}'$ satisfying (6.1). Then $\mathcal{P} \wr \mathcal{P}'$ is a subgroup system over $G \wr G'$.

For $\mathcal{H} = \{H_x \in \mathcal{P} : x \in \Omega\}$ and $H' \in \mathcal{P}'$ satisfying (6.1), we have a bijection

$$\phi_{\mathcal{H}, H'} : \left(\prod_{x \in \Omega} H_x \backslash G \right) \times H' \backslash G' \rightarrow (\mathcal{H} \wr H') \backslash (G \wr G')$$

defined as follows: for $f \in \prod_{x \in \Omega} H_x \backslash G$ whose x -factor is $f_x \in H_x \backslash G$, pick $g_x \in G$ such that $f_x = H_x g_x$. Then define $f' : \Omega \rightarrow G$ sending $x \in \Omega$ to g_x . Define $\phi_{\mathcal{H}, H'}$ such that it sends (f, Hg') to $(\mathcal{H} \wr H')f'g'$ for $g' \in G'$. It can be shown that $\phi_{\mathcal{H}, H'}$ is a well defined bijection. Finally, we define the wreath product of a \mathcal{P} -collection and a \mathcal{P}' -collection as follows.

Definition 6.12. For a \mathcal{P} -collection $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ and a \mathcal{P}' -collection $\mathcal{C}' = \{C'_H : H \in \mathcal{P}'\}$, define the $(\mathcal{P} \wr \mathcal{P}')$ -collection $\mathcal{C} \wr \mathcal{C}' = \{C''_{\mathcal{H} \wr H'} : \mathcal{H} \wr H' \in \mathcal{P} \wr \mathcal{P}'\}$ by

$$C''_{\mathcal{H} \wr H'} = \left\{ \phi_{\mathcal{H}, H'} \left(\left(\prod_{x \in \Omega} B_x \right) \times B' \right) : B_x \in C_{H_x} \text{ for } x \in \Omega, B' \in C'_{H'} \right\},$$

where $\mathcal{H} = \{H_x : x \in \Omega\}$. We call $\mathcal{C} \wr \mathcal{C}'$ the wreath product of \mathcal{C} and \mathcal{C}' .

We have

Lemma 6.22. The wreath product $\mathcal{C} \wr \mathcal{C}'$ is a $(\mathcal{P} \wr \mathcal{P}')$ -scheme if \mathcal{C} is a \mathcal{P} -scheme and \mathcal{C}' is a \mathcal{P}' -scheme. Moreover, if \mathcal{C} and \mathcal{C}' are antisymmetric (resp. strongly antisymmetric), then so is $\mathcal{C} \wr \mathcal{C}'$.

Similarly, we define the wreath product of m -schemes:

Definition 6.13. Let $\Pi = \{P_1, \dots, P_m\}$ and $\Pi' = \{P'_1, \dots, P'_m\}$ be m -schemes on finite sets S and S' respectively, where $m \in \mathbb{N}^+$. Define the m -collection $\Pi \wr \Pi' = \{P''_1, \dots, P''_m\}$ on $S \times S'$ in the following way: for $k \in [m]$, two elements $z = ((x_1, y_1), \dots, (x_k, y_k)), z' = ((x'_1, y'_1), \dots, (x'_k, y'_k)) \in (S \times S')^{(k)}$ are in the same block of P''_k iff the following conditions are satisfied:

1. For $i, j \in [k]$, it holds that $y_i = y_j$ iff $y'_i = y'_j$.
2. For $i \in [k]$, let T_i be set of indices $j \in [k]$ satisfying $y_i = y_j$. Suppose $T_i = \{i_1, \dots, i_\ell\}$, ordered in an arbitrary way. Then $(x_{i_1}, \dots, x_{i_\ell})$ and $(x'_{i_1}, \dots, x'_{i_\ell})$ are in the same block of P_ℓ .
3. Omit a minimal subset T of coordinates in $[k]$ such that all y_i are distinct. Let $k' = k - |T|$. Suppose the remaining y -coordinates of z and z' are $y_{i_1}, \dots, y_{i_{k'}}$ and $y'_{i_1}, \dots, y'_{i_{k'}}$ respectively. Then $(y_{i_1}, \dots, y_{i_{k'}})$ and $(y'_{i_1}, \dots, y'_{i_{k'}})$ are in the same block of $P'_{k'}$.

We have the following analogue of Lemma 6.22.

Lemma 6.23. The m -collection $\Pi \wr \Pi'$ is an m -scheme on $S \times S'$. Moreover, if Π and Π' are antisymmetric (resp. strongly antisymmetric), then so is $\Pi \wr \Pi'$. And if Π and Π' have no matching, then neither does $\Pi \wr \Pi'$.

Remark. The connection between Definition 6.12 and Definition 6.13 is as follows. Given $m \in \mathbb{N}^+$, let \mathcal{P} (resp. $\mathcal{P}', \mathcal{P}''$) be the system of stabilizers of depth m over $G = \text{Sym}(S)$ (resp. $G' = \text{Sym}(S'), G'' = \text{Sym}(S \times S')$) with respect to the natural action of G on S (resp. G' on S', G'' on $S \times S'$). Let $\tilde{\mathcal{P}}$ be the system of stabilizers of depth m with respect to the imprimitive wreath product action of $G \wr G'$ on $S \times S'$.¹³ Then $\tilde{\mathcal{P}} \subseteq \mathcal{P} \wr \mathcal{P}'$.¹⁴ So we obtain a $\tilde{\mathcal{P}}$ -scheme $\tilde{\mathcal{C}}$ from $\mathcal{C} \wr \mathcal{C}'$. Using induction of $\tilde{\mathcal{P}}$ -schemes, we obtain a \mathcal{P}'' -scheme \mathcal{C}'' (see Definition 6.4). Using the connection between m -schemes and \mathcal{P} -schemes (see Theorem 2.1), we see that the construction

¹³The imprimitive wreath product action is defined by $(f, g)(x, x') = (f^{(g x')}x, g x')$ for $(f, g) \in G \wr G'$ and $(x, x') \in S \times S'$.

¹⁴To see this, consider a subset $U \subseteq S \times S'$. For $x' \in S'$, let $U_{x'} = \{x \in S : (x, x') \in U\}$ and $H_{x'} = G_{U_{x'}}$. Let $\mathcal{H} = \{H_{x'} : x' \in S'\}$ and let U' be the projection of U to S' . Then $(G \wr G')_U = \mathcal{H} \wr G'_{U'}$. Moreover, if $x' \in S'$ is not fixed by $G'_{U'}$, then $x' \notin U'$ and hence $U_{x'} = \emptyset$, which implies $H_{x'} = G$.

of \mathcal{C}'' from \mathcal{C} and \mathcal{C}' corresponds to a construction of an m -scheme on $S \times S'$ from those on S and S' . This is exactly Definition 6.13.

Chapter 7

SYMMETRIC GROUPS AND LINEAR GROUPS

Let G be a finite permutation group. Motivated by the \mathcal{P} -scheme algorithms developed in Chapter 3 and Chapter 5, we are interested in the problem of bounding the integer $d(G)$, introduced in Definition 2.8.

In this chapter, we study this problem for symmetric groups and linear groups with various special group actions.

Symmetric groups. For convenience, we introduce the following notation:

Definition 7.1. For $n \in \mathbb{N}^+$, define $d_{\text{Sym}}(n) := d(G)$, where G is the symmetric group $\text{Sym}(S)$ acting naturally on a finite set S of cardinality n .¹

Note that $d_{\text{Sym}}(n)$ is nondecreasing in n by Corollary 6.4. The best known general upper bound for $d_{\text{Sym}}(n)$ is

$$d_{\text{Sym}}(n) \leq \left(\frac{2}{\log 12} \right) \log n + O(1),$$

proven in [Gua09; Aro13] in different notations, based on the work of [Evd94; IKS09]. In Section 7.1, we review this result and interpret it as a result about \mathcal{P} -schemes.

In Section 7.3, we study the more general action of $\text{Sym}(S)$ on the set of k -subsets of S , where $1 \leq k \leq |S|$, and that on (an orbit of) the set of partitions of S . These actions are called the *standard action* of symmetric groups, and play an important role in the study of minimal base sizes of primitive permutation groups (see, e.g., [LS99]). Our results for these group actions will be used in Chapter 8.

Linear groups. Let V be a vector space of dimension $n \in \mathbb{N}^+$ over a finite field \mathbb{F}_q . We have the *general linear group* $\text{GL}(V)$ consisting of all the invertible linear transformations of V over \mathbb{F}_q . It is a subgroup of the *general semilinear group* $\Gamma\text{L}(V)$, which consists of all the invertible *semilinear transformations* of V . Here we say a map $\phi : V \rightarrow V$ is a *semilinear transformation* of V if

¹Clearly $d_{\text{Sym}}(n)$ only depends on n but not on S .

$\phi(x + y) = \phi(x) + \phi(y)$ and $\phi(cx) = \tau_\phi(c)\phi(x)$ hold for all $x, y \in V$ and $c \in \mathbb{F}_q$, where τ_ϕ is an automorphism of the field \mathbb{F}_q . We have the *natural action* of $\text{GL}(V)$ and that of $\Gamma\text{L}(V)$ on $V - \{0\}$, defined in the obvious way.

Denote by $\mathbb{P}V$ the *projective space* associated with V , i.e., $\mathbb{P}V$ is the set of equivalence classes of $V - \{0\}$ where $x, y \in V - \{0\}$ are equivalent iff $x = cy$ for some $c \in \mathbb{F}_q^\times$. Define the *projective linear group* $\text{PGL}(V) := \text{GL}(V)/\mathbb{F}_q^\times$ and the *projective semilinear group* $\text{P}\Gamma\text{L}(V) := \Gamma\text{L}(V)/\mathbb{F}_q^\times$, where \mathbb{F}_q^\times is identified with the subgroup of the *scalar linear transformations* of $\text{GL}(V)$ (resp. $\Gamma\text{L}(V)$) so that $c \in \mathbb{F}_q^\times$ sends $x \in V$ to cx . The natural action of $\text{GL}(V)$ (resp. $\Gamma\text{L}(V)$) on $V - \{0\}$ induces an action of $\text{PGL}(V)$ (resp. $\text{P}\Gamma\text{L}(V)$) on $\mathbb{P}V$, called the *natural action* of $\text{PGL}(V)$ (resp. $\text{P}\Gamma\text{L}(V)$) on $\mathbb{P}V$. Finally, when $V = \mathbb{F}_q^n$, we also use the notations $\text{GL}_n(q)$, $\Gamma\text{L}_n(q)$, $\text{PGL}_n(q)$, and $\text{P}\Gamma\text{L}_n(q)$.

We call the above groups $\text{GL}(V)$, $\Gamma\text{L}(V)$, $\text{PGL}(V)$, and $\text{P}\Gamma\text{L}(V)$ *linear groups*. In Section 7.4, we investigate $d(G)$ for the natural action of a linear group G . For convenience, we introduce the following notations:

Definition 7.2. *Let V be a vector space of dimension $n \in \mathbb{N}^+$ over a finite field \mathbb{F}_q . Define $d_{\text{GL}}(n, q) := d(G)$, where G is the permutation group $\text{GL}(V)$ acting naturally on $V - \{0\}$. Similarly define $d_{\Gamma\text{L}}(n, q)$, $d_{\text{PGL}}(n, q)$, and $d_{\text{P}\Gamma\text{L}}(n, q)$ by choosing G to be the permutation group $\Gamma\text{L}(V)$, $\text{PGL}(V)$, $\text{P}\Gamma\text{L}(V)$ acting naturally on $V - \{0\}$, $\mathbb{P}V$, $\mathbb{P}V$, respectively.²*

We show that the problems of bounding $d_{\text{GL}}(n, q)$, $d_{\Gamma\text{L}}(n, q)$, $d_{\text{PGL}}(n, q)$, and $d_{\text{P}\Gamma\text{L}}(n, q)$ are all equivalent: an upper bound $f(n, q)$ for any one of them implies an upper bound $f(n, q) + O(1)$ for the others. So it suffices to investigate just one of them.

Finally, we prove a bound

$$d_{\text{GL}}(n, q) \leq \left(\frac{\log q}{\log q + (\log 12)/4} \right) n + O(1),$$

slightly improving the trivial bounds.

Self-reduction. The results in Section 7.3 and Section 7.4 require a technique called *self-reduction of discreteness*, which we introduce in Section 7.2. It reduces discreteness of a strongly antisymmetric \mathcal{P} -scheme to discreteness of its restrictions

²Clearly these definitions only depend on n and q but not on V .

to stabilizer subgroups. In many cases, such a reduction greatly simplifies the problem. Our results in Chapter 8 also rely heavily on this technique.

7.1 The natural action of a symmetric group

We introduce the following notations about m -schemes:

Definition 7.3. For $n \in \mathbb{N}^+$, let $m(n)$ (resp. $m'(n)$) be the smallest positive integer such that any non-discrete antisymmetric $m(n)$ -scheme (resp. $m'(n)$ -scheme) on $[n]$ has a matching (resp. is not strongly antisymmetric).

It is easy to see that $m(n)$ and $m'(n)$ are nondecreasing in n . We also have $d_{\text{Sym}}(n) \leq m'(n) \leq m(n)$ by Lemma 2.7 and Lemma 2.10.

It was proven [Gua09] and independently in [Aro13] that $m(n) \leq \left(\frac{2}{\log 12}\right) \log n + O(1)$. We review the proof of this bound, starting from the following lemma:

Lemma 7.1. Let $\Pi = \{P_1, \dots, P_m\}$ be an antisymmetric m -scheme on a finite set S where $m \geq 3$. Suppose $B \in P_1$ satisfies $|B| \geq 3$. Let x be an element of B so that $\Pi|_x = \{P'_1, \dots, P'_{m-1}\}$ is an $(m-1)$ -scheme on $S - \{x\}$ (see Definition 6.3). Then at least one of the two conditions is satisfied.

1. There exists $B' \in P'_1$ contained in B satisfying $|B'| \leq (|B| - 1)/4$.
2. There exist distinct elements $y, z \in B - \{x\}$ such that for the $(m-2)$ -scheme $\Pi|_{x,y} = \{P''_1, \dots, P''_{m-2}\}$ on $S - \{x, y\}$, the block B'' of P''_1 containing z satisfies $|B''| \leq (|B| + 1)/12$. Furthermore, (x, y) , (y, z) , and (z, x) are in the same block of P_2 .

Proof. By replacing Π with $\Pi|_B$, we may assume Π is homogeneous and $S = B$. By antisymmetry, we know $|P_2|$ is even. If $|P_2| \geq 4$, there exists $B_1 \in P_2$ of cardinality at most $|B|(|B| - 1)/4$. Let $B' := \{y \in B : (x, y) \in B_1\}$. Then B' is a block of P'_1 by definition, and its cardinality is $|B_1|/|B| \leq (|B| - 1)/4$ by regularity of Π . And the first condition is met.

So assume $|P_2| = 2$. Then P_2 contains two blocks B_1 and B_2 of the same cardinality $|B|(|B| - 1)/2$. Choose $y \in B - \{x\}$ such that $(x, y) \in B_1$. Such an element y exists by regularity and homogeneity of Π . By Lemma 2.11 and Lemma 2.12, we have an antisymmetric association scheme $P(\Pi) = P_2 \cup \{1_B\}$ that has three blocks. By Lemma 2.20, the number of elements $z \in B - \{x, y\}$ satisfying

$(y, z), (z, x) \in B_1$ is precisely $(|B| + 1)/4 > 0$. The cardinality of the set $T := \{(a, b, c) : (a, b), (b, c), (c, a) \in B_1\}$ is then $|B_1| \cdot (|B| + 1)/4$. Choose $z \in B - \{x, y\}$ such that $(x, y, z) \in T$. Let B'_1, B'_2 , and B'_3 be the blocks of P_3 containing (x, y, z) , (y, z, x) and (z, x, y) respectively, which are all subsets of T . They have the same cardinality by invariance of Π , and are distinct by antisymmetry of Π . So $|B'_1| \leq |T|/3 = |B_1| \cdot (|B| + 1)/12$. By regularity of Π , the cardinality of the set $\{u \in S - \{x, y\} : (x, y, u) \in B'_1\}$ is $|B'_1|/|B_1| \leq (|B| + 1)/12$, and this set is exactly the block B'' of P'_1 containing z by definition. So the second condition is satisfied. \square

Lemma 7.1 implies the following recursive relation:

Lemma 7.2. For $n \geq 3$,

$$m(n) \leq \max \left\{ m \left(\frac{n-1}{4} \right) + 1, m \left(\frac{n+1}{12} \right) + 2 \right\}.$$

The inequality also holds for $m'(\cdot)$ in replaced of $m(\cdot)$.

Proof. Let $\Pi = \{P_1, \dots, P_m\}$ be a non-discrete antisymmetric m -scheme on a finite set S of cardinality n , where $m \geq 3$. Also assume

$$m \geq \max \left\{ m \left(\frac{n-1}{4} \right) + 1, m \left(\frac{n+1}{12} \right) + 2 \right\}.$$

We want to show that Π has a matching.

Choose $B \in P_1$ such that $|B| > 1$. Let x be an element of B and suppose $\Pi|_x = \{P'_1, \dots, P'_{m-1}\}$. Then $\Pi|_x$ is an antisymmetric $(m-1)$ -scheme on $S - \{x\}$. Note that $\Pi|_B$ is a homogeneous antisymmetric m -scheme on B by Lemma 6.14, which implies $|B| \geq 3$. Then either of the two conditions in Lemma 7.1 is satisfied.

If the first condition is satisfied, there exists $B' \in P'_1$ contained in B satisfying $|B'| \leq (|B| - 1)/4 \leq (n - 1)/4$. If $|B'| > 1$, we see $(\Pi|_x)|_{B'}$ is a non-discrete antisymmetric $(m-1)$ -scheme on B' . It has a matching since $m-1 \geq m((n-1)/4) \geq m(|B'|)$. So Π also has a matching by Lemma 6.3 and Lemma 6.14. On the other hand, if $|B'| = 1$, we let y be the unique element in B' and let B_1 be the block of P_2 containing (x, y) . Note that $|B'| = |B_1|/|B|$, which implies $|B_1| = |B|$. As $x, y \in B$, we have $\pi_1^2(B_1) = \pi_2^2(B_1) = B$. Then B_1 is a matching of Π .

Next assume the second condition is satisfied. So there exist distinct elements $y, z \in B - \{x\}$ such that for the $(m-2)$ -scheme $\Pi|_{x,y} = \{P''_1, \dots, P''_{m-2}\}$ on $S - \{x, y\}$, the

cardinality of the block B'' of P_1'' containing z is at most $(|B| + 1)/12 \leq (n + 1)/12$. Furthermore, (x, y) , (y, z) , and (z, x) are in the same block B_0 of P_2 . If $|B''| > 1$, we see $(\Pi|_{x,y})|_{B''}$ is a non-discrete antisymmetric $(m - 2)$ -scheme on B'' . It has a matching since $m - 2 \geq m((n + 1)/12) \geq m(|B''|)$. So Π also has a matching by Lemma 6.3 and Lemma 6.14. On the other hand, if $|B''| = 1$, we let B'_0 be the block of P_3 containing (x, y, z) . We have $\pi_1^3(B'_0) = \pi_3^3(B'_0) = B_0$ since $(x, y), (y, z) \in B_0$. Also note that $|B''| = |B'_0|/|B_0|$, which implies $|B_0| = |B'_0|$. So B'_0 is a matching of Π .

This proves the inequality for $m(\cdot)$. The proof for $m'(\cdot)$ is similar, and we leave it to the reader. \square

Theorem 7.1 ([Gua09; Aro13]). *For all $n \in \mathbb{N}^+$,*

$$m(n) \leq \left(\frac{2}{\log 12} \right) \log n + O(1).$$

More generally, an antisymmetric m -scheme $\Pi = \{P_1, \dots, P_m\}$ on a finite set S always has a matching if P_1 has a block B of cardinality $k > 1$ and $m \geq m(k)$. In particular it holds for sufficiently large $m = \left(\frac{2}{\log 12} \right) \log k + O(1)$.

Proof. Note $m(1) = 1$ and $m(2) = 2$. The first claim then follows from Lemma 7.2 and a simple induction. The second claim follows by considering $\Pi|_B$ and applying Lemma 6.14. \square

Theorem 7.1 implies a bound for $d_{\text{Sym}}(n)$, and also a bound for $d(G)$ by Corollary 6.3, where G is an arbitrary permutation group on a set of cardinality n :

Corollary 7.1. *Let G be a permutation group on a set of cardinality $n \in \mathbb{N}^+$. Then $d(G) \leq d_{\text{Sym}}(n) \leq \left(\frac{2}{\log 12} \right) \log n + O(1)$.*

We conclude this section with the following technical lemma, which is used later in the proof of Theorem 7.5.

Lemma 7.3. *Let G be a permutation group on a finite set S , and let \mathcal{P} be the corresponding system of stabilizers of depth m where $1 \leq m \leq |S|$. Let $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ be a strongly antisymmetric \mathcal{P} -scheme. Suppose \mathcal{C} is non-discrete on G_x for some $x \in S$. Then there exists $(x_1, \dots, x_m) \in S^{(m)}$ such that $C_{G_{x_1, \dots, x_m}}$ has a block of cardinality at least $2^{\left(\frac{\log 12}{4}\right)m^2 - O(m)}$.*

Proof. Let \mathcal{P}' be the system of stabilizers of depth m with respect to the natural action of $G' := \text{Sym}(S)$ on S . Let $\mathcal{C}' = \{C'_H : H \in \mathcal{P}'\}$ be the induction of \mathcal{C} to \mathcal{P}' (see Definition 6.4), which is strongly antisymmetric by Lemma 6.1 and is non-discrete on G'_x for $x \in S$ in the lemma since \mathcal{C} is non-discrete on G_x . Assume the lemma holds for $\text{Sym}(S)$, \mathcal{P}' , and an m -tuple $(y_1, \dots, y_m) \in S^{(m)}$, i.e., there exists $B' \in C'_{G'_{y_1, \dots, y_m}}$ of cardinality at least $2^{\left(\frac{\log 12}{4}\right)m^2 - O(m)}$. By Definition 6.4, we know B' is of the form $\phi_{G'_{y_1, \dots, y_m}, g}(B)$, where $g \in G'$, $\phi_{G'_{y_1, \dots, y_m}, g}$ is an injection from $(G \cap gG'_{y_1, \dots, y_m}g^{-1}) \setminus G$ to $G'_{y_1, \dots, y_m} \setminus G'$, and B is a block of $C_{G \cap gG'_{y_1, \dots, y_m}g^{-1}}$. Let $x_i = {}^g y_i$ for $i \in [m]$. Then $G \cap gG'_{y_1, \dots, y_m}g^{-1} = G_{x_1, \dots, x_m}$. So (x_1, \dots, x_m) and $B \in C_{G_{x_1, \dots, x_m}}$ satisfy the condition in the lemma.

Thus we may assume $G = \text{Sym}(S)$ and it acts naturally on S . By Lemma 2.12, it suffices to show that for any non-discrete strongly antisymmetric m -scheme $\Pi = \{P_1, \dots, P_m\}$ on S , the partition P_m has a block of cardinality at least $2^{\left(\frac{\log 12}{4}\right)m^2 - cm}$, where $c = O(1)$. We prove this claim by induction on m . The case $m = 1$ is trivial. For $m > 1$, assume the claim for $m' < m$. Let B_0 be a block of P_1 of cardinality $k > 1$. By Theorem 7.1, we have $m \leq \left(\frac{2}{\log 12}\right) \log k + c'$ for some $c' = O(1)$, or equivalently $k \geq 2^{\frac{\log 12}{2}(m - c')}$. Choose $x \in B_0$ and consider the $(m - 1)$ -scheme $\Pi' := \Pi|_x = \{P'_1, \dots, P'_{m-1}\}$ on $S - \{x\}$. It is strongly antisymmetric by Lemma 6.3. Let B_1 be a block of P'_1 contained in B_0 , which exists by compatibility of Π and the fact $k > 1$. If $|B_1| = 1$, we have seen in the proof of Lemma 7.2 that Π has matching, contradicting the assumption that Π is strongly antisymmetric. So $|B_1| > 1$. By Lemma 6.14, the homogeneous $(m - 1)$ -scheme $\Pi''|_{B_1} = \{P''_1, \dots, P''_{m-1}\}$ on B_1 is strongly antisymmetric. By the induction hypothesis, the partition P''_{m-1} has a block $B' \subseteq B_1^{(m-1)}$ of cardinality at least $2^{\left(\frac{\log 12}{4}\right)(m-1)^2 - c(m-1)}$. And B' is also a block of $P'_{m-1} \in \Pi'$ by definition and compatibility of Π' . Then $P_m \in \Pi$ has a block B containing (x, x_1, \dots, x_{m-1}) for all $(x_1, \dots, x_{m-1}) \in B'$. By regularity of Π , we have

$$|B| = |B_0||B'| \geq 2^{\frac{\log 12}{2}(m - c')} \cdot 2^{\left(\frac{\log 12}{4}\right)(m-1)^2 - c(m-1)} \geq 2^{\left(\frac{\log 12}{4}\right)m^2 - cm}$$

for sufficiently large $c = O(1)$. □

7.2 Self-reduction of discreteness

In this section, we prove a “self-reduction” lemma, which states that discreteness of a strongly antisymmetric \mathcal{P} -scheme is implied by discreteness of its restrictions to stabilizer subgroups.

We need the following technical lemma.

Lemma 7.4. *Suppose G is a finite group, \mathcal{P} is a subgroup system over G , and $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ is a \mathcal{P} -scheme. Suppose H_0, H_1, H_2 are subgroups in \mathcal{P} such that $H_0 \subseteq H_1 \cap H_2$ and $\mathcal{C}|_{H_1}, \mathcal{C}|_{H_2}$ are both discrete on H_0 . For $i = 0, 1, 2$, let B_i be the block of C_{H_i} containing $H_i e \in H_i \setminus G$. Then $(\pi_{H_0, H_2}|_{B_0}) \circ (\pi_{H_0, H_1}|_{B_0})^{-1}$ is a well-defined bijection from B_1 to B_2 sending $H_1 e \in H_1 \setminus G$ to $H_2 e \in H_2 \setminus G$.*

Proof. Note that $\pi_{H_0, H_1}|_{B_0}$ is a surjective map from B_0 to B_1 sending $H_0 e$ to $H_1 e$, and $\pi_{H_0, H_2}|_{B_0}$ is a surjective map from B_0 to B_2 sending $H_0 e$ to $H_2 e$. So it suffices to prove that these two maps are injective. The set $B_0 \cap (H_0 \setminus H_1)$ contains $H_0 e$ and is a block of $C_{H_0}|_{H_1} \in \mathcal{C}|_{H_1}$ by the definition of restriction (Definition 6.2). By discreteness of $\mathcal{C}|_{H_1}$ on H_0 , this set is just the singleton $\{H_0 e\}$. On the other hand, the set $H_0 \setminus H_1 \subseteq H_0 \setminus G$ is precisely the preimage of $H_1 e$ under $\pi_{H_0, H_1} : H_0 \setminus G \rightarrow H_1 \setminus G$. So $B_0 \cap (H_0 \setminus H_1)$ is the preimage of $H_1 e$ under $\pi_{H_0, H_1}|_{B_0} : B_0 \rightarrow B_1$. By regularity of \mathcal{C} , the map $\pi_{H_0, H_1}|_{B_0}$ is injective. Similarly $\pi_{H_0, H_2}|_{B_0}$ is also injective. \square

The bijection in Lemma 7.4 can be used to separate elements in a strongly antisymmetric \mathcal{P} -scheme:

Lemma 7.5. *Let G be a finite group acting on a finite set S , and let $x \in S$. Let \mathcal{P} be a subgroup system over G such that $G_x \in \mathcal{P}$, and let $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ be a \mathcal{P} -scheme. Suppose $y = {}^g x$ and $z = {}^{g'} x$ in S satisfy (1) $G_y, G_z, G_{y,z} \in \mathcal{P}$ and (2) $\mathcal{C}|_{G_y}$ and $\mathcal{C}|_{G_z}$ are both discrete on $G_{y,z}$. Then there exists a bijection between blocks of C_{G_x} sending $G_x g^{-1}$ to $G_x g'^{-1}$ that can be written as a composition of conjugations, projections and their inverses between blocks of $C_{G_x}, C_{G_y}, C_{G_z}$ and $C_{G_{y,z}}$. In particular, if \mathcal{C} is strongly antisymmetric, then $G_x g^{-1}$ and $G_x g'^{-1}$ are in different blocks of C_{G_x} .*

Proof. Let B_0 (resp. B_1, B_2) be the block of $C_{G_{y,z}}$ (resp. C_{G_y}, C_{G_z}) containing $G_{y,z} e$ (resp. $G_y e, G_z e$). By Lemma 7.4, the map $\pi_{G_{y,z}, G_z}|_{B_0} \circ (\pi_{G_{y,z}, G_y}|_{B_0})^{-1}$ is a bijection from B_1 to B_2 sending $G_y e$ to $G_z e$. Let B'_1 and B'_2 be the blocks of C_{G_x} containing $G_x g^{-1}$ and $G_x g'^{-1}$ respectively. We have the conjugations $c_{G_x, g}|_{B'_1} : B'_1 \rightarrow B_1$ sending $G_x g^{-1}$ to $G_y e$ and $c_{G_z, g'^{-1}}|_{B_2} : B_2 \rightarrow B'_2$ sending $G_z e$ to $G_x g'^{-1}$. Then the map

$$c_{G_z, g'^{-1}}|_{B_2} \circ \pi_{G_{y,z}, G_z}|_{B_0} \circ (\pi_{G_{y,z}, G_y}|_{B_0})^{-1} \circ c_{G_x, g}|_{B'_1}$$

is a bijection from B'_1 to B'_2 sending $G_x g^{-1}$ to $G_x g'^{-1}$. \square

This provides a way of proving discreteness of a strongly antisymmetric \mathcal{P} -scheme using discreteness of its restrictions to stabilizers. For example, if \mathcal{C} is strongly antisymmetric and the conditions in Lemma 7.5 hold for all pairs $(y, z) \in Gx \times Gx$, then \mathcal{C} is discrete on Gx . In fact, we only need to verify the conditions for a subset of pairs that form a connected graph:

Lemma 7.6 (self-reduction lemma). *Let G be a finite group acting on a finite set S , and let $x \in S$. Let \mathcal{P} be a subgroup system over G such that $G_x \in \mathcal{P}$, and let $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ be a strongly antisymmetric \mathcal{P} -scheme. Suppose R is a subset of $S \times S$ satisfying the following conditions:*

1. *For all $(y, z) \in R$, it holds that (1) $G_y, G_z, G_{y,z} \in \mathcal{P}$ and (2) $\mathcal{C}|_{G_y}$ and $\mathcal{C}|_{G_z}$ are both discrete on $G_{y,z}$.*
2. *Let \mathcal{G}_R be the undirected graph on S such that $\{y, z\}$ is an edge iff $(y, z) \in R$ or $(z, y) \in R$. Then Gx is contained in a connected component of \mathcal{G}_R (in particular, this condition is satisfied if \mathcal{G}_R is connected).*

Then \mathcal{C} is discrete on Gx .

Proof. For $y \in S$, denote by B_y the block of C_{G_y} containing $G_y e \in G_y \setminus G$. For $(y, z) \in S \times S$, write $y \sim z$ if there exists a bijection $\tau : B_y \rightarrow B_z$ sending $G_y e$ to $G_z e$ such that τ is a composition of maps of the form $\pi_{H,H'}|_B$ or $(\pi_{H,H'}|_B)^{-1}$ (where $H, H' \in \mathcal{P}$ and B is block of C_H). Then \sim is an equivalence relation on S . By the first condition and Lemma 7.4, we have $y \sim z$ for all $(y, z) \in R$. And by the second condition, we have $y \sim z$ for all $(y, z) \in Gx \times Gx$.

Consider any $g, g' \in G$ and let $y = {}^g x, z = {}^{g'} x \in Gx$. Let $\tau : B_y \rightarrow B_z$ be a bijection sending $G_y e$ to $G_z e$ as above. Let B and B' be the blocks of C_{G_x} containing $G_x g^{-1}$ and $G_x g'^{-1}$ respectively. We have the conjugations $c_{G_x, g}|_B : B \rightarrow B_y$ sending $G_x g^{-1}$ to $G_y e$ and $c_{G_x, g'^{-1}}|_{B_z} : B_z \rightarrow B'$ sending $G_z e$ to $G_x g'^{-1}$. Then the map

$$c_{G_x, g'^{-1}}|_{B_z} \circ \tau \circ c_{G_x, g}|_B$$

is a bijection from B to B' sending $G_x g^{-1}$ to $G_x g'^{-1}$. In particular, if $G_x g^{-1} \neq G_x g'^{-1}$, then $B \neq B'$ by strong antisymmetry of \mathcal{C} . As $g, g' \in G$ are arbitrary, we know \mathcal{C} is discrete on Gx . \square

7.3 The actions of symmetric groups on k -subsets or partitions

Let S be a finite set of cardinality n , and let $G = \text{Sym}(S)$. For $k \in [n]$, the natural action of G on S induces a (transitive) action of G on the set of k -subsets (i.e., subsets of cardinality k) of S . Similarly, it induces an action of G on (an orbit of) the set of partitions P of S , given by ${}^gP := \{{}^gB : B \in P\}$.

In these cases, we expect to have a bound $d(G) = O(\log n)$ as we have in Section 7.1. Let S' be the underlying set on which G acts. The naive approach is to embed G in $\text{Sym}(S')$ and apply Corollary 7.1. In general, however, the cardinality of S' is much larger than n . For example, we have $|S'| = \binom{n}{k}$ for the action of G on the set of k -subsets of S , and hence Corollary 7.1 only implies the bound $d(G) = O(\log |S'|) = O(k \log n)$. The same problem exists for the action of G on an orbit of the set of partitions of S , in which case $|S'|$ is the number of partitions of S into subsets with prescribed cardinalities.

In this section, we extend the result in Section 7.1 and show that in the above cases, we have $d(G) \leq d_{\text{Sym}}(n) + O(1) = O(\log n)$. In fact, we prove more general criteria for a subgroup system \mathcal{P} over G (or more generally, over a subgroup $H \subseteq G$) to have the property that all strongly antisymmetric \mathcal{P} -schemes are discrete on all $x \in S'$.³ It is possible to design a subgroup system \mathcal{P} of complexity $|S'|^{O(1)} n^{O(\log n)}$ that satisfies these criteria. An algorithm of constructing the corresponding collection of number fields will be given in Chapter 8.

The action of $\text{Sym}(S)$ on the set of k -subsets of S . Suppose S is a finite set of cardinality n and consider the action of $G = \text{Sym}(S)$ on the set S' of k -subsets of S . We say two elements $x, y \in S'$ are *adjacent* if there exists $g \in G$ sending x to y and g is a transposition (i.e. 2-cycle) on S . The following technical lemma is needed:

Lemma 7.7. *For all adjacent $x, y \in S$ and $z \in G_{x,y}$ adjacent to y , it holds that $|G_{x,y}z| \leq n$.*

Proof. Choose $h \in G_x$ that sends y to z . As x and y are adjacent, we know $x = {}^h x$ and $z = {}^h y$ are also adjacent. Let $u = x \cap y$ (as the intersection of two k -subsets). As x and y are adjacent, there exist distinct elements $a, b \in S$ such that $x = u \cup \{a\}$ and $y = u \cup \{b\}$. Then $G_{x,y}$ fixes u setwisely as well as a, b . If $b \notin z$, we have

³To derive $d(G) \leq d_{\text{Sym}}(n) + O(1)$, we only need the case $H = G$. The more general setting $H \subseteq G$ is needed for applications in Chapter 8.

$z = u \cup \{c\}$ for some $c \in S$ since y and z are adjacent. In this case, as $G_{x,y}$ fixes the subset u of z of cardinality $k - 1$ setwisely, we have $|G_{x,y}z| \leq |S| = n$, as desired. Next assume $b \in z$. Since x and z are adjacent, we have $z = (x - \{b'\}) \cup \{b\}$ for some $b' \in x$. As $G_{x,y}$ fixes x setwisely as well as $a, b \in S$, the elements in $G_{x,y}z$ are of the form $(x - \{b''\}) \cup \{b\}$ where $b'' \in x$. In this case we have $|G_{x,y}z| \leq |x| = k \leq n$. \square

We state a criterion for a subgroup system \mathcal{P} over a subgroup $H \subseteq G$ to have the property that all strongly antisymmetric \mathcal{P} -schemes are discrete on all $x \in S'$.

Theorem 7.2. *Let G , n , and S' be as above, and let H be a subgroup of G . Suppose \mathcal{P} is a subgroup system over H satisfying the following conditions:*

1. $H_x, H_{x,y} \in \mathcal{P}$ for all $x, y \in S'$.
2. $H_{\{x,y\} \cup T} \in \mathcal{P}$ for all $x, y, z \in S'$ and $T \subseteq H_{x,y}z$ satisfying $|H_{x,y}z| \leq n$ and $1 \leq |T| \leq d_{\text{Sym}}(n)$.

Then all strongly antisymmetric \mathcal{P} -schemes are discrete on H_x for all $x \in S'$.

Proof. Let $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ be a strongly antisymmetric \mathcal{P} -scheme. We want to prove that \mathcal{C} is discrete on H_x for all $x \in S'$. As G is generated by transpositions on S , by Lemma 7.6, we just need to verify for all adjacent $x, y \in S'$ that (1) $H_x, H_y, H_{x,y} \in \mathcal{P}$, and (2) $\mathcal{C}|_{H_x}$ and $\mathcal{C}|_{H_y}$ are discrete on $H_{x,y}$. Fix adjacent $x, y \in S'$. Note that (1) follows from the first condition in the theorem.

So it remains to prove that $\mathcal{C}|_{H_x}$ is discrete on $H_{x,y}$ (the claim for $\mathcal{C}|_{H_y}$ is symmetric). This is trivial if $x = y$. So assume $x \neq y$. We claim that for all $z, w \in H_{x,y} \subseteq G_{x,y}$, there exists a sequence of elements $u_0, \dots, u_t \in G_{x,y}$ such that $u_0 = z$, $u_t = w$, and u_{i-1}, u_i are adjacent for $i \in [t]$. This follows from the fact that $G_x \cong \text{Sym}(x) \times \text{Sym}(S - x)$ is generated by transpositions on S . Let $\mathcal{P}' := \mathcal{P}|_{H_x}$ and $\mathcal{C}' := \mathcal{C}|_{H_x}$. By Lemma 7.6 and the previous claim, it remains to show that for all adjacent $z, w \in G_{x,y}$, it holds that (a) $(H_x)_z, (H_x)_w, (H_x)_{z,w} \in \mathcal{P}'$, or equivalently $H_{x,z}, H_{x,w}, H_{x,z,w} \in \mathcal{P}$, and (b) $\mathcal{C}'|_{(H_x)_z}$ and $\mathcal{C}'|_{(H_x)_w}$ are discrete on $(H_x)_{z,w}$. Fix such $z, w \in G_{x,y}$. Note that z and w are adjacent to x since y is adjacent to x . It follows from Lemma 7.7 that $|H_{x,z}w| \leq |G_{x,z}w| \leq n$. Then (a) follows from the two conditions in the theorem.

It remains to show that $\mathcal{C}'|_{(H_x)_z}$ is discrete on $(H_x)_{z,w}$ (the claim for $\mathcal{C}'|_{(H_x)_w}$ is symmetric). Let $\mathcal{P}'' := \mathcal{P}'|_{(H_x)_z}$. By the second condition of the theorem and the fact $|H_{x,z}w| \leq n$, we have $H_{\{x,z\} \cup T} \in \mathcal{P}''$ for all $T \subseteq H_{x,z}w$ satisfying $1 \leq |T| \leq d_{\text{Sym}}(n)$. This means that \mathcal{P}'' contains the system of stabilizers of depth $d_{\text{Sym}}(n)$ with respect to the action of $H_{x,z}$ on $H_{x,z}w$. By Corollary 6.3 and the fact $|H_{x,z}w| \leq n$, we see all strongly antisymmetric \mathcal{P}'' -schemes are discrete on $(H_{x,z})_w = (H_x)_{z,w}$. Finally, note that $\mathcal{C}'|_{(H_x)_z}$ is strongly antisymmetric by Lemma 6.3, and hence discrete on $(H_x)_{z,w}$, as desired. \square

Choosing $H = G$ in Theorem 7.2, we obtain

Corollary 7.2. $d(G) \leq d_{\text{Sym}}(n) + 2$.

The action of $\text{Sym}(S)$ on the set of partitions of S . Suppose S is a finite set of cardinality n and consider the action of $G = \text{Sym}(S)$ on an orbit S' of the set of partitions of S . We prove an analogue of Theorem 7.2 for this case. The following notations are needed: again, we call two elements $x, y \in S'$ *adjacent* if there exists $g \in G$ sending x to y and g is a transposition on S . For $x, y, z \in S'$, write $y \sim_x z$ if there exists $g \in G_x$ sending y to z such that either (1) g is a transposition on S fixing all the blocks of x setwisely, or (2) $x - y \neq x - z$, and g exchanges two blocks of x while fixing the other blocks of x pointwisely.

We also need the following technical lemma:

Lemma 7.8. *For all adjacent $x, y \in S'$ and $z \in G_x y$ satisfying $y \sim_x z$, it holds that $|G_{x,y} z| \leq 4n$.*

Proof. We may assume $x \neq y$. As x and y are adjacent, there exists a transposition $(a b)$ of S sending x to y where $a \in B_1, b \in B_2$ and B_1, B_2 are distinct blocks of x . So we have

$$y = (x - \{B_1, B_2\}) \cup \{(B_1 - \{a\}) \cup \{b\}, (B_2 - \{b\}) \cup \{a\}\}. \quad (7.1)$$

Fix $h \in G_x$ sending y to z such that either (1) h is a transposition on S fixing all the blocks of x setwisely, or (2) $x - z \neq x - y$ and h exchanges two blocks of x while fixing the other blocks of x pointwisely. We claim that in either case, h fixes at least one of B_1 and B_2 pointwisely. This is obvious in Case (1). And in Case (2), if h fixes neither B_1 nor B_2 pointwisely, it exchanges B_1 and B_2 . But then we have $x - y = x - z = \{B_1, B_2\}$, contradicting the assumption.

So assume h fixes B_1 pointwisely (the other case is symmetric). Consider arbitrary $w = {}^g z \in G_{x,y}z$ where $g \in G_{x,y}$. We have

$$w = {}^{gh(a\ b)}x = {}^{gh(a\ b)(gh)^{-1}}x = (a' \ b')x,$$

where $a' = {}^{gh}a$ and $b' = {}^{gh}b$. So w is determined by the pair (a', b') .

There are at most n choices of $b' \in S$. Now consider the number of choices of a' . Note that $a' = {}^{gh}a = {}^g a$ since h fixes B_1 pointwisely. As $\{a\} \in y|_{B_1}$, we see $\{a'\} \in {}^g y|_{{}^g B_1} = y|_{{}^g B_1}$. As g fixes x and y , it fixes $x - y = \{B_1, B_2\}$ setwisely. So ${}^g B_1 \in \{B_1, B_2\}$. It follows that $\{a'\}$ is in $y|_{B_1}$ or $y|_{B_2}$. By (7.1), we see $\{a'\}$ equals $\{a\}$, $\{b\}$, $B_1 - \{a\}$ or $B_2 - \{b\}$. So the number of choices of a' is at most four. Therefore $|G_{x,y}z| \leq 4n$. \square

We have following criterion for a subgroup system \mathcal{P} over a subgroup $H \subseteq G$ to have the property that all strongly antisymmetric \mathcal{P} -schemes are discrete on all $x \in S'$.

Theorem 7.3. *Let G , n , and S' be as above, and let H be a subgroup of G . Suppose \mathcal{P} is a subgroup system over H satisfying the following conditions:*

1. $H_x, H_{x,y} \in \mathcal{P}$ for all $x, y \in S'$.
2. $H_{\{x,y\} \cup T} \in \mathcal{P}$ for all $x, y, z \in S'$ and $T \subseteq H_{x,y}z$ satisfying $|H_{x,y}z| \leq 4n$ and $1 \leq |T| \leq d_{\text{Sym}}(4n)$.

Then all strongly antisymmetric \mathcal{P} -schemes are discrete on H_x for all $x \in S'$.

Proof. Let $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ be a strongly antisymmetric \mathcal{P} -scheme. We want to prove that \mathcal{C} is discrete on H_x for all $x \in S'$. As G is generated by transpositions on S , by Lemma 7.6, we just need to verify for all adjacent $x, y \in S'$ that (1) $H_x, H_y, H_{x,y} \in \mathcal{P}$, and (2) $\mathcal{C}|_{H_x}$ and $\mathcal{C}|_{H_y}$ are discrete on $H_{x,y}$. Fix adjacent $x, y \in S'$. Note that (1) follows from the first condition in the theorem. So it remains to prove that $\mathcal{C}|_{H_x}$ is discrete on $H_{x,y}$ (the claim for $\mathcal{C}|_{H_y}$ is symmetric). This obviously holds if $x = y$. So assume $x \neq y$.

We claim that for all $z, w \in H_{x,y} \subseteq G_x y$, there exists a sequence of elements $u_0, \dots, u_t \in G_x y$ such that $u_0 = z$, $u_t = w$ and $u_{i-1} \sim_x u_i$ for $i \in [t]$. To see this, note that we can choose distinct elements u_0, \dots, u_t such that $u_0 = z$, $u_t = w$, and for $i \in [t]$, u_{i-1} is sent to u_i by some $g_i \in G_x$ such that g_i is in either of the following two cases:

1. g_i is a transposition on S fixing the blocks of x setwisely, or
2. g_i exchanges two blocks of x while fixing the other blocks of x pointwisely.

This is because G_x is generated by such permutations g_i . Furthermore, if g_i is in the latter case, we may assume $x - u_{i-1} \neq x - u_i$. To see this, note that u_{i-1} and u_i are adjacent to x since y is adjacent to x . So there exist transpositions $(a\ b)$ and $(a'\ b')$ on S sending x to u_{i-1} and u_i respectively. Choose $B_1, B_2, B'_1, B'_2 \in x$ such that $a \in B_1, b \in B_2, a' \in B'_1$ and $b' \in B'_2$. Then $x - u_{i-1} = \{B_1, B_2\}$ and $x - u_i = \{B'_1, B'_2\}$. Suppose $x - u_{i-1} = x - u_i$. Then by exchanging a' with b' and B'_1 with B'_2 if necessary, we may assume $a, a' \in B_1$ and $b, b' \in B_2$. So we have

$$u_{i-1} = x - \{B_1, B_2\} \cup \{(B_1 - \{a\}) \cup \{b\}, (B_2 - \{b\}) \cup \{a\}\}$$

and

$$u_i = x - \{B_1, B_2\} \cup \{(B_1 - \{a'\}) \cup \{b'\}, (B_2 - \{b'\}) \cup \{a'\}\}.$$

If $a = a'$. Then $u_i = {}^{(b\ b')}u_{i-1}$ and we may replace g_i by $(b\ b') \in G_x$ which is in the first case above. Similarly, if $b = b'$, we may replace g_i by $(a\ a') \in G_x$. Finally, if $a \neq a'$ and $b \neq b'$, we insert $u'_i = {}^{(a\ a')}u_{i-1}$ into the sequence between u_{i-1} and u_i , so that $u_i = {}^{(b\ b')}u'_i$. It follows that we may always assume $u_{i-1} \sim_x u_i$ for all $i \in [t]$.

Let $\mathcal{P}' := \mathcal{P}|_{H_x}$ and $\mathcal{C}' := \mathcal{C}|_{H_x}$. By Lemma 7.6 and the previous paragraph, it suffices to show, for all $z, w \in G_x y$ satisfying $z \sim_x w$, that (a) $H_{x,z}, H_{x,w}, H_{x,z,w} \in \mathcal{P}'$, and (b) $\mathcal{C}'|_{(H_x)_z}$ and $\mathcal{C}'|_{(H_x)_w}$ are discrete on $(H_x)_{z,w}$. Fix such $z, w \in G_x y$. Note that z and w are adjacent to x since y is adjacent to x . It follows from Lemma 7.8 that $|H_{x,z,w}| \subseteq |G_{x,z,w}| \leq 4n$. Then (a) follows from the two conditions in the theorem.

It remains to show that $\mathcal{C}'|_{(H_x)_z}$ is discrete on $(H_x)_{z,w}$ (the claim for $\mathcal{C}'|_{(H_x)_w}$ is symmetric). Let $\mathcal{P}'' := \mathcal{P}'|_{(H_x)_z}$. By the second condition of the theorem and the fact $|H_{x,z,w}| \leq 4n$, we have $H_{\{x,z\} \cup T} \in \mathcal{P}''$ for all $T \subseteq H_{x,z,w}$ satisfying $1 \leq |T| \leq d_{\text{Sym}}(4n)$. This means that \mathcal{P}'' contains the system of stabilizers of depth $d_{\text{Sym}}(4n)$ with respect to the action of $H_{x,z}$ on $H_{x,z,w}$. By Corollary 6.3 and the fact $|H_{x,z,w}| \leq 4n$, we see all strongly antisymmetric \mathcal{P}'' -schemes are discrete on $(H_x)_{z,w}$. Finally note that $\mathcal{C}'|_{(H_x)_z}$ is strongly antisymmetric by Lemma 6.3, and hence is discrete on $(H_x)_{z,w}$, as desired. \square

Note $d_{\text{Sym}}(4n) \leq d_{\text{Sym}}(n) + O(1)$ by Lemma 7.2 and Theorem 2.1. Choosing $H = G$ in Theorem 7.3, we obtain

Corollary 7.3. $d(G) \leq d_{\text{Sym}}(4n) + 2 \leq d_{\text{Sym}}(n) + O(1)$.

7.4 The natural actions of linear groups

In this section, we show that $d_{\text{GL}}(n, q)$, $d_{\text{GL}}(n, q)$, $d_{\text{PGL}}(n, q)$, $d_{\text{PGL}}(n, q)$ are equal up to an additive constant. In addition, we prove an upper bound for $d_{\text{GL}}(n, q)$, slightly improving the trivial bounds.

Equivalence between various linear groups. We have the following theorem.

Theorem 7.4. *For $f_1, f_2 \in \{d_{\text{GL}}, d_{\text{GL}}, d_{\text{PGL}}, d_{\text{PGL}}\}$, there exists a constant $c \in \mathbb{N}$ such that $f_1(n, q) \leq f_2(n, q) + c$ holds for all $n \in \mathbb{N}^+$ and prime powers q . And if $f_2 = d_{\text{GL}}$, choosing $c = 6$ suffices.*

We break Theorem 7.4 into six inequalities, corresponding to the the arrows in the following diagram.

$$\begin{array}{ccc} \text{GL}(V) & \longleftrightarrow & \Gamma\text{L}(V) \\ & \updownarrow & \\ \text{PGL}(V) & \longleftrightarrow & \text{P}\Gamma\text{L}(V) \end{array}$$

Fix $n \in \mathbb{N}^+$, a prime power q and a vector space V of dimension n over \mathbb{F}_q from now on. By Lemma 6.3 and the facts $\text{GL}_n(q) \subseteq \Gamma\text{L}_n(q)$ and $\text{PGL}_n(q) \subseteq \text{P}\Gamma\text{L}_n(q)$, we have

Lemma 7.9. $d_{\text{GL}}(n, q) \leq d_{\text{GL}}(n, q)$ and $d_{\text{PGL}}(n, q) \leq d_{\text{PGL}}(n, q)$.

In the other direction, we have

Lemma 7.10. $d_{\text{GL}}(n, q) \leq d_{\text{GL}}(n, q) + 2$.

Proof. Let $G = \Gamma\text{L}(V)$, $S = V - \{0\}$, and $m = d_{\text{GL}}(n, q) + 2 \geq 3$. Let \mathcal{P} be the system of stabilizers of depth m over G with respect to the natural action of G on S . Let \mathcal{C} be a strongly antisymmetric \mathcal{P} -scheme. Fix $x \in S$. We want to show that \mathcal{C} is discrete on G_x . Let α be an element in \mathbb{F}_q^\times not contained in any proper subfield of \mathbb{F}_q . By Lemma 2.3, it suffices to show that \mathcal{C} is discrete on $G_{x, \alpha x}$. Let G act diagonally on $S \times S$ and let O be the G -orbit of $(x, \alpha x)$. By Lemma 7.6, it suffices to prove, for all $u, v \in O$, that (1) $G_u, G_v, G_{u,v} \in \mathcal{P}$ and (2) $\mathcal{C}|_{G_u}$ and $\mathcal{C}|_{G_v}$ are discrete on $G_{u,v}$.

Fix $u, v \in O$. Note we have $u = (y, \beta y)$ and $v = (z, \gamma z)$ for some $y, z \in S$ and $\beta, \gamma \in \mathbb{F}_q^\times$. And β, γ are not contained in any proper subfield of \mathbb{F}_q . Let $g \in G_u$. Then g sends βy to $\tau_g(\beta)y = \beta y$ where τ_g is the automorphism of \mathbb{F}_q^\times determined by g . So τ_g fixes β , i.e., β is in the fixed field of the cyclic group generated by τ_g . As β is not in any proper subfield of \mathbb{F}_q , we conclude that τ_g is the identity. So $G_u \subseteq \text{GL}(V)$. We have $G_u = G_{y, \beta y}$, $G_v = G_{z, \gamma z}$, and $G_{u,v} = G_{y, \beta y, z, \gamma z} = G_{y, \beta y, z}$. As $m \geq 3$, these subgroups are all in \mathcal{P} .

It remains to prove that $\mathcal{C}|_{G_u}$ is discrete on $G_{u,v}$ (the claim for $\mathcal{C}|_{G_v}$ is symmetric). Let \mathcal{P}' be the system of stabilizers of depth $m-2$ over G_u with respect to the action of G_u on S . As $G_u \subseteq \text{GL}(V)$, we have $d(G_u) \leq d_{\text{GL}}(n, q) = m-2$ by Corollary 6.3. So all strongly antisymmetric \mathcal{P}' -schemes are discrete on $(G_u)_z = G_{u,v}$. Also note $\mathcal{P}' \subseteq \mathcal{P}|_{G_u}$ since $G_u = G_{y, \beta y}$. It follows that all strongly antisymmetric $\mathcal{P}|_{G_u}$ -schemes are discrete on $G_{u,v}$. As $\mathcal{C}|_{G_u}$ is strongly antisymmetric by Lemma 6.2, it is discrete on $G_{u,v}$, as desired. \square

Similarly, we have

Lemma 7.11. $d_{\text{PFL}}(n, q) \leq d_{\text{PGL}}(n, q) + 4$.

Proof. Let $G = \text{PFL}(V)$, $S = \mathbb{P}V$, and $m = d_{\text{GL}}(n, q) + 4 \geq 5$. Let \mathcal{P} be the system of stabilizers of depth m over G with respect to the natural action of G on S . Let \mathcal{C} be a strongly antisymmetric \mathcal{P} -scheme. By Lemma 7.6, it suffices to prove, for all $(w, w') \in S^{(2)}$, that $\mathcal{C}|_{G_w}$ is discrete on $G_{w,w'}$. Fix $(w, w') \in S^{(2)}$. Again by Lemma 7.6, it suffices to prove, for all $(x, x') \in (G_w w')^{(2)}$, that $\mathcal{C}|_{G_{w,x}}$ is discrete on $G_{w,x,x'}$ (note $G_{w,x}, G_{w,x'}, G_{w,x,x'} \in \mathcal{P}$ since $m \geq 3$).

Fix $(x, x') \in (G_w w')^{(2)}$. Choose representatives $\tilde{w}, \tilde{x}, \tilde{x}' \in V - \{0\}$ of w, x and x' respectively. Note that \tilde{w}, \tilde{x} and \tilde{x}' are pairwise linearly independent over \mathbb{F}_q since w, x, x' are distinct. Let α be an element in \mathbb{F}_q^\times not contained in any proper subfield of \mathbb{F}_q . Define $\tilde{y} = \tilde{w} + \alpha \tilde{x} \neq 0$ and let y be the element in S represented by \tilde{y} . Consider the diagonal action of $G_{w,x}$ on S^2 and let O be the orbit of (x', y) under this action. We have $(G_{w,x})_{(x',y)} = G_{w,x,x',y} \in \mathcal{P}|_{G_x}$ since $m \geq 4$. By Lemma 2.3, it suffices to prove that $\mathcal{C}|_{G_{w,x}}$ is discrete on $(G_{w,x})_{(x',y)}$. Let $G' = G_{w,x}$. Applying Lemma 7.6 to the action of G' on O , we see that it suffices to prove for all $u, v \in O$ that (1) $G'_u, G'_v, G'_{u,v} \in \mathcal{P}|_{G'}$ and (2) $\mathcal{C}|_{G'_u}$ and $\mathcal{C}|_{G'_v}$ are discrete on $G'_{u,v}$.

Fix $u = (x'_1, y_1) = g^1(x', y)$ and $v = (x'_2, y_2) = g^2(x', y)$ in O , where $g_1, g_2 \in G'$. Lift g_1 to $\tilde{g}_1 \in \Gamma L(V)$. As $g_1 \in G_x$, we have $\tilde{g}_1 \tilde{x} = c \tilde{x}$ for unique $c \in \mathbb{F}_q^\times$. Define

$\tilde{x}'_1 = \tilde{g}_1 \tilde{x}'$ and $\tilde{y}_1 = \tilde{g}_1 \tilde{y}$ so that they are representatives of x'_1 and y_1 respectively. Consider arbitrary $g \in G'_u = G_{w,x,x'_1,y_1}$. We claim $g \in \text{PGL}(V)$. To see this, lift g to $\tilde{g} \in \Gamma\text{L}(V)$. Note that $\tilde{y}_1 = \tilde{g}_1(\tilde{w} + \alpha\tilde{x}) = \tilde{g}_1\tilde{w} + \alpha_1\tilde{g}_1\tilde{x}$ where $\alpha_1 = \tau_{\tilde{g}_1}(\alpha)$ and $\tau_{\tilde{g}_1}$ is the automorphism of \mathbb{F}_q^\times determined by \tilde{g}_1 . Here $\tilde{g}_1\tilde{w}$ and $\tilde{g}_1\tilde{x}$ are collinear with \tilde{w} and \tilde{x} respectively since $g_1 \in G_{w,x}$. And $\tilde{g}_1\tilde{w}, \tilde{g}_1\tilde{x}$ are linearly independent over \mathbb{F}_q since \tilde{w} and \tilde{x} are linearly independent. As $g \in G_{w,x,x',y_1}$, we see that \tilde{g} scales $\tilde{g}_1\tilde{w}, \tilde{g}_1\tilde{x}$ and $\tilde{y}_1 = \tilde{g}_1\tilde{w} + \alpha_1\tilde{g}_1\tilde{x}$. Therefore $\tau_{\tilde{g}}(\alpha_1) = \alpha_1$, where $\tau_{\tilde{g}}$ is the automorphism of \mathbb{F}_q^\times determined by \tilde{g} . So α_1 is in the fixed field of the cyclic group generated by $\tau_{\tilde{g}}$. But $\alpha_1 = \tau_{\tilde{g}_1}(\alpha)$ is not in any proper subfield of \mathbb{F}_q . It follows that $\tau_{\tilde{g}}$ is the identity. So we have $\tilde{g} \in \text{GL}(V)$ and hence $g \in \text{PGL}(V)$. We conclude that $G'_u = \text{PGL}(V)_{w,x,x'_1,y_1}$, and similarly $G'_v = \text{PGL}(V)_{w,x,x'_2,y_2}$. Moreover, observe that \tilde{g} above scales \tilde{w} and \tilde{x} by the same factor since g fixes y_1 . So it also scales any vector in the span of \tilde{w} and \tilde{x} over \mathbb{F}_q . We know \tilde{y}_1 is in this span and by the same argument, so is \tilde{y}_2 . So g fixes y_2 . This shows $G'_{u,v} = \text{PGL}(V)_{w,x,x'_1,y_1,x'_2}$. We then have $G'_u, G'_v, G'_{u,v} \in \mathcal{P}|_{G'}$ since $m \geq 5$.

It remains to prove that $\mathcal{C}|_{G'_u}$ is discrete on $G'_{u,v}$ (the claim for $\mathcal{C}|_{G'_v}$ is symmetric). Let \mathcal{P}' be the system of stabilizers of depth $m-4$ over G'_u with respect to the action of G'_u on S . As $G'_u \subseteq \text{PGL}(V)$, we have $d(G'_u) \leq d_{\text{PGL}}(n, q) = m-4$ by Corollary 6.3. So all strongly antisymmetric \mathcal{P}' -schemes are discrete on $(G'_u)_{x'_2} = G'_{u,v}$. Also note $\mathcal{P}' \subseteq \mathcal{P}|_{G'_u}$ since $G'_u = G_{w,x,x'_1,y_1}$. It follows that all strongly antisymmetric $\mathcal{P}|_{G'_u}$ -schemes are discrete on $G'_{u,v}$. As $\mathcal{C}|_{G'_u}$ is strongly antisymmetric by Lemma 6.2, it is discrete on $G'_{u,v}$, as desired. \square

It remains to show the equivalence between $\text{GL}(V)$ and $\text{PGL}(V)$. To achieve this, we need a lemma about pointwise stabilizers of the natural action of $\text{PGL}(V)$ on $\mathbb{P}V$. Let T be a subset of $\mathbb{P}V$. For each $x \in T$, choose a representative $\tilde{x} \in V - \{0\}$. Call a subset of T *dependent* if the corresponding set of representatives are linearly dependent over \mathbb{F}_q . Clearly, this definition does not depend on the choices of the representatives. Define the relation \sim_T on T such that $x \sim_T y$ iff there exists a minimal dependent subset of T containing both x and y . It is easy to show that this is an equivalence relation.⁴ So it defines a partition of T into the equivalence classes.⁵

⁴To prove transitivity of \sim_T , consider $x, y, z \in T$ such that $x \sim_T y$ and $y \sim_T z$. Then x and y (resp. y and z) are in a dependent subset T_1 (resp. T_2) of T . Then $T_1 \cup T_2$ is a dependent set. We obtain a minimal dependent set containing x and z by removing elements in $T_1 \cup T_2 - \{x, z\}$.

⁵In the language of matroid theory, the dependent sets define a *matroid* on T , and the equivalent classes are known as the *connected components* of this matroid.

Let π denote the quotient map $\mathrm{GL}(V) \rightarrow \mathrm{PGL}(V)$. We have

Lemma 7.12. *Suppose T is a subset of $\mathbb{P}V$ and $T_1, \dots, T_k \subseteq T$ are the equivalence classes with respect to \sim_T . For $i \in [k]$, let V_i be the subspace of V spanned by (the representatives of) the elements in T_i . Then $g \in \mathrm{GL}(V)$ is sent to an element of $\mathrm{PGL}(V)_T$ under π iff g restricts to a scalar linear transformation on each V_i .*

Proof. Suppose $g \in \mathrm{GL}(V)$ restricts to a scalar linear transformation on each V_i . Then obviously $\pi(g)$ fixes each T_i pointwisely. So $\pi(g) \in \mathrm{PGL}(V)_T$. Conversely, suppose $\pi(g) \in \mathrm{PGL}(V)_T$. Then for every $x \in T$ and its representative $\tilde{x} \in V - \{0\}$, there exists a unique scalar $c_x \in \mathbb{F}_q^\times$ such that $g\tilde{x} = c_x\tilde{x}$. We need to show that for x, y in the same equivalence class T_i , it holds that $c_x = c_y$. By definition, there exists a minimal dependent subset of T containing both x and y . So we can write

$$\tilde{x} = \sum_{\tilde{v} \in I} c_v \tilde{v}, \quad c_v \in \mathbb{F}_q^\times \text{ for all } \tilde{v} \in I,$$

where I is a finite set of linearly independent vectors $\tilde{v} \in V_i$, each \tilde{v} represents an element $v \in T_i$, and $\tilde{y} \in I$. As \tilde{x} and all $\tilde{v} \in I$ are scaled by g , they are scaled by the same factor. So $c_x = c_y$, as desired. \square

In one direction, we have

Lemma 7.13. $d_{\mathrm{GL}}(n, q) \leq d_{\mathrm{PGL}}(n, q)$.

Proof. Assume $n > 1$ as otherwise $d_{\mathrm{GL}}(n, q) = d_{\mathrm{PGL}}(n, q) = 1$. Fix $m \in \mathbb{N}^+$ and let \mathcal{P} (resp. \mathcal{P}') be the system of stabilizers of depth m over $\mathrm{GL}(V)$ (resp. $\mathrm{PGL}(V)$) with respect to the natural action of $\mathrm{GL}(V)$ on $V - \{0\}$ (resp. $\mathrm{PGL}(V)$ on $\mathbb{P}V$). Fix $x \in \mathbb{P}V$ and let \tilde{x} be a representative of \tilde{x} in $V - \{0\}$. Suppose \mathcal{C} is a strongly antisymmetric \mathcal{P} -scheme that is not discrete on $\mathrm{GL}(V)_{\tilde{x}}$. We prove that there exists a strongly antisymmetric \mathcal{P}' -scheme that is not discrete on $\mathrm{PGL}(V)_x$.

Define $\mathcal{P}'' = \{\pi^{-1}(H) : H \in \mathcal{P}'\}$ which is a subgroup system over $\mathrm{GL}(V)$. We claim $\mathcal{P}'' \subseteq \mathcal{P}_{\mathrm{cl}}$ (see Definition 6.5). Consider $H = \mathrm{PGL}(V)_T \in \mathcal{P}'$, where $T \subseteq \mathbb{P}V$ satisfies $1 \leq |T| \leq m$. Let $T_1, \dots, T_k \subseteq T$ be the equivalence classes with respect to \sim_T . For $i \in [k]$, let V_i be the subspace of V spanned by (the representatives of) the elements in T_i . And let W be the subspace of V spanned by (the representatives of) those in T , i.e., $W = \sum_{i=1}^k V_i$. Let $H' := \mathrm{GL}(V)_W$. Note that $H' = \mathrm{GL}(V)_B$ for any basis B of W over \mathbb{F}_q , and $\dim_{\mathbb{F}_q} W \leq |T| = m$. So $H' \in \mathcal{P}$. We claim $H' = u_{\mathcal{P}}(\pi^{-1}(H))$ and $\pi^{-1}(H) \subseteq N_{\mathrm{GL}(V)}(H')$.

By Lemma 7.12, the group $\pi^{-1}(H)$ consists of $g \in \mathrm{GL}(V)$ that restricts to a scalar linear transformation on each V_i . So $\pi^{-1}(H)$ fixes W setwisely. Therefore we have $H' \subseteq \pi^{-1}(H) \subseteq N_{\mathrm{GL}(V)}(H')$. Suppose H'' is another subgroup in \mathcal{P} contained in $\pi^{-1}(H)$. It has the form $\mathrm{GL}(V)_{W'}$ where W' is a subspace of V . If $W \not\subseteq W'$, there exists a representative $\tilde{y} \in V - \{0\}$ of some $y \in T$ such that $\tilde{y} \notin W'$. Then there exists $g \in \mathrm{GL}(V)$ that fixes W' pointwisely but sends \tilde{y} to a vector \tilde{y}' such that \tilde{y} and \tilde{y}' are not collinear. Such an element g is in H'' but not in $\pi^{-1}(H)$, contradicting the assumption $H'' \subseteq \pi^{-1}(H)$. So $W \subseteq W'$ and hence $H'' \subseteq H'$. Therefore H' is the unique maximal subgroup in \mathcal{P} contained in $\pi^{-1}(H)$, i.e., $H' = u_{\mathcal{P}}(\pi^{-1}(H))$. By definition, we have $\pi^{-1}(H) \in \mathcal{P}_{\mathrm{cl}}$. So $\mathcal{P}'' \subseteq \mathcal{P}_{\mathrm{cl}}$.

Note that $\mathrm{GL}(V)_{\tilde{x}} = u_{\mathcal{P}}(\pi^{-1}(\mathrm{PGL}(V)_x))$. By Lemma 6.10, the existence of \mathcal{C} implies that there exists a strongly antisymmetric $\mathcal{P}_{\mathrm{cl}}$ -scheme that is not discrete on $\pi^{-1}(\mathrm{PGL}(V)_x)$. As $\mathcal{P}'' \subseteq \mathcal{P}_{\mathrm{cl}}$, there also exists a strongly antisymmetric \mathcal{P}'' -scheme that is not discrete on $\pi^{-1}(\mathrm{PGL}(V)_x)$. Finally, by Lemma 6.4, there exists a strongly antisymmetric \mathcal{P}' -scheme that is not discrete on $\mathrm{PGL}(V)_x$, as desired. \square

In the other direction, we have

Lemma 7.14. $d_{\mathrm{PGL}}(n, q) \leq d_{\mathrm{GL}}(n, q) + 2$.

Proof. Fix $m \in \mathbb{N}^+$ and let \mathcal{P} (resp. \mathcal{P}') be the system of stabilizers of depth $m + 2$ (resp. m) over $\mathrm{PGL}(V)$ (resp. $\mathrm{GL}(V)$) with respect to the natural action of $\mathrm{PGL}(V)$ on $\mathbb{P}V$ (resp. $\mathrm{GL}(V)$ on $V - \{0\}$). Suppose there exists a strongly antisymmetric \mathcal{P} -scheme \mathcal{C} that is not discrete on $\mathrm{PGL}(V)_x$ for some $x \in \mathbb{P}V$, i.e., $d_{\mathrm{PGL}}(n, q) > m + 2$. We prove that there exists a strongly antisymmetric \mathcal{P}' -scheme that is not discrete on $\mathrm{GL}(V)_y$ for some $y \in V - \{0\}$, i.e., $d_{\mathrm{GL}}(n, q) > m$.

By Lemma 6.2 and Lemma 7.6, there exists $u, v \in \mathbb{P}V$ such that the $\mathcal{P}|_{\mathrm{PGL}(V)_u}$ -collection $\mathcal{C}|_{\mathrm{PGL}(V)_u}$ is a strongly antisymmetric $\mathcal{P}|_{\mathrm{PGL}(V)_u}$ -scheme and is not discrete on $\mathrm{PGL}(V)_{u,v}$. Let \tilde{u} be a representative of u in $V - \{0\}$. The map $\pi : \mathrm{GL}(V) \rightarrow \mathrm{PGL}(V)$ restricts to a map $\pi|_{\mathrm{GL}(V)_{\tilde{u}}} : \mathrm{GL}(V)_{\tilde{u}} \rightarrow \mathrm{PGL}(V)_u$. The latter map is surjective (and in fact bijective) since every element in $\mathrm{PGL}(V)_u$ can be lifted to an element in $\mathrm{GL}(V)_{\tilde{u}}$. Define $\mathcal{P}'' := \{(\pi|_{\mathrm{GL}(V)_{\tilde{u}}})^{-1}(H) : H \in \mathcal{P}|_{\mathrm{PGL}(V)_u}\}$, which is a subgroup system over $\mathrm{GL}(V)_{\tilde{u}}$. By Lemma 6.4, there exists a strongly antisymmetric \mathcal{P}'' -scheme \mathcal{C}' that is not discrete on $(\pi|_{\mathrm{GL}(V)_{\tilde{u}}})^{-1}(\mathrm{PGL}(V)_{u,v})$.

Let $\tilde{\mathcal{P}}$ be the system of stabilizers of depth m over $\mathrm{GL}(V)_{\tilde{u}}$ with respect to the action of $\mathrm{GL}(V)_{\tilde{u}}$ on $V - \{0\}$ restricted from that of $\mathrm{GL}(V)$. We claim that

$\tilde{\mathcal{P}} \subseteq \mathcal{P}''$. Consider arbitrary $H \in \tilde{\mathcal{P}}$. It has the form $H = \text{GL}(V)_{\{\tilde{u}\} \cup T}$, where $1 \leq |T| \leq m$. Let W be the subspace of V spanned by \tilde{u} and the elements in T . Extend $\{\tilde{u}\}$ to an \mathbb{F}_q -basis $B = \{\tilde{u}, x_1, \dots, x_k\}$ of W . Then $k \leq m$. Let $w = \tilde{u} + x_1 + \dots + x_k \in V - \{0\}$. Let $B' = B \cup \{w\}$ and let \bar{B}' be the subset of $\mathbb{P}V$ consisting of the elements represented by those in B' . Then $|\bar{B}'| \leq m + 2$ and $u \in \bar{B}'$. So $\text{PGL}(V)_{\bar{B}'} \in \mathcal{P}|_{\text{PGL}(V)_u}$. As B is a basis of W , the set \bar{B}' is a minimal dependent subset and hence is the only equivalence class with respect to $\sim_{\bar{B}'}$. So $\pi^{-1}(\text{PGL}(V)_{\bar{B}'})$ consists of the elements in $\text{GL}(V)$ that restricts to scalar linear transformations on W . Then $(\pi|_{\text{GL}(V)_{\tilde{u}}})^{-1}(\text{PGL}(V)_{\bar{B}'})$ consists of the elements in $\text{GL}(V)$ that fixes W pointwisely, i.e., $(\pi|_{\text{GL}(V)_{\tilde{u}}})^{-1}(\text{PGL}(V)_{\bar{B}'}) = \text{GL}(V)_{\{\tilde{u}\} \cup T} = H$. By definition, we have $H \in \mathcal{P}''$. So $\tilde{\mathcal{P}} \subseteq \mathcal{P}''$.

Recall that \mathcal{C}' is a strongly antisymmetric \mathcal{P}'' -scheme that is not discrete on the subgroup $(\pi|_{\text{GL}(V)_{\tilde{u}}})^{-1}(\text{PGL}(V)_{u,v})$. Let \tilde{v} be a representative of v in $V - \{0\}$. Then $\text{GL}(V)_{\tilde{u}, \tilde{v}} = (\text{GL}(V)_{\tilde{u}})_{\tilde{v}} \in \tilde{\mathcal{P}} \subseteq \mathcal{P}''$. Note $\text{GL}(V)_{\tilde{u}, \tilde{v}} \subseteq (\pi|_{\text{GL}(V)_{\tilde{u}}})^{-1}(\text{PGL}(V)_{u,v})$. By Lemma 2.3, we know \mathcal{C}' is not discrete on $\text{GL}(V)_{\tilde{u}, \tilde{v}}$. As $\tilde{\mathcal{P}} \subseteq \mathcal{P}''$, there exists a strongly antisymmetric $\tilde{\mathcal{P}}$ -scheme that is not discrete on $\text{GL}(V)_{\tilde{u}, \tilde{v}} = (\text{GL}(V)_{\tilde{u}})_{\tilde{v}}$. By Corollary 6.2, there exists a strongly antisymmetric \mathcal{P}' -scheme that is not discrete on $\text{GL}(V)_{\tilde{v}}$, as desired. \square

Theorem 7.4 now follows from Lemma 7.9, Lemma 7.10, Lemma 7.11, Lemma 7.13, and Lemma 7.14.

Upper bounds for $d_{\text{GL}}(n, q)$. It is easy to see that we have two upper bounds for $d_{\text{GL}}(n, q)$:

1. $d_{\text{GL}}(n, q) \leq \left(\frac{2}{\log 12}\right) \log(q^n - 1) + O(1) = \left(\frac{2 \log q}{\log 12}\right) n + O(1)$. This follows from Corollary 7.1.
2. $d_{\text{GL}}(n, q) \leq n$. This follows from Lemma 2.5 and the fact that the natural action of $\text{GL}_n(q)$ has a base of size n .

The first bound is asymptotically better if $q \in \{2, 3\}$. Otherwise the second one is better. Now we prove another upper bound that slightly improves both of the two bounds above.

Theorem 7.5. $d_{\text{GL}}(n, q) \leq \left(\frac{\log q}{\log q + (\log 12)/4}\right) n + O(1)$.

Proof. Let $G = \mathrm{GL}_n(q)$ and $S = \mathbb{F}_q^n - \{0\}$. Fix a positive integer $m \leq n$. Let \mathcal{P} be the system of stabilizers of depth m with respect to the natural action of G on S . Suppose there exists a strongly antisymmetric \mathcal{P} -scheme $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ that is not discrete on G_x for some $x \in S$. We prove that $m \leq \left(\frac{\log q}{\log q + (\log 12)/4}\right)n + O(1)$.

By Lemma 7.3, there exists a subset $T = \{x_1, \dots, x_m\} \subseteq S$ of cardinality m such that C_{G_T} has a block B of cardinality at least $2^{\left(\frac{\log 12}{4}\right)m^2 - O(m)}$. We claim that the elements x_i in T may be assumed to be linearly independent: if they are not, replace T by a set T' of cardinality m such that (1) the elements in T' are linearly independent, and (2) the subspace spanned by T' contains the one spanned by T . Then replace B with a block B' of $C_{G_{T'}}$ such that $\pi_{G_{T'}, G_T}(B') = B$. We have $|B'| \geq |B|$. This proves the claim.

Note that $N_G(G_T)$ is the setwise stabilizer of subspace spanned by T . Therefore $N_G(G_T)/G_T \cong \mathrm{GL}_m(q)$. By antisymmetry, the group $N_G(G_T)/G_T$ acts semiregularly on the set of blocks of C_{G_T} . So we have

$$|G_T \backslash G| \geq |N_G(G_T)/G_T| \cdot |B| \geq 2^{\left(\frac{\log 12}{4}\right)m^2 - O(m)} \cdot \prod_{i=0}^{m-1} (q^m - q^i).$$

On the other hand, note that G_T is the stabilizer of $u := (x_1, \dots, x_m) \in S^{(m)}$ under the diagonal action of G on $S^{(m)}$. By the orbit-stabilizer theorem, we have $|G_T \backslash G| = |Gu|$, which is the number of m -tuples of linearly independent vectors in $V - \{0\}$. Therefore $|G_T \backslash G| = \prod_{i=0}^{m-1} (q^n - q^i)$. So we have

$$2^{\left(\frac{\log 12}{4}\right)m^2 - O(m)} \cdot \prod_{i=0}^{m-1} (q^m - q^i) \leq \prod_{i=0}^{m-1} (q^n - q^i).$$

Solving the inequality yields the desired bound. □

As $q \geq 2$, we have $\log q + (\log 12)/4 < (\log 12)/2$. So Theorem 7.5 is indeed an improvement of the bound $d_{\mathrm{GL}}(n, q) \leq \left(\frac{2 \log q}{\log 12}\right)n + O(1)$ above.

Chapter 8

GROUPS WITH RESTRICTED NONCYCLIC COMPOSITION FACTORS

In this chapter, we consider the problem of factoring a polynomial $f(X) \in \mathbb{F}_q[X]$ using a lifted polynomial \tilde{f} where the Galois group of \tilde{f} has *restricted noncyclic composition factors*.

Simple groups, composition factors, and CFSG. To formally state our result, we first review some definitions and facts in group theory. A *simple group* is a nontrivial group whose only normal subgroups are the trivial group and the group itself. A *composition series* of a group G is a finite chain of subgroups

$$\{e\} = H_0 \subseteq H_1 \subseteq \cdots \subseteq H_k = G$$

such that for every $i \in [k]$, H_{i-1} is a maximal normal subgroup of H_i , so that H_i/H_{i-1} is simple. Such a series always exists when G is finite. The groups H_i/H_{i-1} are called the *composition factors* of G . It is a consequence of the *Jordan-Hölder theorem* that the set of the composition factors of G does not depend on the choice of composition series (see, e.g., [Lan02]).

Now suppose G is a finite group. The composition factors of G are finite simple groups, which are classified by the *classification of finite simple groups* (CFSG):

Theorem 8.1 (classification of finite simple groups). *A finite simple group is isomorphic to one of the following groups: a cyclic group of prime order, an alternating group $\text{Alt}(n)$ ($n \geq 5$), a classical group, an exceptional group of Lie type, or one of the 26 sporadic simple groups.*

See, e.g., [GLS94]. We do not describe these families of finite simple groups, except mentioning that a finite simple group is a classical group if it has one of the following forms (see, e.g., [KL90]):

$$\text{PSL}_n(q), \text{PSU}_n(q), \text{PSp}_n(q) \ (n \text{ even}), \text{P}\Omega_n^\pm(q) \ (n \text{ even}), \ \Omega_n(q) \ (n \text{ odd}).$$

We denote by $k(G)$ the maximum degree of the alternating groups that appear as noncyclic composition factors of G , and let $k(G) = 1$ if such alternating groups do

not exist. Similarly, denote by $r(G)$ the maximum order of the classical groups that appear as noncyclic composition factors of G , and let $r(G) = 1$ if such classical groups do not exist.

Main result. Let \mathbb{F}_q , A_0 and K_0 be as in Chapter 5. The main result of this chapter is a GRH-based deterministic algorithm that factorizes $f(X) \in \mathbb{F}_q[X]$ using a lifted polynomial $\tilde{f}(X) \in A_0[X]$, such that the running time of the algorithm is controlled by $k(G)$ and $r(G)$, where $G = \text{Gal}(\tilde{f}/K_0)$ is the Galois group of \tilde{f} over K_0 .

Theorem 8.2. *Under GRH, there exists a deterministic algorithm that, given a polynomial $f(X) \in \mathbb{F}_q[X]$ of degree $n \in \mathbb{N}^+$ and a lifted polynomial $\tilde{f}(X) \in A_0[X]$ of f with the Galois group $G := \text{Gal}(\tilde{f}/K_0)$ over K_0 , computes the complete factorization of f over \mathbb{F}_q in time polynomial in n , $\log q$, $k(G)^{\log k(G)}$ and $r(G)$.*

For $k \in \mathbb{N}^+$, denote by Γ_k the family of finite groups whose noncyclic composition factors are all isomorphic to subgroups of $\text{Sym}(k)$. It is known that a classical group H is isomorphic to a subgroup of $\text{Sym}(k)$ only if $|H| = k^{O(\log k)}$ [Coo78]. Therefore we have

Theorem 8.3. *Under GRH, there exists a deterministic algorithm that, given a polynomial $f(X) \in \mathbb{F}_q[X]$ of degree $n \in \mathbb{N}^+$ and a lifted polynomial $\tilde{f}(X) \in A_0[X]$ of f , computes the complete factorization of f over \mathbb{F}_q in time polynomial in n , $\log q$ and $k^{\log k}$, where k is the smallest positive integer satisfying $\text{Gal}(\tilde{f}/K_0) \in \Gamma_k$. In particular, the algorithm runs in polynomial time if $k = 2^{O(\sqrt{\log n})}$.*

Table 8.1: Known deterministic polynomial-time factoring algorithms for Γ_k

k	Reference
4	[Evd92]
$O(1)$	[Evd92] + [BCP82]
$2^{O(\sqrt{\log n})}$	Our result
n	Goal

By Theorem 8.3, we have a deterministic polynomial-time algorithm that given $\tilde{f}(X)$, completely factorizes $f(X)$ under GRH, provided that $\text{Gal}(\tilde{f}/K_0) \in \Gamma_k$ for some $k = 2^{O(\sqrt{\log n})}$ (note that achieving $k = n$ would fully resolve the problem of deterministic polynomial factoring under GRH). Previously, such an algorithm was

known only for bounded k : for $k \leq 4$ this follows directly from the deterministic polynomial-time factoring algorithm for solvable Galois groups [Evd92] (see Theorem 4.3 and Theorem 5.13). For $k = O(1)$, it follows from the proof in [Evd92] together with the bound in [BCP82] for the orders of primitive permutation groups. See Table 8.1 for a summary.

Overview of the proof. We prove Theorem 8.2 using the generalized \mathcal{P} -scheme algorithm in Chapter 5. If the input polynomial f is assumed to satisfy Condition 3.1, we may also use the simpler algorithm in Chapter 3. These algorithms reduce the problem of factoring f to the one of constructing a collection of (relative) number fields such that the associated subgroup system \mathcal{P} has the property that all strongly antisymmetric \mathcal{P} -schemes are discrete on a certain subgroup (see Theorem 3.9 and Theorem 5.9).

We further reduce the latter problem to the case that the Galois group $\text{Gal}(\tilde{f}/K_0)$ is a primitive permutation group on the set of roots of \tilde{f} , using Theorem 4.2 and some facts from group theory. Next we consider the following special kind of subgroup systems.

Definition 8.1. *Let G be a finite permutation group on a finite set S . For $N \in \mathbb{N}^+$, define the subgroup system $\mathcal{P}_{G,N}$ over G by*

$$\mathcal{P}_{G,N} := \left\{ G_{U \cup U'} : \begin{array}{l} \emptyset \neq U \subseteq S, x \in S, U' \subseteq G_U x \\ |S|^{|U|}, |G_U x|^{|U'|} \leq N \end{array} \right\}.$$

We prove a sufficient condition for a subgroup system $\mathcal{P}_{G,N}$ over a primitive permutation group to have the desired property:

Theorem 8.4. *Let G be a primitive permutation group on a finite set S . For sufficiently large $N = \text{poly}(k(G)^{d_{\text{sym}}(k(G))}, r(G), |S|) \geq |S|$, all strongly antisymmetric $\mathcal{P}_{G,N}$ -schemes are discrete on $G_x \in \mathcal{P}_{G,N}$ for all $x \in S$.*

It is easy to see that the complexity $c(\mathcal{P}_{G,N})$ of $\mathcal{P}_{G,N}$ is polynomial in N . We modify the algorithm in Lemma 4.10 to construct a collection of (relative) number fields in time polynomial in n , $\log q$ and $c(\mathcal{P}_{G,N})$ such that the associated subgroup system is precisely $\mathcal{P}_{G,N}$. Theorem 8.2 then follows from Theorem 8.4.

Finally, to prove Theorem 8.4, we apply the *O’Nan-Scott theorem* [LPS88] in permutation group theory, which states that a finite primitive permutation group is in

exactly one of the following five categories: almost simple type, affine type, diagonal type, product type, and twisted wreath type. We prove Theorem 8.4 by verifying it in these five cases separately.

Outline of the chapter. In Section 8.1, we derive Theorem 8.2 from Theorem 8.4 using an algorithm that constructs the collection of (relative) number fields corresponding to $\mathcal{P}_{G,N}$. The rest of the chapter focuses on the proof of Theorem 8.4: Section 8.2 describes the O’Nan-Scott theorem [LPS88] and the five categories of primitive permutation groups. In Sections 8.3–8.6, we prove Theorem 8.4 for primitive permutation groups of almost simple type, affine type, diagonal type and product type respectively. We also address twisted wreath type at the end of Section 8.6 by reducing to the case of product type using an argument in [Pra90]. Finally, we discuss possible directions for future research in Section 8.7.

8.1 Proof of the main theorem

We start by describing an algorithm `SubgroupSystem` that computes a (K_0, g) -subfield system \mathcal{F} given a number field K_0 , an integer $N \in \mathbb{N}^+$, and a polynomial $g(X) \in K_0[X]$ irreducible over K_0 , such that the subgroup system associated with \mathcal{F} is exactly $\mathcal{P}_{G,N}$, where $G = \text{Gal}(g/K_0)$.

The pseudocode is given in Algorithm 16. First compute the greatest integer $d \in \{0, \dots, \deg(g)\}$ subject to $\deg(g)^d \leq N$. Run the algorithm `Stabilizers` in Lemma 4.10 on the input (K_0, d, g) to obtain a (K_0, g) -subfield system \mathcal{F}' , and let $\mathcal{F} = \mathcal{F}'$. Next enumerate $K \in \mathcal{F}'$ and the irreducible factors g_0 of g over K . For each (K, g_0) , let d' be the greatest integer in $\{0, \dots, \deg(g_0)\}$ subject to $\deg(g_0)^{d'} \leq N$, run the algorithm `Stabilizers` on the input (K, d', g_0) to obtain a (K, g_0) -subfield system \mathcal{F}'' , and add the fields in \mathcal{F}'' to \mathcal{F} .¹

The following lemma states that the subgroup system associated with \mathcal{F} is precisely $\mathcal{P}_{G,N}$.

Lemma 8.1. *Given a number field K_0 , an integer $N \in \mathbb{N}^+$, and a polynomial $g(X) \in K_0[X]$ irreducible over K_0 , the algorithm `SubgroupSystem` computes a (K_0, g) -subfield system \mathcal{F} , such that the subgroup system associated with \mathcal{F} is precisely $\mathcal{P}_{G,N}$ over $G := \text{Gal}(g/K_0)$, where G is regarded as a permutation group*

¹We add a field to \mathcal{F} only if it is non-isomorphic to all fields in \mathcal{F} over K_0 , so that the fields in \mathcal{F} are always mutually non-isomorphic over K_0 .

Algorithm 16 SubgroupSystem

Input: number field K_0 , $N \in \mathbb{N}^+$, and $g(X) \in K_0[X]$ irreducible over K_0

Output: (K_0, g) -subfield system \mathcal{F}

- 1: $d \leftarrow \max\{i \in \mathbb{N} : 0 \leq i \leq \deg(g), \deg(g)^i \leq N\}$
 - 2: run Stabilizers on (K_0, d, g) to compute a (K_0, g) -subfield system \mathcal{F}'
 - 3: $\mathcal{F} \leftarrow \mathcal{F}'$
 - 4: **for** $K \in \mathcal{F}'$ **do**
 - 5: factorize g over K
 - 6: **for** irreducible factor g_0 of g over K **do**
 - 7: $d' \leftarrow \max\{i \in \mathbb{N} : 0 \leq i \leq \deg(g_0), \deg(g_0)^i \leq N\}$
 - 8: run Stabilizers on (K, d', g_0) to compute a (K, g_0) -subfield system \mathcal{F}''
 - 9: **for** $K' \in \mathcal{F}''$ **do**
 - 10: compute a relative number field \tilde{K}' over K_0 such that $\tilde{K}' \cong_{K_0} K'$
 - 11: **if** \tilde{K}' is non-isomorphic to all fields in \mathcal{F} over K_0 **then**
 - 12: $\mathcal{F} \leftarrow \mathcal{F} \cup \{\tilde{K}'\}$
 - 13: **return** \mathcal{F}
-

on the set of roots of g in the splitting field of g over K_0 . Moreover, the algorithm runs in time polynomial in $c(\mathcal{P}_{G,N}) = N^{O(1)}$ and the size of the input.

Proof. Let S be the set of roots of g in the splitting field of g over K_0 . Let $d = \max\{i \in \mathbb{N} : 0 \leq i \leq \deg(g), \deg(g)^i \leq N\}$. By definition, the subgroup system $\mathcal{P}_{G,N}$ consists of the pointwise stabilizers $G_{U \cup U'}$, such that U is a nonempty subset of S of cardinality at most d , and U' is a subset of a G_U -orbit $O \subseteq S$ satisfying $|O|^{|U'|} \leq N$.

Note that when we fix $U' = \emptyset$, the groups $G_{U \cup U'} = G_U$ are precisely those in the system of stabilizers of depth d with respect to the action of G on S . We construct the corresponding fields by running the algorithm Stabilizers on (K_0, d, g) .

Next consider the groups $G_{U \cup U'}$ where $U' \neq \emptyset$. We enumerate $K = L^{G_U}$ and the irreducible factor g_0 of g over K . By Galois theory, the set of roots of g_0 is a G_U -orbit $O \subseteq S$. Let $d' = \max\{i \in \mathbb{N} : 0 \leq i \leq \deg(g_0), \deg(g_0)^i \leq N\}$. We run the algorithm Stabilizers on (K, d', g_0) to construct the fields corresponding to the subgroups $(G_U)_{U'} = G_{U \cup U'}$, where $U' \subseteq O$ and $1 \leq |U'| \leq d'$. Moreover, all the groups G_U and the G_U -orbits in S are enumerated. It follows that the subgroup

system associated with \mathcal{F} is precisely $\mathcal{P}_{G,N}$.

Finally, the fact $c(\mathcal{P}_{G,N}) = N^{O(1)}$ and the claim about the running time follow from Lemma 4.2 and Lemma 4.10. \square

We also need the following lemma, which states that restricting to a subgroup does not increase the quantities $k(G)^{\log k(G)}$ and $r(G)$ by much.

Lemma 8.2. *Let G be a permutation group on a finite set S , and let G' be a subquotient of G . Then $k(G')^{\log k(G')}$ and $r(G')$ are polynomial in $k(G)^{\log k(G)}$, $r(G)$ and $|S|$.*

Proof. Let H' be a noncyclic composition factor G' . Then H' is isomorphic to a subquotient of a noncyclic composition factor H of G , i.e., there exists a subgroup H'' of H and a normal subgroup N of H'' such that $H' \cong H''/N$. Fix such H , H'' and N . We want to prove (1) if H' is an alternating group $\text{Alt}(k')$, then $k'^{\log k'}$ is polynomial in $k(G)^{\log k(G)}$, $r(G)$ and $|S|$, and (2) if H' is a classical group, then $|H'|$ is polynomial in $k(G)^{\log k(G)}$, $r(G)$ and $|S|$.

By CFSG, the group H is either an alternating group or a group of Lie type (i.e. a classical group or an exceptional group of Lie type). First assume H is an alternating group of degree $k \leq k(G)$. If H' is also an alternating group, its degree k' is obviously bounded by k . So $k'^{\log k'} \leq k(G)^{\log k(G)}$. Now consider the case that H' is a classical group of the form $\text{PSL}_n(q)$, $\text{PSU}_n(q)$, $\text{PSp}_n(q)$, $\text{P}\Omega_n^\pm(q)$, or $\Omega_n(q)$ over a finite field \mathbb{F}_q for some $n \in \mathbb{N}^+$. We have $|H'| = q^{\Theta(n^2)}$. Denote by $\mu(T)$ the minimal degree of a faithful permutation representation of a finite group T . It was proven in [KP00] that if \bar{T} is a quotient group of T with no nontrivial abelian normal subgroup, then $\mu(\bar{T}) \leq \mu(T)$. As $H' \cong H''/N$ is simple and noncyclic, we have

$$\mu(H') = \mu(H''/N) \leq \mu(H'') \leq \mu(H) \leq k.$$

On the other hand, it was shown in [Coo78] that $\mu(H') = q^{\Theta(n)}$ (see also [KL90, Table 5.2.A]). So we have $n = O(\log k / \log q)$ and

$$|H'| = q^{\Theta(n^2)} = k^{O(\log k / \log q)} = k(G)^{O(\log k(G))}.$$

Next assume H is a group of Lie type over a finite field \mathbb{F}_q , and has Lie rank ℓ .² Then $|H| = q^{\Theta(\ell^2)}$ [KL90, Table 5.1.A]. It is also known that H has a faithful projective

²Each finite simple group of Lie type has an associated *Lie rank*. See, e.g., [KL90, Section 5.1].

linear representation $H \hookrightarrow \mathrm{PGL}_d(\overline{\mathbb{F}}_q)$ of degree $d = O(\ell)$, where $\overline{\mathbb{F}}_q$ is the algebraic closure of \mathbb{F}_q (see [KL90, Proposition 5.4.13]). As H is finite, this also holds for some finite field F in place of $\overline{\mathbb{F}}_q$. Identify $H'' \subseteq H$ with a subgroup of $\mathrm{PGL}_d(F)$. Then $H' \cong H''/N$ is a subquotient of $\mathrm{PGL}_d(F)$, and hence also a subquotient of $\mathrm{GL}_d(F)$. Choose the largest $s \in \mathbb{N}^+$ such that H' has a subquotient isomorphic to $\mathrm{Alt}(s)$. Then $\mathrm{Alt}(s)$ is isomorphic to a subquotient of $\mathrm{GL}_d(F)$. On the other hand, it is known that $\mathrm{Alt}(s)$ has a finite preimage in $\mathrm{GL}_d(F)$ only if $s = O(d)$ [DM96, Theorem 5.7A]. So we have $s = O(d) = O(\ell)$.

Suppose H is a classical group. If $H' = \mathrm{Alt}(s)$, we have $s^{\log s} = \ell^{O(\log \ell)} = |H|^{O(1)} = r(G)^{O(1)}$. And if H' is a classical group, we have the obvious bound $|H'| \leq |H| \leq r(G)$.

Finally, suppose H is an exceptional group of Lie type. Then $s = O(\ell) = O(1)$. In the case $H' = \mathrm{Alt}(s)$, we have $s^{\log s} = O(1)$. So assume H' is a classical group of the form $\mathrm{PSL}_n(q')$, $\mathrm{PSU}_n(q')$, $\mathrm{PSp}_n(q')$, $\mathrm{P}\Omega_n^\pm(q')$ or $\Omega_n(q')$ over a finite field \mathbb{F}_q for some $n \in \mathbb{N}^+$. It is easy to see that H' has a subquotient isomorphic to an alternating group of degree $\Omega(n)$ (see, e.g., [LS03, Proposition 16.4.4]). So $s = \Omega(n)$, which implies $n = O(1)$. Then $\mu(H') = q'^{\Theta(n)} = q'^{\Theta(n^2)} = |H'|^{\Theta(1)}$. On the other hand, we see above that $\mu(H') \leq \mu(H)$ since H' is a subquotient of H and is a noncyclic simple group. For the same reason, we have $\mu(H) \leq \mu(G) \leq |S|$. It follows that $|H'| = |S|^{O(1)}$. \square

Now we are ready to prove Theorem 8.2 under the assumption of Theorem 8.4.

Proof of Theorem 8.2. The first step is to reduce to the case that \tilde{f} is irreducible over K_0 , as in Chapter 5: Let $p = \mathrm{char}(\mathbb{F}_q)$. Using Lemma 5.1, we compute an integer D satisfying $D \equiv 1 \pmod{p}$ and a factorization of $D \cdot \tilde{f}$ into irreducible factors $\tilde{f}_1, \dots, \tilde{f}_k \in A_0[X]$ over K_0 . Then we have $f(X) = \prod_{i=1}^k \tilde{\psi}_0(f_i)(X)$. The Galois groups $\mathrm{Gal}(\tilde{f}_i(X)/K_0)$ are quotient groups of $G = \mathrm{Gal}(\tilde{f}/K_0)$. So the set of the composition factors of each \tilde{f}_i is a subset of that of G . By replacing $\tilde{f}(X)$ with $\tilde{f}_i(X)$ and $f(X)$ with $\tilde{\psi}_0(f_i) \in \mathbb{F}_q[X]$ for each $i \in [k]$, we reduce to the case that \tilde{f} is irreducible over K_0 .

Choose sufficiently large $N = \mathrm{poly}(k(G)^{\log k(G)}, r(G), \deg(f)) \geq \deg(f)$. Assume for a moment that the value of N is known to the algorithm. First consider the case that G acts primitively on the set of roots of \tilde{f} . We compute a (K_0, \tilde{f}) -subfield system \mathcal{F} using the algorithm `SubgroupSystem` above. By Lemma 8.1,

the associated subgroup system \mathcal{P} over G equals $\mathcal{P}_{G,N}$. Then by Theorem 8.4 and the fact $d_{\text{Sym}}(k(G)) = O(\log k(G))$ (see Corollary 7.1), all strongly antisymmetric \mathcal{P} -schemes are discrete on $G_x \in \mathcal{P}$ for all roots x of \tilde{f} .

Now consider the general case, where the action of G may be imprimitive. We run the algorithm `GeneralAction` in Theorem 4.2 to compute \mathcal{F} , as well as a tower of relative number fields $K_0 \subseteq K_1 \subseteq \cdots \subseteq K_{k-1} \subseteq K_k$ over K_0 and $g_i(X) \in K_{i-1}[X]$ for $i \in [k]$, such that

1. K_i is isomorphic to $K_{i-1}[X]/(g_i(X))$ over K_{i-1} , and
2. the Galois group $G_i := \text{Gal}(L_i/K_{i-1})$ acts primitively on the set of roots of g_i in L_i , where L_i is the Galois closure of K_i/K_{i-1} .

We implement the algorithm `PrimitiveAction` required in Theorem 4.2 using the algorithm `SubgroupSystem`. The latter has an extra parameter N , which is chosen as above. For $i \in [k]$, let \mathcal{F}_i be the (K_{i-1}, g_i) -subfield system computed by `SubgroupSystem` on the input (K_{i-1}, N, g_i) , and let \mathcal{P}_i be the associated subgroup system over G_i . Note that the groups G_i are subquotients of G . Then by Theorem 8.4, Lemma 8.1, and Lemma 8.2, for all $i \in [k]$, all strongly antisymmetric \mathcal{P}_i -schemes are discrete on $(G_i)_x$ for all roots x of g_i , provided that $N = \text{poly}(k(G)^{\log k(G)}, r(G), \deg(f))$ is sufficiently large. In this case, by Theorem 4.2, all strongly antisymmetric \mathcal{P} -schemes are discrete on $G_x \in \mathcal{P}$ for all roots x of \tilde{f} , where \mathcal{P} is the subgroup system associated with \mathcal{F} .

Finally, we run the generalized \mathcal{P} -scheme algorithm in Chapter 5 using the (K_0, \tilde{f}) -subfield system \mathcal{F} computed above, so that \tilde{f} is completely factorized by Theorem 3.9. If f satisfies Condition 3.1, we may also use the simpler algorithm in Chapter 3 and apply Theorem 3.9 instead.

The above algorithm assumes that the value of a sufficiently large integer N is known. We may avoid this assumption by running the algorithm multiple times, where N is initially a constant and is doubled each time, until f is completely factorized. It only causes an extra factor of $O(\log N)$ in the running time. \square

8.2 The O’Nan-Scott theorem for finite primitive permutation groups

The O’Nan-Scott theorem for finite primitive permutation groups [LPS88] is one of the most influential theorems in permutation group theory. In this section, we describe this theorem and the related definitions.

We start with the notion of the *socle* of a finite group:

Definition 8.2 (socle). *The socle of a finite group G , denoted by $\text{soc}(G)$, is the subgroup generated by the minimal (nontrivial) normal subgroups of G .*

Next we define the five categories of finite primitive permutation groups appeared in the O’Nan-Scott theorem.

Almost simple type. Let T be a noncyclic finite simple group, so that its center $Z(T)$ is trivial. We identify T with the inner automorphism group $\text{Inn}(T) \subseteq \text{Aut}(T)$ via the isomorphism sending $g \in T$ to the conjugation $h \mapsto ghg^{-1}$ (this map is indeed an isomorphism since its kernel equals $Z(T)$ and hence is trivial).

We say a finite group is *almost simple* if it is isomorphic to a group G satisfying $T \subseteq G \subseteq \text{Aut}(T)$ for some noncyclic finite simple group T . It is known that in this case $T = \text{soc}(G)$ holds.

A finite permutation group of *almost simple type* is simply a finite primitive permutation group that is also almost simple as an abstract group:

Definition 8.3 (almost simple type). *A finite permutation group is said to be of almost simple type if it is primitive and almost simple.*

Affine type. Finite permutation groups of *affine type* are primitive groups arising as subgroups of general affine groups that contain all the translations:

Definition 8.4 (affine type). *A finite permutation group is said to be of affine type if it is primitive and permutation isomorphic to a subgroup G of a general affine group $\text{AGL}(V)$ acting naturally on a finite-dimensional vector space V over a prime field \mathbb{F}_p , and G contains the subgroup of translations $V^\# := \{x \mapsto x + u : u \in V\} \subseteq \text{AGL}(V)$.*

For example, Lemma 6.19 states that finite primitive solvable permutation groups are of affine type.

Diagonal type. Let T be a noncyclic finite simple group and let $k \geq 2$ be an integer. Consider the subgroup A of $\text{Aut}(T)^k$, defined by

$$A := \{(a_1, \dots, a_k) \in \text{Aut}(T)^k : a_i \text{Inn}(T) = a_j \text{Inn}(T) \text{ for all } i, j \in [k]\}.$$

The group $\text{Sym}(k)$ acts on A by permuting the k coordinates, sending $(a_1, \dots, a_k) \in A$ to $(a_{\pi^{-1}(1)}, \dots, a_{\pi^{-1}(k)})$. So we can form the semidirect product

$$W := A \rtimes \text{Sym}(k).$$

Also define the subgroups $M, D \subseteq W$ by

$$M := \text{Inn}(T)^k \subseteq A \subseteq W$$

and

$$D := \{(a, \dots, a)\pi : a \in \text{Aut}(T), \pi \in \text{Sym}(k)\} \subseteq W.$$

Then W acts on the right coset space $D \backslash W$ by inverse right translation. Permutation groups of diagonal type arise as subgroups of W :

Definition 8.5 (diagonal type). *A finite permutation group is said to be of diagonal type if it is primitive and is permutation isomorphic to a group G satisfying $M \subseteq G \subseteq W$ acting on $D \backslash W$ by inverse right translation, where D, M, W are as above.*

Example 8.1 (holomorph of a noncyclic finite simple group). Let T be a noncyclic finite simple group. We may form the semidirect product $\text{Hol}(T) := T \rtimes \text{Aut}(T)$ with respect to the natural action of $\text{Aut}(T)$ on T . The group $\text{Hol}(T)$ is called the *holomorph* of T . By identifying T (as a set) with the left coset space $\text{Hol}(T)/\text{Aut}(T)$ via the bijection $T \rightarrow \text{Hol}(T)/\text{Aut}(T)$ sending $g \in T$ to $g\text{Aut}(T)$, we see that the action of $\text{Hol}(T)$ on $\text{Hol}(T)/\text{Aut}(T)$ by left translation is equivalent to its action on the set T defined by ${}^{hg}h' = h^g h'$ for $h, h' \in T, g \in \text{Aut}(T)$. By the following lemma, this is an example of finite primitive permutation groups of diagonal type.

Lemma 8.3. $\text{Hol}(T)$ is a finite primitive permutation group of diagonal type on T .³

Proof. The action of $\text{Hol}(T)$ on T is obviously transitive. It is faithful since $\text{Hol}(T)_e = \text{Aut}(T)$ acts faithfully on T . To prove that $\text{Hol}(T)$ is primitive, we want to show that $\text{Aut}(T)$ is maximal in $\text{Hol}(T)$. Consider any group G satisfying $\text{Aut}(T) \subseteq G \subseteq \text{Hol}(T)$. The kernel of G under the quotient map $\text{Hol}(T) \rightarrow \text{Aut}(T)$ is a normal subgroup of T , and hence is either $\{e\}$ or T . So G equals either $\text{Aut}(T)$ or $\text{Hol}(T)$. Therefore $\text{Hol}(T)$ acts primitively on T .

³Lemma 8.3 holds more generally for any group G satisfying $T \rtimes \text{Inn}(T) \subseteq G \subseteq \text{Hol}(T)$. In an alternative formulation of the O’Nan-Scott theorem, such a group G is said to have type HS (holomorph of a simple group). See, e.g., [PLN97]. We do not use this notation in the thesis.

Now define the groups D, M, W, A as above with respect to T and $k = 2$. For $g \in T$, denote by $\tau_g \in \text{Inn}(T)$ the conjugation by g which sends $x \in T$ to gxg^{-1} . Define the map $\rho : \text{Hol}(T) \rightarrow A$ via $\rho(gh) = (\tau_g h, h)$ for $g \in T, h \in \text{Aut}(T)$. It is straightforward to check that ρ is a well defined injective group homomorphism and $M \subseteq \rho(\text{Hol}(T))$. The action of A on $D \setminus W$ by inverse right translation thus induces an action of $\text{Hol}(T)$ on $D \setminus W$, which is transitive since $M \subseteq \rho(\text{Hol}(T))$. The stabilizer of $De \in D \setminus W$ with respect to this action is $\rho^{-1}(D \cap A) = \text{Aut}(T) \subseteq \text{Hol}(T)$, which is exactly the stabilizer of $e \in T$ with respect to the action of $\text{Hol}(T)$ on T . By Lemma 2.1, the action of $\text{Hol}(T)$ on T and that on $D \setminus W$ are equivalent. The lemma follows by definition. \square

Product type. Let H be a primitive permutation group on a finite set Γ of almost simple type or diagonal type. Let $k \geq 2$ be an integer. Define the wreath product

$$W := H \wr \text{Sym}(k) = H^k \rtimes \text{Sym}(k),$$

where $\text{Sym}(k)$ permutes the k factors of H^k . The group W has a natural *primitive wreath product action* on Γ^k where H^k acts coordinatewise and $\text{Sym}(k)$ permutes the coordinates. Also define

$$M := \text{soc}(H)^k \subseteq W.$$

Permutation groups of product type arise as subgroups of W :

Definition 8.6 (product type). *A finite permutation group is said to be of product type if it is primitive and is permutation isomorphic to a group G satisfying $M \subseteq G \subseteq W$ acting on Γ^k via the primitive wreath product action, where M, W, Γ, k are as above.*

Twisted wreath type. Let T be a noncyclic finite simple group. Let P be a transitive permutation group on $[k]$ where $k \geq 2$. Denote by $\text{Map}(P, T)$ the set of the maps from P to T . Suppose $\varphi : P_1 \rightarrow \text{Aut}(T)$ is a group homomorphism from the stabilizer P_1 of $1 \in [k]$ to $\text{Aut}(T)$. Define

$$B := \{f \in \text{Map}(P, T) : f(pq^{-1}) = \varphi(q)(f(p)) \text{ for all } p \in P, q \in P_1\},$$

which is a group under coordinatewise multiplication. The group P acts on B via $({}^p f)(px) = f(x)$, or equivalently

$$({}^p f)(x) = f(p^{-1}x) \quad \text{for all } p, x \in P, f \in B.$$

It is easy to check that this is a well defined action.⁴ So we can form the semidirect product $G := B \rtimes P$ with respect to this action. The group G is also called the *twisted wreath product* with respect to the data (T, P, φ) , denoted by $T \text{ twr}_\varphi P$ [Neu63; DM96].

Finite permutation groups of twisted wreath type are defined as follows.

Definition 8.7 (twisted wreath type). *A finite permutation group is said to be of twisted wreath type if it is primitive and is permutation isomorphic to a group $G = T \text{ twr}_\varphi P$ acting on the left coset space G/P via left translation, where T , P , and φ are as above.*

The O’Nan-Scott theorem. Now we are ready to state the O’Nan-Scott theorem for finite primitive permutation groups [LPS88].

Theorem 8.5 (O’Nan-Scott theorem). *A finite primitive permutation group is of exactly one of the following types: almost simple type, affine type, diagonal type, product type, and twisted wreath type.*

Schreier conjecture. We conclude this section by mentioning the fact that the outer automorphism group of every finite simple group is solvable. This is known as the *Schreier conjecture*, and is now known to be true as a result of CFSG. See, e.g., [DM96].

Theorem 8.6. *The outer automorphism group $\text{Out}(T)$ of every finite simple group T is solvable.*

8.3 Almost simple type

In this section, we prove Theorem 8.4 for finite primitive permutation groups of almost simple type. Our proof is based on the work on the minimal base sizes of such permutation groups, including the work on Pyber’s base size conjecture, and the constant bounds for non-standard actions.

Pyber’s base size conjecture. Recall that a base of a permutation group G on a finite set S is a subset $T \subseteq S$ satisfying $G_T = \{e\}$, and the minimal base size $b(G)$ is the minimum cardinality of a base of G . By the orbit-stabilizer theorem, we

⁴For example, the map ${}^p f$ is indeed in B for $p \in P$ and $f \in B$ since $({}^p f)(p'q^{-1}) = f(p^{-1}p'q^{-1}) = \varphi^{(q)}(f(p^{-1}p')) = \varphi^{(q)}(({}^p f)(p'))$ for all $p' \in P$ and $q \in P_1$.

have the lower bound $b(G) \geq \log |G| / \log |S|$. *Pyber's base size conjecture* [Pyb93] asserts that this is asymptotically tight if G is primitive:

Conjecture 8.1 (Pyber's base size conjecture). *Let G be a finite primitive permutation group on a finite set S . Then $b(G) = \Theta(\log |G| / \log |S|)$.*

There has been extensive work on Pyber's conjecture [Ser96; GM98; GSS98; LS02; Ben05; Faw13; LS14; BS15]. Recently, Duyan, Halasi, and Maróti announced a proof of this conjecture [DHM16].

We only need the special case of the conjecture for almost simple type, which is verified in [Ben05].

Theorem 8.7 ([Ben05]). *Let G be a finite primitive permutation group of almost simple type on a finite set S . Then $b(G) = \Theta(\log |G| / \log |S|)$.*

Bounds for non-standard actions. We also need a result on non-standard actions of primitive permutation groups of almost simple type. Recall that an action of a symmetric group $\text{Sym}(n)$ is *standard* if it is equivalent to the action on the set of k -subsets of $[n]$ for some $k \in [n]$, or the action on an orbit of the set of partitions of $[n]$, induced from the natural action of $\text{Sym}(n)$ on $[n]$ (see Chapter 7). And we say an action of $\text{Alt}(n)$ on a finite set S is standard if it is restricted from a standard action of $\text{Sym}(n)$ on S . Analogously, one can define standard actions of a classical group which, roughly speaking, are actions that permute subspaces (or pairs of subspaces of complementary dimension) of the natural module. See [LS99; Bur07] for the rigorous definition. Finally, an action of a primitive permutation group of almost simple type is *non-standard* if it is not a standard action.

It was conjectured in [Cam92; CK93] that the minimal base sizes of non-standard actions are bounded by an absolute constant $c \in \mathbb{N}$. This conjecture was proved by Liebeck and Shalev [LS99].⁵ We state the following weaker form of this result, where we do not distinguish standard and non-standard actions of classical groups. This weaker form is sufficient for our goal.

Theorem 8.8. *Let G be a finite primitive permutation group G of almost simple type, and let $T = \text{soc}(G)$. Then one of the following holds:*

⁵In addition, a chain of papers [Bur07; BLS09; BOW10; BGS11] shows that the minimum possible value of the constant c is 7.

1. G is permutation isomorphic to a symmetric group or an alternating group with a standard action.
2. T is a classical simple group.
3. $b(G) \leq c$, where $c \in \mathbb{N}$ is some absolute constant.

See [LS99, Theorem 1.3] for the original statement.

Proof of Theorem 8.4 for almost simple type. Now we are ready to prove Theorem 8.4 for a primitive permutation group G of almost simple type. In fact, we prove it in the following general form which applies to any subgroup $H \subseteq G$.

Lemma 8.4. *Let G be a primitive permutation group of almost simple type on a finite set S , and let H be a subgroup of G on S . Then for sufficiently large $N = \text{poly}(k(G)^{d_{\text{Sym}}(k(G))}, r(G), |S|) \geq |S|$, all strongly antisymmetric $\mathcal{P}_{H,N}$ -schemes are discrete on $H_x \in \mathcal{P}_{H,N}$ for all $x \in S$.*

Proof. Let $T = \text{soc}(G)$ and $\mathcal{P} = \mathcal{P}_{H,N}$. Consider the three cases in Theorem 8.8. First assume G is permutation isomorphic to a symmetric group $\text{Sym}(k)$ or an alternating group $\text{Alt}(k)$ with a standard action. Note $k \leq k(G)$. We have $H_{x,y} \in \mathcal{P}$ for all $x, y \in S$ provided $N \geq |S|^2$. We also have $H_{\{x,y\} \cup U} \in \mathcal{P}$ for all $x, y, z \in S$ and $U \subseteq H_{x,y}z$ satisfying $|H_{x,y}z| \leq k$ and $1 \leq |U| \leq d_{\text{Sym}}(k)$, provided that $N = k(G)^{\Omega(d_{\text{Sym}}(k(G)))}$ is sufficiently large. The lemma holds by Theorem 7.2 in this case.

Next assume T is a classical simple group. Then $|T| \leq r(G)$. It is also known by CFSG that $|\text{Out}(T)| = O(\log |T|)$ (see [Con+85]) and hence $|G| \leq |\text{Aut}(T)| = |T|^{O(1)} = r(G)^{O(1)}$. By Lemma 2.5 and Theorem 8.7, we have

$$d(H) \leq b(H) \leq b(G) = \Theta(\log |G| / \log |S|)$$

and hence $|S|^{d(G)} = |G|^{\Theta(1)} = r(G)^{O(1)}$. It follows that for sufficiently large $N = r(G)^{\Omega(1)}$, the subgroup system \mathcal{P} contains the system of stabilizers of depth $d(H)$. So the lemma also holds in this case.

Finally, in the last case of Theorem 8.8, we have $d(H) \leq b(H) \leq b(G) \leq c$, and the lemma holds for $N \geq |S|^c$. \square

Choosing $H = G$ in Lemma 8.4, we have

Corollary 8.1. *Theorem 8.4 holds for finite primitive permutation groups of almost simple type.*

8.4 Affine type

In this section, we prove Theorem 8.4 for finite primitive permutation groups of affine type. The following definitions are needed.

Definition 8.8 (irreducible / primitive linear group). *A group $H \subseteq \text{GL}(V)$ is said to be an irreducible linear group on V if H does not fix any subspace $W \subseteq V$ other than $\{0\}$ and V . And $H \subseteq \text{GL}(V)$ is said to be a primitive linear group on V if it is an irreducible linear group, and V cannot be written as a direct sum $V = \bigoplus_{i=1}^k V_i$ such that $k > 1$ and H permutes the direct summands V_i .*

The following fact is well known (see, e.g., [Sup76, Section I.4]).

Lemma 8.5. *Let G be a finite primitive permutation group G of affine type on a vector space V over a prime field \mathbb{F}_p . Then the stabilizer $G_0 \subseteq \text{GL}(V)$ of the origin $0 \in V$ is an irreducible linear group on V .*

We prove Theorem 8.4 for affine type by studying the stabilizer of the origin. In the following, we first discuss the case that this stabilizer is a primitive linear group (over \mathbb{F}_p), and then the case of (possibly imprimitive) irreducible linear groups.

Primitive linear groups. Our analysis is based on the work [LS02; LS14] on bases of primitive linear groups. We start with the following definitions.

Definition 8.9 (fully deleted permutation module [KL90]). *Fix a finite field \mathbb{F}_q and $k \in \mathbb{N}^+$. Define*

$$\begin{aligned} E(k, q) &:= \{(a, \dots, a) : a \in \mathbb{F}_q\} \subseteq \mathbb{F}_q^k, \\ M(k, q) &:= \{(a_1, \dots, a_k) \in \mathbb{F}_q^k : a_1 + \dots + a_k = 0\}, \\ U(k, q) &:= M(k, q) / (M(k, q) \cap E(k, q)). \end{aligned}$$

Let $\text{Sym}(k)$ act on \mathbb{F}_q^k by permuting the k coordinates, which induces an action on $U(k, q)$. We call $U(k, q)$ the fully deleted permutation module for $\text{Sym}(k)$ over \mathbb{F}_q .

Definition 8.10 (tensor product of linear groups). *Let V_1, \dots, V_k be vector spaces over a finite field \mathbb{F}_q . Let G_1, \dots, G_k be finite groups where $G_i \subseteq \text{GL}(V_i)$ for $i \in [k]$.*

Define an action of $G_1 \times \cdots \times G_k$ on the tensor product $U := V_1 \otimes \cdots \otimes V_k$ (over \mathbb{F}_q) by letting

$$(g_1, \dots, g_k) a_1 \otimes \cdots \otimes a_k = g_1 a_1 \otimes \cdots \otimes g_k a_k$$

and extending to all tensors multilinearly. This gives a linear representation $\rho : G_1 \times \cdots \times G_k \rightarrow \text{GL}(U)$. Write $g_1 \otimes \cdots \otimes g_k$ for $\rho(g_1, \dots, g_k) \in \text{GL}(U)$. And write $G_1 \otimes \cdots \otimes G_k$ for $\rho(G_1 \times \cdots \times G_k) \subseteq \text{GL}(U)$, called the tensor product of G_1, \dots, G_k (over \mathbb{F}_q).

We need the following structure theorem in [LS14] on primitive linear groups. See [LS14, Theorem 1] for a more detailed statement.

Theorem 8.9 ([LS14]). *Let p be a prime number, V a finite-dimensional vector space over \mathbb{F}_p , and G a primitive linear group on V . Choose the largest power q of p such that V can be identified with a vector space $V(q)$ over \mathbb{F}_q and $G \subseteq \Gamma\text{L}(V(q))$. Let $H := G \cap \text{GL}(V(q))$ act on $V(q)$. Then there exists an absolute constant $C \in \mathbb{N}^+$ such that either $b(H) \leq C$, or $V(q)$ can be identified with a tensor product over \mathbb{F}_q*

$$V(q) = \bigotimes_{i=1}^s U(k_i, q) \otimes W_0 \otimes \bigotimes_{j=1}^t W_j,$$

where $k_i \geq 5^6$ and $U(k_i, q)$ is the fully deleted permutation module for $\text{Sym}(k_i)$ over \mathbb{F}_q for $i \in [s]$, and W_j is vector space of dimension $d_j \in \mathbb{N}^+$ over \mathbb{F}_q for $0 \leq j \leq t$. Moreover, in the latter case, the group H is a subgroup of

$$\bigotimes_{i=1}^s \text{Sym}(k_i) \otimes D_0 \otimes \bigotimes_{j=1}^t D_j$$

acting on $V(q)$ that satisfies the following conditions:

1. For $i \in [s]$, the group $\text{Sym}(k_i)$ acts faithfully on $U(k_i, q)$ (see Definition 8.9).⁷
2. $D_0 \subseteq \text{GL}(W_0)$ acts on W_0 and $b(D_0) \leq C$.
3. For $j \in [t]$, the group D_j acting on W_j is the normalizer in $\text{GL}(W_j)$ of one of the quasisimple classical groups $\text{SL}_{d_j}(q_j)$, $\text{SU}_{d_j}(q_j^{1/2})$, $\text{Sp}_{d_j}(q_j)$, $\Omega_{d_j}(q_j) \subseteq$

⁶The condition $k_i \geq 5$ is implicit in [LS14]. If $k_i < 5$, we may always remove the factor $U(k_i, q)$ by replacing W_0 with $U(k_i, q) \otimes W_0$ (see [LS02, Lemma 3.3]).

⁷We regard $\text{Sym}(k)$ as a subgroup of $\text{GL}(U(k, q))$ via the faithful linear representation $\text{Sym}(k) \hookrightarrow \text{GL}(U(k, q))$.

$\mathrm{GL}_{d_j}(q_j)$.⁸ Here \mathbb{F}_{q_j} is a subfield of \mathbb{F}_q , and we identify $\mathrm{GL}_{d_j}(q_j)$ with a subgroup $\mathrm{GL}(W'_j) \subseteq \mathrm{GL}(W_j)$ for some vector space $W'_j \subseteq W_j$ over \mathbb{F}_{q_j} by fixing an \mathbb{F}_{q_j} -basis of W'_j that is also an \mathbb{F}_q -basis of W_j .

4. H contains the group $\bigotimes_{i=1}^s \mathrm{Alt}(k_i) \otimes \{e\} \otimes \bigotimes_{j=1}^t D_j^{(\infty)}$, where $D_j^{(\infty)}$ denotes the last term in the derived series of D_j .

The following lemma implies that the group D_j in Definition 8.9 for each $j \in [t]$ is a subgroup of $\mathbb{F}_q^\times \mathrm{GL}_{d_j}(q_j)$. For its proof, see [KL90, Proposition 4.5.1].

Lemma 8.6. *Suppose $\mathbb{F}_{q_0} \subseteq \mathbb{F}_q$, and $G \subseteq \mathrm{GL}_d(q)$ is one of the quasisimple classical groups $\mathrm{SL}_d(q_0), \mathrm{SU}_d(q_0^{1/2}), \mathrm{Sp}_d(q_0), \Omega_d(q_0) \subseteq \mathrm{GL}_d(q_0) \subseteq \mathrm{GL}_d(q)$. Then $N_{\mathrm{GL}_d(q)}(G) \subseteq \mathbb{F}_q^\times \mathrm{GL}_d(q_0)$.*

For convenience, we also make the following definition.

Definition 8.11 (primary tensor). *Use the notations in Theorem 8.9 and assume $b(G) > C$. So W is identified with the tensor product*

$$\bigotimes_{i=1}^s U(k_i, q) \otimes W_0 \otimes \bigotimes_{j=1}^t W_j$$

over \mathbb{F}_q by Theorem 8.9. We say an element $x \in V - \{0\}$ is a primary tensor if x is a pure tensor, i.e., $x = u_1 \otimes \cdots \otimes u_s \otimes w_0 \otimes w_1 \otimes \cdots \otimes w_t$, where $u_i \in U(k_i, q)$ for $i \in [s]$ and $w_j \in W_j$ for $0 \leq j \leq t$, and in addition,

1. for $i \in [s]$, $u_i \in U(k_i, q)$ is represented by an element in $M(k_i, q) \subseteq \mathbb{F}_{q_i}^{k_i}$ (see Definition 8.9) that has exactly two nonzero coordinates, and
2. for $j \in [t]$, $w_j \in W_j$ has the form $w_j = cw'_j$ where $c \in \mathbb{F}_q^\times$ and $w'_j \in W'_j$ (see Definition 8.9).

In addition, for two primary tensors $x, y \in V - \{0\}$, we write $x \sim y$ if x and y can be written as tensor products of vectors satisfying the above conditions and they differ at no more than one vector u_i or w_j . In other words, we can write

$$x = u_1 \otimes \cdots \otimes u_s \otimes w_0 \otimes \cdots \otimes w_t$$

⁸ For the definitions of these classical groups, see, e.g., [KL90; Asc00]. A group G is quasisimple if it equals its commutator subgroup and its inner automorphism group is simple, or equivalently, if it is a perfect central extension of a simple group [Asc00].

and either

$$y = u_1 \otimes \cdots \otimes u_{i-1} \otimes u'_i \otimes u_{i+1} \otimes \cdots \otimes u_s \otimes w_0 \otimes \cdots \otimes w_t$$

for some $i \in [s]$ and $u'_i \in U(k_i, q)$, or

$$y = u_1 \otimes \cdots \otimes u_s \otimes w_0 \otimes \cdots \otimes w_{j-1} \otimes w'_j \otimes w_{j+1} \otimes \cdots \otimes w_t$$

for some $0 \leq j \leq t$ and $w_j \in W_j$, so that the vectors u_i (resp. u'_i) and w_j (resp. w'_j) satisfy the above defining conditions of primary tensors.

Note that in Definition 8.11, a vector space $M(k_i, q)$ is spanned by vectors with exactly two nonzero coordinates, and W_j is spanned by vectors in W'_j over \mathbb{F}_q . So any $x \in V$ can be written as a finite sum of primary tensors. Also note that for any two primary tensors $x, y \in V - \{0\}$, there exists a finite sequence of primary tensors $x_0, \dots, x_k \in V - \{0\}$ such that $x_0 = x$, $x_k = y$, and $x_{i-1} \sim x_i$ for all $i \in [k]$.

Now we are ready to prove the following analogue of Theorem 8.4 for subgroups of primitive linear groups over \mathbb{F}_p .

Lemma 8.7. *Let G be a primitive linear group on a vector space V over \mathbb{F}_p as in Theorem 8.9, and let G' be a subgroup of G on V . Then for sufficiently large $N = \text{poly}(r(G), |V|) \geq |V|$, all strongly antisymmetric $\mathcal{P}_{G',N}$ -schemes are discrete on $G'_x \in \mathcal{P}_{G',N}$ for all $x \in V$.*

Proof. Use the notations in Theorem 8.9. Fix $\alpha \in \mathbb{F}_q^\times$ that does not lie in any proper subfield of \mathbb{F}_q . First assume $b(H) \leq C$. Let $B \subseteq V$ be a base of H of cardinality at most C . Pick a nonzero element $z \in B$. Then $B \cap \{\alpha z\}$ is a base of G since $G_{z,\alpha z} \subseteq G \cap \text{GL}(V(q)) = H$. So $d(G) \leq b(G) \leq C + 1$. Then for $N \geq |V|^{C+1}$, all strongly antisymmetric $\mathcal{P}_{G',N}$ -schemes are discrete on $G'_x \in \mathcal{P}_{G',N}$ for all $x \in V$, as desired.

So assume $b(H) > C$. Then we have $V(q) = \bigotimes_{i=1}^s U(k_i, q) \otimes W_0 \otimes \bigotimes_{j=1}^t W_j$ and $H \subseteq \bigotimes_{i=1}^s \text{Sym}(k_i) \otimes D_0 \otimes \bigotimes_{j=1}^t D_j$ as in Theorem 8.9. Let $\mathcal{P} = \mathcal{P}_{G',N}$ and let \mathcal{C} be a strongly antisymmetric \mathcal{P} -scheme. Let $N \geq |V|^4$ so that $G'_{x,y,z,w} \in \mathcal{P}$ for all $x, y, z, w \in V$. Fix $x \in V$. We want to prove that \mathcal{C} is discrete on G'_x . By Lemma 2.3, it suffices to prove that \mathcal{C} is discrete on $G'_{x,\alpha x}$.

Consider the diagonal action of G' on $V \times V$, and let O be the G' -orbit of $(x, \alpha x)$. The elements in O are of the form $(y, \beta y)$, where $y \in V$ and $\beta \in \mathbb{F}_q^\times$ is a conjugate

of α , i.e., $\beta = {}^g\alpha$ for some $g \in \text{Gal}(\mathbb{F}_q/\mathbb{F}_p)$. Also note that for any distinct $y, z \in V$, the difference $z - y$ can be written as a finite sum of primary tensors. By Lemma 7.6, it suffices to prove, for all distinct $y, z \in V$ whose difference $z - y$ is a primary tensor and conjugates $\beta, \gamma \in \mathbb{F}_q^\times$ of α , that $\mathcal{C}|_{G'_{y,\beta y}}$ is discrete on $G'_{y,\beta y,z,\gamma z} = G'_{y,\beta y,z}$. Fix such y, z, β, γ .

Let $H' = G' \cap H$. Then $G'_{y,\beta y} = H'_y$ and $G'_{y,\beta y,z} = H'_{y,z} = H'_{y,z-y}$. So we want to prove that $\mathcal{C}|_{H'_y}$ is discrete on $H'_{y,z-y}$. Note that every element in the H'_y -orbit of $z - y$ is a primary tensor. As noted after Definition 8.11, for any two primary tensors $u, v \in V - \{0\}$, there exists a finite sequence of primary tensors $x_1, \dots, x_k \in V - \{0\}$ such that $x_1 = u$, $x_k = v$, and $x_{i-1} \sim x_i$ for all $i \in [k]$. Again by Lemma 7.6, it suffices to prove, for all primary tensors $u, v \in V - \{0\}$ satisfying $u \sim v$, that $\mathcal{C}|_{H'_{y,u}}$ is discrete on $H'_{y,u,v}$ (note $H'_{y,u} = G'_{y,\beta y,u}$, $H'_{y,u,v} = G'_{y,\beta y,u,v} \in \mathcal{P}$). Fix such u, v . Suppose $u = u_1 \otimes \dots \otimes u_s \otimes w_0 \otimes \dots \otimes w_t$ where $u_i \in U(k_i, q)$ for $i \in [s]$ and $w_j \in W_j$ for $0 \leq j \leq t$ satisfy the conditions in Definition 8.11.

First consider the case that v has the form

$$v = u_1 \otimes \dots \otimes u_{r-1} \otimes u'_r \otimes u_{r+1} \otimes \dots \otimes u_s \otimes w_0 \otimes \dots \otimes w_t,$$

where $r \in [s]$ and $u'_r \in U(k_r, q)$ is represented by a vector $\tilde{u}'_r \in M(k_r, q)$ with exactly two nonzero coordinates. Let $n := |H'_{y,u}v|$. We prove a bound on n . Consider an element $g \in H'_{y,u}$. By Theorem 8.9, we may write $g = g_1 \otimes \dots \otimes g_s \otimes h_0 \otimes \dots \otimes h_t$ where $g_i \in \text{Sym}(k_i)$ for $i \in [s]$ and $h_j \in D_j$ for $0 \leq j \leq t$. As g fixes u , we know $g_r \in \text{Sym}(k_r)$ sends u_r to cu_r for some $c \in \mathbb{F}_q^\times$. As $k_r \geq 5$ and \tilde{u}'_r has exactly two nonzero coordinates, either g_r fixes \tilde{u}'_r and $c = 1$, or g_r swaps the two nonzero coordinates of \tilde{u}'_r and $c = -1$. From ${}^g u = u$, it is easy to see that

$${}^g v = c^{-1}(u_1 \otimes \dots \otimes u_{r-1} \otimes {}^{g_r}u'_r \otimes u_{r+1} \otimes \dots \otimes u_s \otimes w_0 \otimes \dots \otimes w_t).$$

So c and ${}^{g_r}u'_r$ determine ${}^g v$. The number of possible values of ${}^{g_r}u'_r$ is bounded by $|\text{Sym}(k_r)u'_r| \leq |\text{Sym}(k_r)\tilde{u}'_r| \leq k_r^2$. It follows that $n = |H'_{y,u}v| \leq 2k_r^2$. Also note that $|V| \geq |U(k_r, q)| \geq q^{k_r-2}$ and hence $k_r = O(\log |V|)$. Let $N \geq n^{d_{\text{Sym}}(n)} = |V|^{O(1)}$. Then $\mathcal{P}|_{H'_{y,u}}$ contains the system of stabilizers of depth $d_{\text{Sym}}(n)$ with respect to the action of $H'_{y,u}$ on $H'_{y,u}v$. So $\mathcal{C}|_{H'_{y,u}}$ is discrete on $H'_{y,u,v}$, as desired.

Next consider the case that v has the form

$$v = u_1 \otimes \dots \otimes u_s \otimes w'_0 \otimes w_1 \otimes \dots \otimes w_t$$

for some $w'_0 \in W_r$. Let $B \subseteq W_0$ be a base of D_0 of cardinality at most C , which exists by Theorem 8.9. For any subset T of W_0 , define

$$\tilde{T} := \{u_1 \otimes \cdots \otimes u_s \otimes a \otimes w_1 \otimes \cdots \otimes w_t : a \in T\} \subseteq V.$$

Consider $g = g_1 \otimes \cdots \otimes g_s \otimes h_0 \otimes \cdots \otimes h_t \in (H'_{y,u})_{\tilde{B}}$ where $g_i \in \text{Sym}(k_i)$ for $i \in [s]$ and $h_j \in D_j$ for $0 \leq j \leq t$. As g fixes every element in \tilde{B} , we see $h_0 \in D_0$ scales every element in B by the same factor $c \in \mathbb{F}_q^\times$. Then $c^{-1}h_0 \in (D_0)_B = \{e\}$ and hence $h_0 = c$. Therefore h_0 scales every element in W_0 by the factor c . So $g \in (H'_{y,u})_{\tilde{W}_0}$. It follows that $(H'_{y,u})_{\tilde{B}} = (H'_{y,u})_{\tilde{W}_0}$. Also note that $H'_{y,u}$ fixes \tilde{W}_0 setwisely. So $(H'_{y,u})_{\tilde{W}_0}$ is normal in $H'_{y,u}$. Let $N \geq |V|^{C+3}$ so that $(H'_{y,u})_{\tilde{W}_0} = H_{\{y,\beta y,u\} \cup \tilde{B}} \in \mathcal{P}$. By antisymmetry of $\mathcal{C}|_{H'_{y,u}}$, we know $\mathcal{C}|_{H'_{y,u}}$ is discrete on $(H'_{y,u})_{\tilde{W}_0}$. By Lemma 2.3, it is also discrete on $H'_{y,u,v} \supseteq (H'_{y,u})_{\tilde{W}_0}$, as desired.

Finally, consider the case that v has the form

$$v = u_1 \otimes \cdots \otimes u_s \otimes w_0 \otimes \cdots \otimes w_{r-1} \otimes w'_r \otimes w_{r+1} \otimes \cdots \otimes w_t$$

for some $0 \leq r \leq t$ and $w'_r \in W_r$ such that $w'_r = c_0 w''_r$ for some $c_0 \in \mathbb{F}_q^\times$ and $w''_r \in W'_r$. We claim that $|H'_{y,u}v| \leq q_r^{d_r}$. To see this, consider $g = g_1 \otimes \cdots \otimes g_s \otimes h_0 \otimes \cdots \otimes h_t \in H'_{y,u}$ where $g_i \in \text{Sym}(k_i)$ for $i \in [s]$ and $h_j \in D_j$ for $0 \leq j \leq t$. By Lemma 8.6, we have $h_r \in \mathbb{F}_q^\times \text{GL}(W'_r)$. As g fixes u , we have $h^r w_r = c_1 w_r$ for some $c_1 \in \mathbb{F}_q^\times$. Then it is easy to see that

$$\begin{aligned} g v &= c_1^{-1} (u_1 \otimes \cdots \otimes u_s \otimes w_0 \otimes \cdots \otimes w_{r-1} \otimes h^r w'_r \otimes w_{r+1} \otimes \cdots \otimes w_t) \\ &= u_1 \otimes \cdots \otimes u_s \otimes w_0 \otimes \cdots \otimes w_{r-1} \otimes c_1^{-1} h^r w'_r \otimes w_{r+1} \otimes \cdots \otimes w_t. \end{aligned}$$

As $h_r \in \mathbb{F}_q^\times \text{GL}(W'_r)$, we may write $h_r = c_2 h'_r$ for some $c_2 \in \mathbb{F}_q^\times$ and $h'_r \in \text{GL}(W'_r)$. Note that $h'_r = c_2^{-1} h_r \in \text{GL}(W'_r) \subseteq \text{GL}(W_r)$ sends w_r to $c_2^{-1} c_1 w_r$. Then $c_2^{-1} c_1 \in \mathbb{F}_q^\times$. Therefore $c_1^{-1} h_r = c_1^{-1} c_2 h'_r \in \text{GL}(W'_r)$. It follows that

$$|H'_{y,u}v| \leq |\text{GL}(W'_r)w'_r| = |\text{GL}(W'_r)w''_r| \leq |W'_r| = q_r^{d_r}$$

as claimed. Let $V' \subseteq V$ be the vector space over \mathbb{F}_q spanned by the elements in $H'_{y,u}v$. Let $B \subseteq H'_{y,u}v$ be an \mathbb{F}_q -basis of V' . Then $|B| = \dim_{\mathbb{F}_q} V' \leq \dim_{\mathbb{F}_q} W_r = d_r$. Note $q_r^{d_r} = |D_r|^{O(1)} = r(H)^{O(1)} = r(G)^{O(1)}$. Let $N \geq q_r^{d_r} \geq |H'_{y,u}v|^{d_r}$, so that $\mathcal{P}|_{H'_{y,u}}$ contains the system of stabilizers of depth d_r with respect to the action of $H'_{y,u}$ on $H'_{y,u}v$. Then $(H'_{y,u})_{V'} = (H'_{y,u})_B \in \mathcal{P}$. Note that $H'_{y,u}$ fixes V' setwisely, and hence $(H'_{y,u})_{V'}$ is normal in $H'_{y,u}$. By antisymmetry of $\mathcal{C}|_{H'_{y,u}}$, we know $\mathcal{C}|_{H'_{y,u}}$ is discrete on $(H'_{y,u})_{V'}$. By Lemma 2.3, it is also discrete on $H'_{y,u,v} \supseteq (H'_{y,u})_{V'}$, as desired. \square

Irreducible linear groups. Next we extend Lemma 8.7 to irreducible linear groups over \mathbb{F}_p . For a group $G \subseteq \text{GL}(V)$ and a subspace $W \subseteq V$, the setwise stabilizer $G_{\{W\}}$ acts on W , which gives a linear representation $\pi_W : G_{\{W\}} \rightarrow \text{GL}(W)$. Write $G|_W$ for its image $\pi_W(G_{\{W\}}) \subseteq \text{GL}(W)$.

We need the following lemma, whose proof can be found in, e.g., [Sup76, Section IV.15].

Lemma 8.8. *Let G be an irreducible linear group on a finite-dimensional vector space $V \neq \{0\}$. Then there exists a nonzero subspace $W \subseteq V$ such that $G|_W$ is a primitive linear group on W , and G permutes the subspaces in the set $\{^gW : g \in G\}$.*

We have the following generalization of Lemma 8.7.

Lemma 8.9. *Let G be an irreducible linear group on a vector space V over \mathbb{F}_p , and let G' be a subgroup of G on V . Then for sufficiently large $N = \text{poly}(r(G), |V|) \geq |V|$, all strongly antisymmetric $\mathcal{P}_{G',N}$ -schemes are discrete on $G'_x \in \mathcal{P}_{G',N}$ for all $x \in V$.*

Proof. Assume $V \neq \{0\}$ as otherwise the claim is trivial. By Lemma 8.8, we may choose a nonzero subspace $W \subseteq V$ such that $G|_W$ is a primitive linear group on W , and G permutes the subspaces in the set $S_W := \{^gW : g \in G\}$. Note $|S_W| = \log |V| / \log |W| = O(\log |V|)$. We claim $r(G|_W) = \text{poly}(r(G), |V|)$. To see this, consider a classical group H that is a composition factor of $G_{\{W\}}$. The group G permutes the subspaces in S_W , which gives a permutation representation $\rho : G \rightarrow \text{Sym}(S_W)$. Then H is either a composition factor of $\rho(G_{\{W\}})$ or that of $\text{Ker}(\rho) \cap G_{\{W\}} = \text{Ker}(\rho)$. In the former case, the group H is a subquotient of $\text{Sym}(S_W)$. And Lemma 8.2 implies that $|H| = r(H)$ is polynomial in $|S_W|^{\log |S_W|} = |V|^{O(1)}$. In the latter case, we have $|H| \leq r(G)$ since $\text{Ker}(\rho) \trianglelefteq G$. So in either case, we have $r(G|_W) \leq r(G_{\{W\}}) = \text{poly}(r(G), |V|)$.

Let $\mathcal{P} = \mathcal{P}_{G',N}$. Let $N \geq |V|^3$ so that $G'_{x,y,z} \in \mathcal{P}$ for all $x, y, z \in V$. Suppose $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ is a strongly antisymmetric \mathcal{P} -scheme. We want to show that all strongly antisymmetric \mathcal{P} -schemes are discrete on G'_x for all $x \in V$. Note that for any $x, y \in V$, we may choose a sequence of elements $z_0, \dots, z_t \in S$ such that $z_0 = x, z_t = y$, and for all $i \in [t]$, the vector $z_i - z_{i-1}$ is in gW for some $g \in G$. By Lemma 7.6, it suffices to prove, for all $x, y \in V$ and $g \in G$ satisfying $x - y \in ^gW$, that $\mathcal{C}|_{G'_x}$ is discrete on $G'_{x,y}$. Fix such $x, y \in V$ and $g \in G$.

Let $z = y - x$. Note that $G'_{x,y} = G'_{x,z}$. Every element in $G'_x z$ is in a subspace $g'W$ for some $g' \in G$. Consider distinct $u, v \in G'_x z$ that are in the same subspace $g'W$. Pick $h, h' \in G'_x$ such that $u = hz$ and $v = h'z$. We claim that $G'_{x,y}h^{-1}, G'_{x,y}h'^{-1} \in G'_{x,y} \setminus G'_x$ are in different blocks of $C_{G'_{x,y}}|_{G'_x}$. By Lemma 7.5, it suffices to show that $\mathcal{C}|_{G'_{x,u}}$ and $\mathcal{C}|_{G'_{x,v}}$ are discrete on $G'_{x,u,v}$. We only prove it for $\mathcal{C}|_{G'_{x,u}}$ as the claim for $\mathcal{C}|_{G'_{x,v}}$ is symmetric.

Note that $G'_{x,u}$ is a subgroup of $G|_{\{g'W\}}$ since $u \in g'W$ and G permutes the subspaces in the set S_W . Define

$$\mathcal{P}' := \{(G'_{x,u}|_{g'W})_B : B \subseteq g'W, (G'_{x,u})_B \in \mathcal{P}\},$$

which is a subgroup system over $G'_{x,u}|_{g'W}$. By Lemma 6.4, it suffices to show that all strongly antisymmetric \mathcal{P}' -schemes are discrete on $(G'_{x,u}|_{g'W})_v$. Let $N' := \lfloor N/|V|^2 \rfloor$. Then \mathcal{P}' contains the subgroup system $\mathcal{P}_{G'_{x,u}|_{g'W}, N'}$ with respect to the faithful action of $G'_{x,u}|_{g'W}$ on $g'W$. Note that $G'_{x,u}|_{g'W}$ is a subgroup of $G|_{g'W}$ and the latter is a primitive linear group since $G|_W$ is a primitive linear group. Also note $G|_{g'W} \cong G|_W$ and hence $r(G|_{g'W}) = r(G|_W) = \text{poly}(r(G), |V|)$. Applying Lemma 8.7 to $G'_{x,u}|_{g'W} \subseteq G|_{g'W}$, we see that all strongly antisymmetric $\mathcal{P}_{G'_{x,u}|_{g'W}, N'}$ -schemes are discrete on $(G'_{x,u}|_{g'W})_v$, and hence all strongly antisymmetric \mathcal{P}' -schemes are discrete on $(G'_{x,u}|_{g'W})_v$, as desired. This proves the claim that $G'_{x,y}h^{-1}$ and $G'_{x,y}h'^{-1}$ are in different blocks of $C_{G'_{x,y}}|_{G'_x}$ given that $u = hz$ and $v = h'z$ are distinct elements in the same subspace $g'W$.

Consider an arbitrary block $\{G'_{x,y}g_1^{-1}, \dots, G'_{x,y}g_s^{-1}\}$ of $C_{G'_{x,y}}|_{G'_x}$ of cardinality $s \in \mathbb{N}^+$. By the claim just proved, the elements $g_1 z, \dots, g_s z$ are in distinct subspaces in the set S_W . So $s \leq |S_W| = O(\log |V|)$. Therefore we have $m(s) = O(\log s) = O(\log \log |V|)$ by Theorem 7.1 (see Definition 7.3 for the definition of $m(\cdot)$). Choose the largest $m \in \mathbb{N}$ satisfying $|G'_x z|^m \leq N$. By definition, the subgroup system $\mathcal{P}|_{G'_x}$ contains the system of stabilizers of depth m over G'_x (with respect to the action of G'_x on $G'_x z$). Lemma 2.7 and Theorem 6.1 then imply the existence of a strongly antisymmetric m -scheme $\Pi = \{P_1, \dots, P_m\}$ on $G'_x z$ such that P_1 has a block of cardinality s . Note $|G'_x z| \leq |S_W| \cdot |W|$. And we have $|S_W|^{m(s)} = (\log |V|)^{O(\log \log |V|)} = |V|^{O(1)}$ and $|W|^{m(s)} = |W|^{O(\log |S_W|)} = |V|^{O(1)}$. Then for sufficiently large $N = |V|^{\Omega(1)}$, we have

$$|G'_x z|^{m(s)} \leq (|S_W| \cdot |W|)^{m(s)} \leq N,$$

and hence $m \geq m(s)$. Theorem 7.1 then forces $s = 1$. So $\mathcal{C}|_{G'_x}$ is discrete on $G'_{x,z} = G'_{x,y}$, as desired. \square

Now we are ready to prove Theorem 8.4 for finite primitive permutation groups of affine type.

Lemma 8.10. *Theorem 8.4 holds for finite primitive permutation groups of affine type.*

Proof. Let G be a finite primitive permutation groups of affine type on a vector space V over a prime field \mathbb{F}_p . Then the stabilizer $G_0 \subseteq \text{GL}(V)$ of the origin $0 \subseteq V$ is an irreducible linear group by Lemma 8.5. Let $V^\# \subseteq G$ be the group of translations. Then $G \cong V^\# \rtimes G_0$ and hence $r(G_0) \leq r(G)$.

Let \mathcal{C} be a strongly antisymmetric $\mathcal{P}_{G,N}$ -scheme. By Lemma 7.6, it suffices to prove for all $x, y \in V$ that $\mathcal{C}|_{G_x}$ is discrete on $G_{x,y}$. Fix such $x, y \in V$. By invariance of \mathcal{C} and the fact that G acts transitively on V , we may assume $x = 0$. So we want to show that $\mathcal{C}|_{G_0}$ is discrete on $G_{0,y}$. This follows from Lemma 8.9 applied to the irreducible linear group G_0 on V and the subgroup system $\mathcal{P}_{G,N}|_{G_0}$ over G_0 . \square

8.5 Diagonal type

In this section, we verify Theorem 8.4 for a finite primitive permutation group G of diagonal type. By Definition 8.5, we may assume G is a permutation group satisfying $M \subseteq G \subseteq W$ and acting on a set $S := D \setminus W$ by inverse right translation, where

$$A = \{(a_1, \dots, a_k) \in \text{Aut}(T)^k : a_i \text{Inn}(T) = a_j \text{Inn}(T) \text{ for all } i, j \in [k]\},$$

$$W = A \rtimes \text{Sym}(k),$$

$$M = \text{Inn}(T)^k \subseteq A \subseteq W,$$

$$D = \{(a, \dots, a)\pi : a \in \text{Aut}(T), \pi \in \text{Sym}(k)\} \subseteq W$$

for a noncyclic finite simple group T and an integer $k \geq 2$. The cardinality of S is $|W|/|D| = |T|^{k-1}$.

Let x_0 denote the element $De \in S$, so that $G_{x_0} = D \cap G$. It is a consequence of CFSG that every finite simple group is generated by at most two elements ([AG84]). So we can choose $r, s \in \text{Inn}(T) - \{e\}$ that generate $\text{Inn}(T) \cong T$. For $g \in \text{Inn}(T)$, define $a_g := (g, e, \dots, e) \in M \subseteq G$. We have the following lemma.

Lemma 8.11. *For $U = \{x_0, {}^{a_r}x_0, {}^{a_s}x_0, {}^{a_{rs}}x_0\}$, it holds that $W_U = \text{Sym}(k)_1$.*

Proof. Note that

$$W_U = D \cap a_r D a_r^{-1} \cap a_s D a_s^{-1} \cap a_{rs} D a_{rs}^{-1}$$

from which it is straightforward to see $\text{Sym}(k)_1 \subseteq W_U$.

For the other direction, consider $g = (a, \dots, a)\pi \in D \subseteq W_U$, where $a \in \text{Aut}(T)$ and $\pi \in \text{Sym}(k)$. We have

$$a_r^{-1}ga_r = a_r^{-1}(a, \dots, a)\pi a_r = a_r^{-1}(a, \dots, a)^\pi a_r \pi \in D \quad (8.1)$$

since $a_r^{-1}W_U a_r \subseteq D$.

First assume $k > 2$. Suppose π sends 1 to $i \in [k]$. Note that all coordinates of a_r (resp. $^\pi a_r$) are identity except that the first (resp. i th) coordinate is $r \neq e$. As $k > 2$ and $a_r^{-1}(a, \dots, a)^\pi a_r \pi \in D$, we must have $i = 1$ and $r^{-1}ar = a$. So $\pi \in \text{Sym}(k)_1$. The same argument using the fact $a_s^{-1}W_U a_s \subseteq D$ implies $s^{-1}as = a$. Then a commutes with $\langle r, s \rangle = \text{Inn}(T)$. Note that the isomorphism $T \cong \text{Inn}(T)$ sending $h \in T$ to the inner automorphism $x \mapsto h x h^{-1}$ is an equivalence between the action of $\text{Aut}(T)$ on T and that on $\text{Inn}(T)$ by conjugation. So a fixes T pointwisely, which implies $a = e$. Then we have $g = \pi \in \text{Sym}(k)_1$, as desired.

Next assume $k = 2$. If $\pi = e$, we have $a_r^{-1}ga_r = (r^{-1}ar, a) \in D$ by (8.1). So $r^{-1}ar = a$, and the same argument using the fact $a_s^{-1}W_U a_s \subseteq D$ implies $s^{-1}as = a$. Again we conclude that a commutes with $\langle r, s \rangle = \text{Inn}(T)$, which implies $a = e \in \text{Sym}(k)_1$. Now consider the case $\pi \neq e$, i.e., $\pi = (1\ 2) \in \text{Sym}(2)$. Note that the proof for the previous case $\pi = e$ shows $W_U \cap A = \{e\}$. Therefore

$$|W_U| = [W_U : W_U \cap A] \leq [W : A] = |\text{Sym}(k)| = 2.$$

The lemma is trivial if $|W_U| = 1$. So assume $|W_U| = 2$. Then $W_U = \{e, g\}$, where $g = (a, a)\pi$ is as above. By (8.1), we have $(r^{-1}a, ar)\pi \in D$. So $ara^{-1} = r^{-1}$. The same argument using the facts $a_s^{-1}W_U a_s \subseteq D$ and $a_{rs}^{-1}W_U a_{rs} \subseteq D$ implies $asa^{-1} = s^{-1}$ and $arsa^{-1} = (rs)^{-1} = s^{-1}r^{-1}$. On the other hand, we have $arsa^{-1} = (ara^{-1})(asa^{-1}) = r^{-1}s^{-1}$. So r commutes with s . Then $T = \langle r, s \rangle$ is abelian, contradicting the assumption that T is a noncyclic finite simple group. \square

We prove Theorem 8.4 for a finite primitive permutation group G of diagonal type in the following general form that applies to any subgroup $H \subseteq G$.

Lemma 8.12. *Let G be a finite primitive permutation group of diagonal type on $S = D \setminus W$ as above, and let H be a subgroup of G on S . Then for sufficiently large $N = \text{poly}(|S|) \geq |S|$, all strongly antisymmetric $\mathcal{P}_{H,N}$ -schemes are discrete on $H_x \in \mathcal{P}_{H,N}$ for all $x \in S$.*

Proof. Let $\mathcal{P} = \mathcal{P}_{H,N}$. By choosing $N \geq |S|^2$, we may assume $H_{x,y} \in \mathcal{P}$ for all $x, y \in S$. Let $\mathcal{C} = \{C_{H'} : H' \in \mathcal{P}\}$ be a strongly antisymmetric \mathcal{P} -scheme. Define Z to be the set of elements $g \in M = \text{Inn}(T)^k$ such that g has exactly one coordinate different from the identity. Note that $g \in Z$ iff $g^{-1} \in Z$, and the elements in Z generate M . Also note that M acts transitively on S . Then by Lemma 7.6, it suffices to show that for all $x \in S$ and $g \in Z$, the \mathcal{P} -scheme $\mathcal{C}|_{H_x}$ is discrete on $H_{x,gx} \in \mathcal{P}|_{H_x}$. Fix $x \in S$ and $g \in Z$. As M acts transitively on S , there exists $h \in M$ sending x to x_0 . Let $y := {}^h(gx) = hgh^{-1}x_0$. By invariance of \mathcal{C} , it suffices to show that $\mathcal{C}|_{H_{x_0}}$ is discrete on $H_{x_0,y} \in \mathcal{P}|_{H_{x_0}}$.

Let $g' := hgh^{-1}$, so that $y = g'x_0$. Note $g' \in Z$. Suppose the i th coordinate of g' is different from the identity. Choose $U = \{x_0, {}^{a_r}x_0, {}^{a_s}x_0, {}^{a_{rs}}x_0\}$ as in Lemma 8.11, and let $U' := ({}^1 i)U$. As $(1 i) \subseteq D$ fixes x_0 , we have $x_0 \in U'$. By Lemma 8.11, we have

$$H_{U'} = \text{Sym}(k)_i \cap H = \text{Sym}(k)_i \cap H_{x_0}. \quad (8.2)$$

Note $g' \in N_G(\text{Sym}(k)_i)$, and hence

$$\text{Sym}(k)_i = g' \text{Sym}(k)_i g'^{-1} \subseteq g' D g'^{-1} = g' W_{x_0} g'^{-1} = W_y.$$

So $H_{U'} \subseteq H_{x_0,y}$. We have $H_{U'} \in \mathcal{P}$ provided that $N \geq |S|^{|U'|} = |S|^{O(1)}$. By Lemma 2.3, it suffices to prove that $\mathcal{C}|_{H_{x_0}}$ is discrete on $H_{U'}$. By Lemma 7.6, it suffices to show, for all $h, h' \in H_{x_0}$, that (1) $H_{h_{U'} \cup h'_{U'}} \in \mathcal{P}|_{H_{x_0}}$ and (2) $\mathcal{C}|_{H_{h_{U'}}}$ is discrete on $H_{h_{U'} \cup h'_{U'}}$. Fix $h, h' \in H_{x_0}$. We have $H_{h_{U'} \cup h'_{U'}} \in \mathcal{P}|_{H_{x_0}}$ provided that $N \geq |S|^{|h_{U'} \cup h'_{U'}|} = |S|^{O(1)}$.

So it remains to prove that $\mathcal{C}|_{H_{h_{U'}}}$ is discrete on $H_{h_{U'} \cup h'_{U'}}$. Write $h = b\pi$ and $h' = b'\pi'$ where $b, b' \in A$ and $\pi, \pi' \in \text{Sym}(k)$. As $h \in H_{x_0} \subseteq D$, the k coordinates of b are equal. So b commutes with $\text{Sym}(k)$. By (8.2), we have

$$H_{h_{U'}} = h(\text{Sym}(k)_i \cap H_{x_0})h^{-1} = \text{Sym}(k)_{\pi_i} \cap H_{x_0}.$$

Similarly, we have $H_{h'_{U'}} = \text{Sym}(k)_{\pi'_i} \cap H_{x_0}$ and

$$H_{h_{U'} \cup h'_{U'}} = \text{Sym}(k)_{\pi_i, \pi'_i} \cap H_{x_0}.$$

Let $n := [H_{h_{U'}} : H_{h_{U'} \cup h'_{U'}}]$. We have

$$n \leq [\text{Sym}(k)_{\pi_i} : \text{Sym}(k)_{\pi_i, \pi'_i}] \leq k.$$

Also note $k = \log |S| / \log |T| + 1 = O(\log |S|)$. Consider the action of $H_{h_{U'}}$ on $H_{h_{U'} \cup h'_{U'}} \setminus H_{h_{U'}}$ by inverse right translation. Each one-point stabilizer with

respect to this action is a pointwise stabilizer of a set $S' \subseteq S$ of cardinality at most $|U'| = O(1)$. Choose sufficiently large $N \geq n^{|U'|d_{\text{Sym}}(n)} = k^{O(\log k)} = |S|^{O(1)}$ so that $\mathcal{P}|_{H_{h_{U'}}$ contains the system of stabilizers of depth $d_{\text{Sym}}(n)$ with respect to this action. Then all strongly antisymmetric $\mathcal{P}|_{H_{h_{U'}}$ -schemes, including $\mathcal{C}|_{H_{h_{U'}}$, are discrete on $H_{h_{U' \cup h'_{U'}}$, as desired. \square

Choosing $H = G$ in Lemma 8.12, we have

Corollary 8.2. *Theorem 8.4 holds for finite primitive permutation groups of diagonal type.*

8.6 Product type and twisted wreath type

In this section, we verify Theorem 8.4 for finite primitive permutation groups of product type and those of twisted wreath type.

Product type. Suppose G is a finite primitive permutation group of product type. By Definition 8.6, there exist an integer $k \geq 2$ and a primitive permutation group H on a finite set Γ that is of almost simple type or diagonal type such that G is a subgroup of $W := H \wr \text{Sym}(k) = H^k \rtimes \text{Sym}(k)$ acting on $S := \Gamma^k$, and $M := \text{soc}(H)^k \subseteq W$ is a subgroup of G .

We prove Theorem 8.4 for a finite primitive permutation group G of product type in the following general form that applies to any subgroup $G' \subseteq G$.

Lemma 8.13. *Let G be a finite primitive permutation group of product type on S as above. Let G' be a subgroup of G on S . Then for sufficiently large $N = \text{poly}(k(G)^{d_{\text{Sym}}(k(G))}, r(G), |S|) \geq |S|$, all strongly antisymmetric $\mathcal{P}_{G',N}$ -schemes are discrete on $G'_x \in \mathcal{P}_{G',N}$ for all $x \in S$.*

Proof. Let $\mathcal{P} = \mathcal{P}_{G',N}$. Choose $N \geq |S|^3$ so that $G'_{x,y,z} \in \mathcal{P}$ for all $x, y, z \in S$. Suppose $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ is a strongly antisymmetric \mathcal{P} -scheme. Fix $x \in S$. We prove that \mathcal{C} is discrete on G'_x . Note that for any $y, z \in S$, we may choose a sequence of elements $y_0, \dots, y_t \in S$ such that $y_0 = y$, $y_t = z$, and for all $i \in [t]$, the elements $y_{i-1}, y_i \in S = \Gamma^k$ differ at exactly one coordinate. By Lemma 7.6, it suffices to prove, for all $y, z \in S$ differing at exactly one coordinate, that $\mathcal{C}|_{G'_y}$ is discrete on $G'_{y,z}$. Fix such $y, z \in S$. Also note that all elements in $G'_y z$ differ from y at exactly one coordinate. In particular, we have $|G'_y z| \leq k|\Gamma|$.

Consider $u, v \in G'_y z$ differing from y at the same coordinate whose index is denoted by $i \in [k]$. Pick $g, g' \in G'_y$ such that $u = g z$ and $v = g' z$. We claim that $G'_{y,z} g^{-1}, G'_{y,z} g'^{-1} \in G'_{y,z} \setminus G'_y$ are in different blocks of $C_{G'_{y,z}}|_{G'_y} \in \mathcal{P}|_{G'_y}$. By Lemma 6.3 and Lemma 7.5, it suffices to verify that $\mathcal{C}|_{G'_{y,u}}$ and $\mathcal{C}|_{G'_{y,v}}$ are discrete on $G'_{y,u,v}$. We only prove it for $\mathcal{C}|_{G'_{y,u}}$ since the claim for $\mathcal{C}|_{G'_{y,v}}$ is symmetric. Note that $\mathcal{C}|_{G'_{y,u}}$ is a strongly antisymmetric $\mathcal{P}|_{G'_{y,u}}$ -scheme by Lemma 6.3. We show that in fact all strongly antisymmetric $\mathcal{P}|_{G'_{y,u}}$ -schemes are discrete on $G'_{y,u,v}$. As $G'_{y,u}$ fixes y and u which differ at the i th coordinate, the image of $G'_{y,u}$ under the quotient map $H \wr \text{Sym}(k) \rightarrow \text{Sym}(k)$ is contained in $\text{Sym}(k)_i$. Define

$$P := \{(g_1, \dots, g_k)\pi \in G'_{y,u} : g_i = e\} \subseteq G'_{y,u}.$$

Then P is a normal subgroup of $G'_{y,u}$. Suppose $v = (v_1, \dots, v_k) \in S = \Gamma^k$. Define

$$S' := \{(v_1, \dots, v_{i-1}, v'_i, v_{i+1}, \dots, v_k) \in S : v'_i \in \Gamma\}.$$

The action of $G'_{y,u}$ on S restricts to an action on S' which factors through $\bar{G} := G'_{y,u}/P$. And the action of \bar{G} on S' is permutation isomorphic to H' on Γ , where $H' \subseteq H$ is defined by

$$H' := \{g \in H : (g_1, \dots, g_k)\pi \in G'_{y,u}, g_i = g\}.$$

Let $N' = \lfloor N/|S|^2 \rfloor$. Note $\mathcal{P}_{G'_{y,u}, N'} \subseteq \mathcal{P}|_{G'_{y,u}}$. By Lemma 6.4, we just need to prove that all strongly antisymmetric $\mathcal{P}_{\bar{G}, N'}$ -schemes are discrete on \bar{G}_v . Equivalently, we want to prove all strongly antisymmetric $\mathcal{P}_{H', N'}$ -schemes (defined with respect to the action of H' on Γ) are discrete on H'_{v_i} . Note $H' \subseteq H$ where H is a primitive permutation group of almost simple type or diagonal type on Γ . If H is of almost simple type, we have $k(H) = k(\text{soc}(H)) \leq k(G)$ by Theorem 8.6 and similarly $r(H) = r(\text{soc}(H)) \leq r(G)$. It follows from Lemma 8.4 that all strongly antisymmetric $\mathcal{P}_{H', N'}$ -schemes are discrete on H'_{v_i} for sufficiently large $N = \text{poly}(k(G)^{d_{\text{Sym}(k(G))}}, r(G), |S|)$. If H is of diagonal type, then we apply Lemma 8.4 instead to conclude that all strongly antisymmetric $\mathcal{P}_{H', N'}$ -schemes are discrete on H'_{v_i} for sufficiently large $N = \text{poly}(|\Gamma|)$. So $\mathcal{C}|_{G'_{y,u}}$ is discrete on $G'_{y,u,v}$. Therefore $G'_{y,z} g^{-1}$ and $G'_{y,z} g'^{-1}$ are in different blocks of $C_{G'_{y,z}}|_{G'_y}$, as claimed.

Consider an arbitrary block $\{G'_{y,z} g_1^{-1}, \dots, G'_{y,z} g_s^{-1}\}$ of $C_{G'_{y,z}}|_{G'_y}$ of cardinality $s \in \mathbb{N}^+$. By the claim just proved, the elements $g_1 z, \dots, g_s z$ differ from y at distinct coordinates. So $s \leq k$. Then $m(s) = O(\log s) = O(\log k)$ by Theorem 7.1 (see Definition 7.3 for the definition of $m(\cdot)$). Choose the largest $m \in \mathbb{N}$ satisfying

$|G'_y z|^m \leq N$. By definition, the subgroup system $\mathcal{P}|_{G'_y}$ contains the system of stabilizers of depth m over G'_y (with respect to the action of G'_y on $G'_y z$). Lemma 2.7 and Theorem 6.1 then imply the existence of a strongly antisymmetric m -scheme $\Pi = \{P_1, \dots, P_m\}$ on $G'_y z$ such that P_1 has a block of cardinality s . Note $|G'_y z| \leq k|\Gamma|$, $|S| = |\Gamma|^k$, and $k = \log |S| / \log |\Gamma| \leq \log |S|$. Then for sufficiently large $N = |S|^{\Omega(1)}$, we have

$$|G'_y z|^{m(s)} = (k|\Gamma|)^{O(\log k)} \leq N$$

and hence $m \geq m(s)$. Theorem 7.1 then forces $s = 1$. So $\mathcal{C}|_{G'_y}$ is discrete on $G'_{y,z}$, as desired. \square

Choosing $G' = G$ in Lemma 8.13, we have

Corollary 8.3. *Theorem 8.4 holds for finite primitive permutation groups of product type.*

Twisted wreath type. Suppose G is a finite primitive permutation group of twisted wreath type. By Definition 8.7, we may assume $G = B \rtimes P$ acting on $S := G/P$ by left translation, where

- T is a noncyclic finite simple group,
- $P \subseteq \text{Sym}(k)$ is a transitive permutation group on $[k]$ for some integer $k \geq 2$,
- φ is a group homomorphism from P_1 to $\text{Aut}(T)$,
- B is the group $\{f \in \text{Map}(P, T) : f(pq^{-1}) = \varphi^{(q)}(f(p)) \text{ for all } p \in P, q \in P_1\}$ under coordinatewise multiplication, and
- P acts on B via $({}^p f)(x) = f(p^{-1}x)$ for $p, x \in P, f \in B$.

It turns out that G can be embedded in a finite primitive permutation group of product type on S . This is explained in [Pra90, Section 3.6]. We provide a detailed proof of this fact.

Lemma 8.14. *Let G be a finite primitive permutation group of twisted wreath type on $S = G/P$ as above. Then G is permutation isomorphic to a subgroup of a finite primitive permutation group $\text{Hol}(T) \wr P$ of product type on S .*

Proof. Identifying $S = G/P$ with the set B via the bijection $B \rightarrow G/P$ sending $g \in B$ to $gP \in G/P$, we may regard $G = B \rtimes P$ as a permutation group on the set B where $B \subseteq G$ acts on B by left translation and $P \subseteq G$ acts by $({}^p f)(x) = f(p^{-1}x)$ for $p, x \in P, f \in B$. Pick $g_1, \dots, g_k \in P$ such that ${}^{g_i}1 = i \in [k]$. Then g_1, \dots, g_k form a complete set of representatives of P/P_1 . We further regard G as a permutation group on T^k by identifying the set B with T^k via the bijection $B \rightarrow T^k$ sending $f \in B$ to $(f(g_1), \dots, f(g_k)) \in T^k$.

The holomorph $\text{Hol}(T)$ of T is a primitive permutation group of diagonal type on T where the action is defined by ${}^{hg}h' = h^g h'$ for $h, h' \in T$ and $g \in \text{Aut}(T)$ (see Example 8.1 and Lemma 8.3). Denote by G' the wreath product $\text{Hol}(T) \wr P$ acting faithfully on the set T^k by the primitive wreath product action, i.e., $\text{Hol}(T)^k$ acts on T^k coordinatewisely and $P \subseteq \text{Sym}(k)$ permutes the k coordinates. We claim that G is permutation isomorphic to a subgroup of G' on T^k . To see this, note that a permutation $f \in B \subseteq G$ of T^k is the same as the permutation $(f(g_1), \dots, f(g_k)) \in T^k \trianglelefteq \text{Hol}(T)^k \trianglelefteq G'$. Now consider $\pi \in P \subseteq G$ and we show that it is also a permutation in G' . For $i \in [k]$, the permutations $\pi^{-1}g_i$ and $g_{\pi^{-1}i}$ of $[k]$ both send 1 to $\pi^{-1}i$, and hence $\pi^{-1}g_i P_1 = g_{\pi^{-1}i} P_1$. So we can choose $h_1, \dots, h_k \in P_1$ such that $\pi^{-1}g_i = g_{\pi^{-1}i} h_i^{-1}$ holds for all $i \in [k]$. We claim that $\pi \in P \subseteq G$, as a permutation of T^k , equals $(\varphi(h_1), \dots, \varphi(h_k))\pi \in G'$. This is because for $f \in B$, we have

$$\begin{aligned} (\varphi(h_1), \dots, \varphi(h_k))\pi (f(g_1), \dots, f(g_k)) &= (\varphi(h_1), \dots, \varphi(h_k)) \left(f \left(g_{\pi^{-1}1} \right), \dots, f \left(g_{\pi^{-1}k} \right) \right) \\ &= \left(f \left(g_{\pi^{-1}1} h_1^{-1} \right), \dots, f \left(g_{\pi^{-1}k} h_k^{-1} \right) \right) \\ &= (f(\pi^{-1}g_1), \dots, f(\pi^{-1}g_k)) \\ &= ((\pi f)(g_1), \dots, (\pi f)(g_k)) \\ &= \pi (f(g_1), \dots, f(g_k)). \end{aligned}$$

Here π in the last equation acts as an element of G , whereas $(\varphi(h_1), \dots, \varphi(h_k))\pi$ is an element of G' . It follows that $G = B \rtimes P$ is permutation isomorphic to a subgroup of G' on T^k . As G acts primitively on T^k , so does G' . By definition, the group G' is a finite primitive permutation group of product type on T^k . The lemma follows. \square

For the groups $G = T \text{ twr}_\varphi P$ and $G' = \text{Hol}(T) \wr P$ in Lemma 8.14, we have $k(G) = k(G')$ and $r(G) = r(G')$ by Theorem 8.6. Then by Lemma 8.13 and Lemma 8.14, we have

Corollary 8.4. *Theorem 8.4 holds for finite primitive permutation groups of twisted wreath type.*

8.7 Future research

In this section, we suggest some possible directions for future research.

Dependence on classical groups. As we have shown, the running time of the factoring algorithm in this chapter is controlled by the alternating groups and the classical groups among the composition factors of the Galois group. Nevertheless, the exact relation between the running time and the classical groups is not fully investigated. The bound we use for classical simple groups is simply the group order $r(G)$, and a natural problem is to improve this bound. In the case of the natural action of a general linear group $G := \mathrm{GL}_n(q)$ on $S := \mathbb{F}_q^n - \{0\}$, it yields the bound $r(G) = |\mathrm{PSL}_n(q)| = q^{O(n^2)}$. Note that Corollary 3.2 or Corollary 5.2 gives the same bound $|S|^{d(\mathrm{GL}_n(q))} = |S|^{d_{\mathrm{GL}}(n,q)} = q^{O(n^2)}$ if we use the trivial $O(n)$ bound for $d_{\mathrm{GL}}(n, q)$ (see Section 7.4). This observation suggests that proving $d_{\mathrm{GL}}(n, q) = o(n)$ is possibly the first step towards a faster factoring algorithm for classical groups.

Factoring algorithms and \mathcal{P} -schemes for various permutation groups. The main results of this chapter demonstrate that the problem of deterministic polynomial factoring may be much easier when the Galois group has a relatively simple structure. In particular, the results are obtained for Galois groups with restricted composition factors. It is an interesting problem to see if similar results can be obtained for other families of permutation groups under possibly different restrictions.

A related problem is proving the schemes conjectures (Conjecture 6.3) for more general families of permutation groups. As we observed in Section 6.3, proving these conjectures for various permutation groups are intermediate steps towards proving the original schemes conjecture in [IKS09].

Connections with association schemes. Another approach is to exploit the connections between our notion of \mathcal{P} -schemes and association schemes. For example, by drawing connections between m -schemes [IKS09] and association schemes, the work [Aro+14] gave a factoring algorithm that finds a nontrivial factor of a reducible polynomial $f(X) \in \mathbb{F}_q[X]$ of prime degree n in time $\mathrm{poly}(\log q, n^{r+\log \ell})$ under GRH, provided that $n - 1$ has an r -smooth divisor s satisfying $s \geq \sqrt{n/\ell} + 1$.

We have shown that \mathcal{P} -schemes generalize m -schemes, in the sense that an m -scheme is essentially a \mathcal{P} -scheme with \mathcal{P} chosen to be the system of stabilizers of depth m over a multiply transitive group (see Theorem 2.1). Thus it is a curious question if the theory of association schemes can find more applications in deterministic polynomial factoring within our framework of \mathcal{P} -schemes.

Appendix A

A UNIFYING DEFINITION OF \mathcal{P} -SCHEMES

We present an alternative ring-theoretic definition of \mathcal{P} -schemes, such that the three defining properties (compatibility, invariance, regularity) are given in a unifying way.

Ring $\text{Ind}^G K$. Let G be a finite group and K be an arbitrary field of characteristic zero. Define $\text{Ind}^G K$ to be the set of all the functions $\phi : G \rightarrow K$. We make it into a commutative ring by defining addition and multiplication entry-wisely. Let G act on it by $({}^g\phi)(gh) = \phi(h)$, or equivalently

$$({}^g\phi)(h) = \phi(g^{-1}h)$$

for $g, h \in G$ and $\phi \in \text{Ind}^G K$.

For a subgroup $H \subseteq G$, the subring $(\text{Ind}^G K)^H$ of H -invariant elements consists of functions $\phi : G \rightarrow K$ taking a constant value on each right coset in $H \backslash G$. So $(\text{Ind}^G K)^H$ is identified with the commutative ring consisting of all the functions from $H \backslash G$ to K where addition and multiplication are defined entry-wisely.

We define inclusions, conjugations and trace maps between $(\text{Ind}^G K)^H$ for various subgroups $H \subseteq G$:

- (inclusion) for $H \subseteq H' \subseteq G$, the ring $(\text{Ind}^G K)^{H'}$ is a subring of $(\text{Ind}^G K)^H$. Define the map $i_{H,H'} : (\text{Ind}^G K)^{H'} \hookrightarrow (\text{Ind}^G K)^H$ to be the natural inclusion.
- (conjugation) for $g \in G$ and $H' = gHg^{-1}$, define $c_{H,g}^* : (\text{Ind}^G K)^{H'} \rightarrow (\text{Ind}^G K)^H$ to be the map sending ϕ to ${}^{g^{-1}}\phi$.
- (trace map) for $H \subseteq H'$, define $\text{Tr}_{H,H'} : (\text{Ind}^G K)^{H'} \rightarrow (\text{Ind}^G K)^H$ to be the map sending ϕ to $\sum_{gH \in H'/H} {}^g\phi$.

Note that trace maps are indeed well defined: as $\phi \in (\text{Ind}^G K)^{H'}$ is fixed by H , the function ${}^g\phi$ depends only on the left coset gH , and the image $\text{Tr}_{H,H'}(\phi)$ does lie in $(\text{Ind}^G K)^H$, since for $h \in H'$ we have

$${}^h\text{Tr}_{H,H'}(\phi) = \left(\sum_{gH \in H'/H} {}^g\phi \right) = \sum_{gH \in H'/H} {}^{hg}\phi = \sum_{gH \in H'/H} {}^g\phi = \text{Tr}_{H,H'}(\phi).$$

The third equality holds since if g ranges over a complete set of representatives for H'/H , so does hg .

Subring R_P associated with a partition P . For a subgroup $H \subseteq G$ and a partition P of $H \backslash G$, define R_P as the subring of $(\text{Ind}^G K)^H$ consisting of functions $\phi : H \backslash G \rightarrow K$ taking a constant value on each block B of P .

The connection between these subrings and \mathcal{P} -schemes is described by the following theorem.

Theorem A.1. For a \mathcal{P} -collection $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$,

- \mathcal{C} is compatible iff $i_{H,H'}(R_{C_{H'}}) \subseteq R_{C_H}$ holds for all $H, H' \in \mathcal{P}$ with $H \subseteq H'$,
- \mathcal{C} is invariant iff $c_{H,g}^*(R_{C_{H'}}) \subseteq R_{C_H}$ holds for all $H, H' \in \mathcal{P}$, $g \in G$ with $H' = gHg^{-1}$, and
- \mathcal{C} is regular iff $\text{Tr}_{H,H'}(R_{C_H}) \subseteq R_{C_{H'}}$ holds for all $H, H' \in \mathcal{P}$ with $H \subseteq H'$.

Proof. Make every ring $(\text{Ind}^G K)^H$ as well as R_{C_H} into a K -algebra by defining scalar multiplication of K entry-wisely. Note that maps $i_{H,H'}$, $c_{H,g}^*$ and $\text{Tr}_{H,H'}$ are K -linear. For $H \in G$ and $B \in C_H$, define the function $\delta_B : H \backslash G \rightarrow K$ by

$$\delta_B(x) = \begin{cases} 1 & x \in B, \\ 0 & x \notin B. \end{cases}$$

Then R_{C_H} is spanned by the functions δ_B over K where $B \in C_H$. So by K -linearity, we have $i_{H,H'}(R_{C_{H'}}) \subseteq R_{C_H}$ iff $i_{H,H'}(\delta_B) \in R_{C_H}$ for all $B \in C_{H'}$, and similar claims hold for $c_{H,g}^*$ and $\text{Tr}_{H,H'}$.

Suppose \mathcal{C} is compatible. Fix $H, H' \in \mathcal{P}$ with $H \subseteq H'$, and we check $i_{H,H'}(\delta_B) \in R_{C_H}$ for all $B \in C_{H'}$, i.e., the function $i_{H,H'}(\delta_B)$ takes a constant value on each block of C_H for all $B \in C_{H'}$. By definition, its value at $Hh \in H \backslash G$ equals $\delta_B(H'h) = \delta_B(\pi_{H,H'}(Hh))$, which equals one if $\pi_{H,H'}(Hh) \in B$ and zero otherwise. The claim follows by compatibility of \mathcal{C} .

Conversely, assume \mathcal{C} is not compatible, i.e., for some $H, H' \in \mathcal{P}$ with $H \subseteq H'$, $B \in C_H$, $B' \in C_{H'}$ and elements $Hh, Hh' \in B$, we have $H'h = \pi_{H,H'}(Hh) \in B'$ but $H'h' = \pi_{H,H'}(Hh') \notin B'$. We show that $i_{H,H'}(\delta_{B'}) \notin R_{C_H}$. By definition, we have $(i_{H,H'}(\delta_{B'}))(Hh) = \delta_{B'}(H'h) = 1$ but $(i_{H,H'}(\delta_{B'}))(Hh') = \delta_{B'}(H'h') = 0$.

So the value of $i_{H,H'}(\delta_{B'})$ is not a constant on the block B . Therefore $i_{H,H'}(\delta_{B'}) \notin R_{C_H}$.

The proof for invariance is similar. Suppose \mathcal{C} is invariant. Fix $H, H' \in \mathcal{P}$, $g \in G$ with $H' = gHg^{-1}$, and we check $c_{H,g}^*(\delta_B) \in R_{C_H}$ for all $B \in C_{H'}$, i.e., the function $c_{H,g}^*(\delta_B)$ takes a constant value on each block of C_H for all $B \in C_{H'}$. By definition, we have $c_{H,g}^*(\delta_B) = {}^{g^{-1}}\delta_B$ with respect to the action of G on $\text{Ind}^G K$ defined at the beginning, where δ_B is regarded as an element of $\text{Ind}^G K$. Then for $Hh \in H \setminus G$, we have

$$(c_{H,g}^*(\delta_B))(Hh) = ({}^{g^{-1}}\delta_B)(h) = \delta_B(gh) = \delta_B(H'gh) = \delta_B(c_{H,g}(Hh))$$

which equals one if $c_{H,g}(Hh) \in B$ and zero otherwise. The claim follows by invariance of \mathcal{C} .

Conversely, assume \mathcal{C} is not invariant, i.e., for some for $H, H' \in \mathcal{P}$, $g \in G$ with $H' = gHg^{-1}$, $B \in C_H$, $B' \in C_{H'}$ and elements $Hh, Hh' \in B$, we have $H'gh = c_{H,g}(Hh) \in B'$ but $H'gh' = c_{H,g}(Hh') \notin B'$. We show that $c_{H,g}^*(\delta_{B'}) \notin R_{C_H}$. By definition, we have $(c_{H,g}^*(\delta_{B'}))(Hh) = \delta_{B'}(H'gh) = 1$ but $(c_{H,g}^*(\delta_{B'}))(Hh') = \delta_{B'}(H'gh') = 0$. So the value of $c_{H,g}^*(\delta_{B'})$ is not a constant on the block B . Therefore $c_{H,g}^*(\delta_{B'}) \notin R_{C_H}$.

Now suppose \mathcal{C} is regular. Fix $H, H' \in \mathcal{P}$ with $H \subseteq H'$, and we check $\text{Tr}_{H,H'}(\delta_B) \in R_{C_{H'}}$ for all $B \in C_H$, i.e., the map $\text{Tr}_{H,H'}(\delta_B)$ takes a constant value on each block of $C_{H'}$ for all $B \in C_H$. By definition, we have $\text{Tr}_{H,H'}(\delta_B) = \sum_{gH \in H'/H} {}^g\delta_B$ with respect to the action of G on $\text{Ind}^G K$ defined at the beginning, where δ_B is regarded as an element of $\text{Ind}^G K$. Then for $H'h \in H' \setminus G$, we have

$$\begin{aligned} (\text{Tr}_{H,H'}(\delta_B))(H'h) &= \sum_{gH \in H'/H} ({}^g\delta_B)(h) = \sum_{Hg \in H \setminus H'} ({}^{g^{-1}}\delta_B)(h) = \sum_{Hg \in H \setminus H'} \delta_B(gh) \\ &= |\{Hg \in H \setminus H' : Hgh \in B\}| \\ &= |\{Hg \in H \setminus G : Hgh \in B, \pi_{H,H'}(Hgh) = H'h\}| \\ &= |\{Hg \in H \setminus G : Hg \in B, \pi_{H,H'}(Hg) = H'h\}|, \end{aligned}$$

which counts the number of elements in B mapped to $H'h$ by $\pi_{H,H'}$. By regularity, this value is a constant when $H'h$ ranges over a block of $C_{H'}$, as desired.

Conversely, assume \mathcal{C} is not regular, i.e., for some $H, H' \in \mathcal{P}$ with $H \subseteq H'$, $B \in C_H$, $B' \in C_{H'}$, and $H'h, H'h' \in B'$, the number of elements in B mapped to $H'h$ is different from the number of those mapped to $H'h'$. As shown in

the previous paragraph, these two numbers are precisely $(\mathrm{Tr}_{H,H'}(\delta_B))(H'h)$ and $(\mathrm{Tr}_{H,H'}(\delta_B))(H'h')$ respectively. So the value of $\mathrm{Tr}_{H,H'}(\delta_B)$ is not a constant on the block B' . Therefore $\mathrm{Tr}_{H,H'}(\delta_B) \notin R_{C_{H'}}$. \square

By Theorem A.1, we have the following alternative definition for \mathcal{P} -schemes, which is equivalent to the original one (Definition 2.4).

Definition A.1 (\mathcal{P} -scheme, alternative definition). *A \mathcal{P} -collection $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ is a \mathcal{P} -scheme if it has the following properties:*

- (compatibility) $i_{H,H'}(R_{C_{H'}}) \subseteq R_{C_H}$ holds for all $H, H' \in \mathcal{P}$ with $H \subseteq H'$.
- (invariance) $c_{H,g}^*(R_{C_{gHg^{-1}}}) \subseteq R_{C_H}$ holds for all $H \in \mathcal{P}$ and $g \in G$.
- (regularity) $\mathrm{Tr}_{H,H'}(R_{C_H}) \subseteq R_{C_{H'}}$ holds for all $H, H' \in \mathcal{P}$ with $H \subseteq H'$.

Remark. The reader familiar with the notion of *affine schemes* [Mum99] may recognize the right coset space $H \backslash G$ as (the underlying set of) the affine scheme associated with the commutative ring $(\mathrm{Ind}^G K)^H$. More generally, each a partition P of $H \backslash G$ determines a quotient set of $H \backslash G$ which is (the underlying set of) the affine scheme associated with the subring R_P . It is known that the language of affine schemes and that of commutative rings are equivalent.¹ Theorem A.1 is a manifestation of this equivalence.

Therefore in principle, statements and proofs about \mathcal{P} -schemes may be carried out either set-theoretically or ring-theoretically. We stick to the more elementary set-theoretic language in this thesis.

¹Formally, this is known as the fact that the category of affine schemes is anti-equivalent to the category of commutative rings. See, e.g., [Mum99, Section II.2, Corollary 1].

Appendix B

PROOFS OMITTED FROM CHAPTER III

This chapter contains proofs that are omitted from Chapter 3.

Lemma 3.5. *The partitions $P(I)$ and the idempotent decompositions $I(P)$ are well defined. And for any idempotent decomposition I of $\bar{\mathcal{O}}_K$, the idempotents $\delta \in I$ correspond one-to-one to the blocks of $P(I)$ via the map $\delta \mapsto B_\delta := \{Hg \in H \setminus G : g^{-1}(i_{K,L}(\delta)) \equiv 1 \pmod{\bar{\mathfrak{Q}}_0}\}$ with the inverse map $B \mapsto \delta_B$.*

Proof. We first show that $P(I)$ and $I(P)$ are well defined. For $P(I)$ we note that $g^{-1}(i_{K,L}(\delta))$ depends only on the coset Hg , since $i_{K,L}(\delta) \in i_{K,L}(\bar{\mathcal{O}}_K)$ is fixed by H . The relation $g^{-1}(i_{K,L}(\delta)) \equiv g'^{-1}(i_{K,L}(\delta)) \pmod{\bar{\mathfrak{Q}}_0}$ for all $\delta \in I$ is obviously an equivalence relation on $H \setminus G$, and hence defines a partition of $H \setminus G$.

For $I(P)$, we fix $B \subseteq H \setminus G$ and show that $t := \sum_{g \in G: Hg \in B} {}^g\delta_{\bar{\mathfrak{Q}}_0}$ does lie in the image of $i_{K,L}$ so that $\delta_B = i_{K,L}^{-1}(t)$ is well defined. By Corollary 3.1, each coset $x = Hg$ corresponds to a maximal ideal $\mathfrak{P}_x := ({}^g\bar{\mathfrak{Q}}_0 \cap \mathcal{O}_K)/p\mathcal{O}_K$ of $\bar{\mathcal{O}}_K$. By Lemma 3.3, there exists a unique idempotent δ of $\bar{\mathcal{O}}_K$ satisfying $\delta \equiv 1 \pmod{\mathfrak{P}_x}$ for $x \in B$, and $\delta \equiv 0 \pmod{\mathfrak{P}_x}$ for $x \notin B$. It follows that for $g \in G$, the residue of $i_{K,L}(\delta)$ modulo ${}^g\bar{\mathfrak{Q}}_0$ equals one if $Hg \in B$ and zero otherwise. The same holds for t by definition: for $g \in G$, the residue of t modulo ${}^g\bar{\mathfrak{Q}}_0$ equals one if $Hg \in B$ and zero otherwise. As all the maximal ideals of the semisimple ring $\bar{\mathcal{O}}_L$ have the form ${}^g\bar{\mathfrak{Q}}_0$ where $g \in G$, we have $t = i_{K,L}(\delta)$, as desired. Furthermore, by choosing $B = H \setminus G$ and $t = i_{K,L}(1) = 1$, we see that $\sum_{g \in G} {}^g\delta_{\bar{\mathfrak{Q}}_0} = 1$. It follows that $I(P)$ is a well defined idempotent decomposition of $\bar{\mathcal{O}}_K$.

For the second claim, we first check that the sets B_δ form a partition of $H \setminus G$ and the map $\delta \mapsto B_\delta$ is injective. To see this, note that if an element Hg lies in both B_δ and $B_{\delta'}$ for distinct $\delta, \delta' \in I$, then $g^{-1}(i_{K,L}(\delta)) \equiv g^{-1}(i_{K,L}(\delta')) \equiv 1 \pmod{\bar{\mathfrak{Q}}_0}$ by definition. But then $g^{-1}(i_{K,L}(\delta\delta')) \equiv 1 \pmod{\bar{\mathfrak{Q}}_0}$, contradicting the fact that $\delta\delta' = 0$. So the sets B_δ are disjoint and the map $\delta \mapsto B_\delta$ is injective. Furthermore, each $Hg \in H \setminus G$ lies in at least one set B_δ since

$$\sum_{\delta \in I} g^{-1}(i_{K,L}(\delta)) \equiv \left(i_{K,L} \left(\sum_{\delta \in I} \delta \right) \right) \equiv 1 \pmod{\bar{\mathfrak{Q}}_0}. \quad (\text{B.1})$$

So the sets B_δ form a partition of $H \setminus G$.

It remains to show that the partition formed by the sets B_δ is $P(I)$. Assume two elements Hg, Hg' are in the same set B_{δ_0} for some $\delta_0 \in I$, then $g^{-1}(i_{K,L}(\delta_0)) \equiv g'^{-1}(i_{K,L}(\delta_0)) \equiv 1 \pmod{\bar{\mathfrak{Q}}_0}$ by definition. By (B.1), we have $g^{-1}(i_{K,L}(\delta)) \equiv g'^{-1}(i_{K,L}(\delta)) \equiv 0 \pmod{\bar{\mathfrak{Q}}_0}$ for $\delta \in I - \{\delta_0\}$. So Hg, Hg' are in the same block of $P(I)$.

Conversely, assume two elements Hg, Hg' are in the same block of $P(I)$. So $g^{-1}(i_{K,L}(\delta)) \equiv g'^{-1}(i_{K,L}(\delta)) \pmod{\bar{\mathfrak{Q}}_0}$ for all $\delta \in I$. And by (B.1), we have $g^{-1}(i_{K,L}(\delta_0)) \equiv g'^{-1}(i_{K,L}(\delta_0)) \equiv 1 \pmod{\bar{\mathfrak{Q}}_0}$ for some $\delta_0 \in I$. Then $Hg, Hg' \in B_{\delta_0}$. So the partition formed by the sets B_δ is exactly $P(I)$. \square

Lemma 3.6. *The map $I \mapsto P(I)$ is a one-to-one correspondence between the idempotent decompositions of $\bar{\mathcal{O}}_K$ and the partitions of $H \setminus G$, with the inverse map $P \mapsto I(P)$.*

Proof. It suffices to show that the map $B \mapsto \delta_B$ is a one-to-one correspondence between the subsets of $H \setminus G$ and the idempotents of $\bar{\mathcal{O}}_K$, with the inverse map $\delta \mapsto B_\delta$. Fix $B \subseteq H \setminus G$ and let $\delta = \delta_B = i_{K,L}^{-1} \left(\sum_{g \in G: Hg \in B} g \delta_{\bar{\mathfrak{Q}}_0} \right)$. We verify that $B_\delta = B$: for $Hh \in H \setminus G$, we have

$$h^{-1}(i_{K,L}(\delta)) = \sum_{g \in G: Hg \in B} h^{-1}g \delta_{\bar{\mathfrak{Q}}_0}.$$

Note that the residue of $h^{-1}g \delta_{\bar{\mathfrak{Q}}_0}$ modulo $\bar{\mathfrak{Q}}_0$ equals one if $h = g$, and zero otherwise. So the residue of $h^{-1}(i_{K,L}(\delta))$ modulo $\bar{\mathfrak{Q}}_0$ equals one if $Hh \in B$ and zero otherwise. It follows by definition that $B_\delta = B$.

It then suffices to show that the map $\delta \mapsto B_\delta$ is injective. Consider two distinct idempotents δ, δ' of $\bar{\mathcal{O}}_K$. As $\delta \neq \delta'$, there exists a maximal ideal ${}^g\bar{\mathfrak{Q}}_0$ of $\bar{\mathcal{O}}_L$ for which $i_{K,L}(\delta) \not\equiv i_{K,L}(\delta') \pmod{{}^g\bar{\mathfrak{Q}}_0}$ holds. Equivalently, we have $g^{-1}(i_{K,L}(\delta)) \not\equiv g^{-1}(i_{K,L}(\delta')) \pmod{\bar{\mathfrak{Q}}_0}$. Then exactly one of $g^{-1}(i_{K,L}(\delta)), g^{-1}(i_{K,L}(\delta'))$ equals one modulo $\bar{\mathfrak{Q}}_0$. So Hg is in one of $B_\delta, B_{\delta'}$ but not in the other. Therefore $B_\delta \neq B_{\delta'}$. \square

Lemma 3.9. *There exists a polynomial-time algorithm `ComputeResidue` that takes the following data as the input*

- a number fields K , a prime number p , and $\alpha \in \mathcal{O}_K$ given as an element of K ,

- the outputs of `ComputeQuotientRing` (see Lemma 3.8) on the inputs (K, p) , i.e., the quotient ring $\bar{\mathcal{O}}_K$, a maximal p -orders \mathcal{O}'_K , the inclusion $\mathcal{O}'_K \hookrightarrow K$, and the quotient map $\mathcal{O}'_K \rightarrow \bar{\mathcal{O}}_K$,

and computes $\alpha + p\mathcal{O}_K \in \bar{\mathcal{O}}_K$.

Proof. Let $d = [K : \mathbb{Q}]$. Suppose the structure constants of K and \mathcal{O}'_K are given in the \mathbb{Q} -basis B of K and the \mathbb{Z} -basis $B' = \{x_1, \dots, x_d\}$ of \mathcal{O}'_K respectively. Then we may assume the structure constants of $\bar{\mathcal{O}}_K$ is given in the \mathbb{F}_p -basis $\{x_1 + p\mathcal{O}_K, \dots, x_d + p\mathcal{O}_K\}$ of $\bar{\mathcal{O}}_K$. The goal is computing the constants $c_1, \dots, c_d \in \mathbb{F}_p$ determined by

$$\alpha + p\mathcal{O}_K = \sum_{i=1}^d c_i(x_i + p\mathcal{O}_K). \quad (\text{B.2})$$

Note that B' is also a \mathbb{Q} -basis of K . The change-of-basis matrix M from B' to B is given by the inclusion $\mathcal{O}'_K \hookrightarrow K$, whose entries are rational numbers of polynomial size. So the entries of M^{-1} are also rational numbers of polynomial size. We apply M^{-1} and write α in the basis B' :

$$\alpha = \sum_{i=1}^d r_i x_i, \quad r_i \in \mathbb{Q}.$$

For $i \in [d]$, write r_i in the form a_i/b_i where a_i, b_i are coprime integers and $b_i > 0$. Let m be the least common multiple of all the denominators b_i . Then we have $m\alpha = \sum_{i=1}^d mr_i x_i$ with the coefficients $mr_i \in \mathbb{Z}$. So $m\alpha \in \mathcal{O}'_K \subseteq \mathcal{O}_K$. Passing to the quotient ring $\bar{\mathcal{O}}_K$, we obtain

$$m\alpha + p\mathcal{O}_K = \sum_{i=1}^d c'_i(x_i + p\mathcal{O}_K), \quad c'_i = mr_i \bmod p \in \mathbb{F}_p.$$

Suppose $m = p^e m'$ where $e \in \mathbb{N}$, $m' \in \mathbb{Z}$ and $p \nmid m'$. We claim $e = 0$. Assume to the contrary that $e > 0$. For some $i_0 \in [d]$, we have $p^e | b_{i_0}$ but $p^{e+1} \nmid b_{i_0}$. Then $p \nmid a_{i_0}$ since a_{i_0}, b_{i_0} are coprime. So $p \nmid mr_{i_0}$. Then $c'_{i_0} \neq 0$ and hence $m\alpha + p\mathcal{O}_K \neq 0$. But as $\alpha + p\mathcal{O}_K \in \bar{\mathcal{O}}_K$, we have $m\alpha + p\mathcal{O}_K \in m\bar{\mathcal{O}}_K = 0$, which is a contradiction. So $e = 0$ and $p \nmid m$. Let s be the multiplicative inverse of $m \bmod p \in \mathbb{F}_p$. We compute s and let $c_i = sc'_i$ for $i \in [d]$, which satisfy (B.2). \square

Lemma 3.11. *There exists a polynomial-time algorithm `ComputeRingHom` that takes the following data as the input*

- number fields K, K' , an embedding $\phi : K \rightarrow K'$, and a prime number p ,
- the outputs of `ComputeQuotientRing` (see Lemma 3.8) on the inputs (K, p) and (K', p) respectively,¹

and computes the ring homomorphism $\bar{\phi} : \bar{\mathcal{O}}_K \rightarrow \bar{\mathcal{O}}_{K'}$ induced from ϕ .

Proof. Let $d = [K : \mathbb{Q}]$. Suppose the structure constants of \mathcal{O}'_K is given in the \mathbb{Z} -bases $\{x_1, \dots, x_d\}$ of \mathcal{O}'_K . Then we may assume the structure constants of $\bar{\mathcal{O}}_K$ is given in the \mathbb{F}_p -bases $\{x_1 + p\mathcal{O}, \dots, x_d + p\mathcal{O}\}$ of $\bar{\mathcal{O}}_K$.

For $i \in [d]$, we need to compute $\bar{\phi}(x_i + p\mathcal{O}_K) \in \bar{\mathcal{O}}_{K'}$. Note that $\bar{\phi}(x_i + p\mathcal{O}_K) = \phi(x_i) + p\mathcal{O}_{K'}$. First compute $\phi(x_i) \in K'$ using the inclusion $\mathcal{O}'_K \hookrightarrow K$ and the embedding $\phi : K \rightarrow K'$ given in the input. Here $\phi(x_i)$ is actually in $\mathcal{O}_{K'}$ since $x_i \in \mathcal{O}'_K \subseteq \mathcal{O}_K$. Use the algorithm `ComputeResidue` to compute $\phi(x_i) + p\mathcal{O}_{K'} \in \bar{\mathcal{O}}_{K'}$, and we are done. \square

Lemma 3.16 ([IKS09; Iva+12]). *There exists an algorithm `FreeModuleTest` that given a semisimple \mathbb{F}_p -algebra A and a finitely generated A -module M , returns a zero divisor a of A in polynomial time, such that a is zero only if M is a free A -module.*

Proof. We maintain a submodule N of M that is free over A . Initially N equals $\{0\}$ and we iteratively enlarge it. Each time we pick $x \in M - N$ and check if the sum $N + Ax$ is a direct sum, i.e., if $N \cap Ax = \{0\}$. If so, we replace N with $N + Ax$. Otherwise we find a nonzero element $y \in N \cap Ax$ and $a \in A$ satisfying $y = ax$, and return a . Note that in the latter case, the element a is indeed a zero divisor: otherwise a would be invertible and hence $x = a^{-1}y$ is in N , contradicting the assumption $x \notin N$.

If N eventually becomes M , we conclude that M is free over A , in which case we return zero. The algorithm clearly runs in polynomial time. \square

Lemma 3.17. *There exists an algorithm `SplitByZeroDivisor` that given*

- a semisimple \mathbb{F}_p -algebra R , an idempotent decomposition I of R , and an idempotent $\gamma \in I$,

¹That is, the quotient rings $\bar{\mathcal{O}}_K, \bar{\mathcal{O}}_{K'}$, the maximal p -orders $\mathcal{O}'_K, \mathcal{O}'_{K'}$, the inclusions $\mathcal{O}'_K \hookrightarrow K, \mathcal{O}'_{K'} \hookrightarrow K'$, and the quotient maps $\mathcal{O}'_K \rightarrow \bar{\mathcal{O}}_K, \mathcal{O}'_{K'} \rightarrow \bar{\mathcal{O}}_{K'}$.

- the ring $\bar{R} := R/(1 - \gamma)$, the quotient map $\pi : R \rightarrow \bar{R}$, and a zero divisor $a \neq 0$ of \bar{R} ,

replaces $\gamma \in I$ with two nonzero idempotents γ_1, γ_2 satisfying $\gamma = \gamma_1 + \gamma_2$ in polynomial time.

Proof. We pick an element $\tilde{a} \in R$ lifting a , i.e., $\pi(\tilde{a}) = a$. Compute the ideal (\tilde{a}) of R generated by \tilde{a} . As R is semisimple, we have $(\tilde{a}) = (\gamma')$ for some idempotent γ' of R . Compute γ' by solving a system of linear equations using the fact that γ' is the unique element in (\tilde{a}) satisfying $\gamma'x = x$ for all $x \in (\tilde{a})$. Finally we replace γ with $\gamma'\gamma$ and $(1 - \gamma')\gamma$. It remains to show that $\gamma'\gamma \notin \{0, \gamma\}$.

Note that $\pi(\gamma') \in \bar{R}$ generates the ideal (a) of \bar{R} , and hence $\pi(\gamma')$ is also a nonzero zero divisor of \bar{R} . But $\pi(\gamma') = \gamma' + (1 - \gamma) = \gamma'\gamma + (1 - \gamma)$. So $\gamma'\gamma \neq 0$. It also follows that $\gamma'\gamma \neq \gamma$ since otherwise we would have $\pi(\gamma') = \gamma'\gamma + (1 - \gamma) = \gamma + (1 - \gamma) = 1 + (1 - \gamma)$, which is the unity of \bar{R} and not a zero divisor. \square

Lemma 3.18 ([Rón92]). *Under GRH, there exists an algorithm Automorphism that given a ring A isomorphic to a finite product of \mathbb{F}_p and a nontrivial ring automorphism σ of A , returns a zero divisor $a \neq 0$ of A in polynomial time.*

To prove Lemma 3.18, we need the following lemma.

Lemma B.1 ([Rón92; Iva+12]). *There exists an algorithm IteratedExp that, given a semisimple \mathbb{F}_p -algebra A , a prime number $\ell \neq p$, and elements x, y in the multiplicative group A^\times of order n_x and n_y respectively such that n_x, n_y are powers of ℓ and $n_x \geq n_y$, returns a zero divisor of the form $x^k - y \in A$, $k \in \mathbb{N}$, in time polynomial in $\log |A|$ and ℓ . In particular, zero is returned only if y is a power of x .*

Proof. The algorithm is as follows: try to find $k \in \{0, \dots, \ell - 1\}$ such that $x^k - y$ is a zero divisor. If such an integer k is found, simply return $x^k - y$. Otherwise raise x to its ℓ th power and repeat.

To analyze the algorithm, note that there exists a maximal ideal \mathfrak{m} of A such that the order of $x + \mathfrak{m} \in (A/\mathfrak{m})^\times$ is n_x , and the order of $y + \mathfrak{m} \in (A/\mathfrak{m})^\times$, which we denote by n'_y , divides n_y . Then $x^{n_x/n'_y} + \mathfrak{m}$ and $y + \mathfrak{m}$ are both primitive n'_y -th roots of unity in $(A/\mathfrak{m})^\times \cong \mathbb{F}_p$. Then there exists $k \in \{0, \dots, \ell - 1\}$ such that $x^{kn_x/n'_y} - y$ is in \mathfrak{m} and hence is a zero divisor. Such a zero divisor is guaranteed to be found when x is raised to x^{n_x/n'_y} (or earlier). \square

Proof of Lemma 3.18. For $x, y \in A$ linearly independent over \mathbb{F}_p , at least one element in the set $\{y - cx : c \in \mathbb{F}_p\}$ is a nonzero zero divisor. If $p \leq \dim_{\mathbb{F}_p} A$, we can find such an element in polynomial time by choosing x, y and enumerating c . So assume $p > \dim_{\mathbb{F}_p} A$. In this case, the pseudocode of the algorithm is given in Algorithm 17. Here id denotes the identity map on A .

Algorithm 17 Automorphism

Input: ring A isomorphic to a finite product of \mathbb{F}_p , automorphism $\sigma \neq \text{id}$ of A

Output: zero divisor $a \neq 0$ of A

```

1:  $n \leftarrow 1$ 
2: repeat
3:   find  $z \in A$  satisfying  $\sigma^n(z) \neq z$ 
4:   if  $\sigma^n(z) - z$  is a zero divisor of  $A$  then
5:     return  $\sigma^n(z) - z$ 
6:   until  $\sigma^n = \text{id}$ 
7: compute the least prime factor  $\ell$  of  $n$ 
8:  $\sigma \leftarrow \sigma^{n/\ell}$ 
9: compute  $\mathbb{F}_{p^d}$ , where  $d$  is the smallest positive integer satisfying  $\ell | p^d - 1$ 
10: compute  $A \otimes \mathbb{F}_{p^d}$  and the inclusion  $i : A \hookrightarrow A \otimes \mathbb{F}_{p^d}$  sending  $t \in A$  to  $t \otimes 1$ 
11: compute the automorphism  $\sigma \otimes 1$  of  $A \otimes \mathbb{F}_{p^d}$  sending  $t \otimes u \in A$  to  $\sigma(t) \otimes u$ 
12: pick an  $\ell$ th power non-residue  $\gamma$  of  $\mathbb{F}_{p^d}$ 
13:  $\xi \leftarrow \gamma^{(p^d-1)/\ell}$ 
14: compute a nonzero element  $x \in A \otimes \mathbb{F}_{p^d}$  satisfying  $(\sigma \otimes 1)(x) = \xi x$ 
15: if  $x$  is a not zero divisor of  $A \otimes \mathbb{F}_{p^d}$  then
16:    $k \leftarrow$  the largest factor of  $p^d - 1$  coprime to  $\ell$ 
17:   call IteratedExp with the input  $A \otimes \mathbb{F}_{p^d}$ ,  $\ell$ ,  $\gamma^k$ , and  $x^k$  to obtain a zero
     divisor  $b$  of  $A \otimes \mathbb{F}_{p^d}$ 
18:    $x \leftarrow b$ 
19: choose  $a \in A - \{0\}$  such that  $i(a)$  is in the ideal  $(x)$  of  $A \otimes \mathbb{F}_{p^d}$ 
20: return  $a$ 

```

The loop in Lines 2–6 of the algorithm computes the powers σ^n of σ for $n = 1, 2, \dots$ and tries to find $z \in A$ satisfying $\sigma^n(z) \neq z$. The loop exits either when such an element z is found, or when the condition $\sigma^n = \text{id}$ is satisfied. In the former case, the algorithm returns the zero divisor $\sigma^n(z) - z \neq 0$, and in the latter case, the algorithm proceeds. Note that initially $n = 1$ and we have $\sigma \neq \text{id}$ by assumption.

By assumption, we may identify A with a product $\prod_{i=1}^m \mathbb{F}_p$ where $m = \dim_{\mathbb{F}_p} A$. For $i \in [m]$, let δ_i be the element of A whose i th coordinate is one and the other components are zero. So $\delta_1, \dots, \delta_m$ are the primitive idempotents of A . The automorphism σ of A permutes these primitive idempotents, i.e., it is associated with a permutation π of $[m]$ such that $\sigma(\delta_i) = \delta_{\pi(i)}$ for $i \in [m]$. By \mathbb{F}_p -linearity of σ (which is automatic since \mathbb{F}_p is a prime field), we know σ sends $(x_1, \dots, x_m) \in A$ to $(x_{\pi^{-1}(1)}, \dots, x_{\pi^{-1}(m)})$.

Let H be the cyclic group generated by π , and it acts on $[m]$. Assume the H -orbits of $[m]$ do not have the same cardinality. We claim that in this case a zero divisor $\sigma^n(z) - z \neq 0$ is returned at Line 5 for some $n \leq m$. To see this, suppose O_1 and O_2 are two H -orbits of distinct cardinalities $n_1, n_2 \leq m$ respectively. We may assume $n_1 \leq n_2$. Then σ^{n_1} fixes all elements in O_1 but not all in O_2 . So $\sigma^{n_1} \neq \text{id}$. If the loop returns a (nonzero) zero divisor at Line 5 in the n th iteration for some $n < n_1$ then we are done. Otherwise, an element z satisfying $\sigma^{n_1}(z) - z \neq 0$ is found at Line 3 in the n_1 th iteration. Note that for any $i \in O_1$, the i th coordinate of $\sigma^{n_1}(z) - z$ is a zero, and hence $\sigma^{n_1}(z) - z$ is annihilated by δ_i . It follows that $\sigma^{n_1}(z) - z$ is a zero divisor and is returned at Line 5.

So assume all the H -orbits of $[m]$ have the same cardinality and the algorithm reaches Line 7. Then the order of σ equals n . Line 8 replaces σ with its (n/ℓ) th power where ℓ is the least prime factor of n . Then the order of σ becomes the prime number ℓ . Note that $\ell < p$ since $n \leq m < p$.

At Line 9, we compute the finite field \mathbb{F}_{p^d} where d is the smallest positive integer satisfying $\ell | p^d - 1$. Equivalently, the integer d is the (multiplicative) order of p in the group $(\mathbb{Z}/\ell\mathbb{Z})^\times$. So we have $d \leq |(\mathbb{Z}/\ell\mathbb{Z})^\times| = \ell - 1$. Under GRH (or Hypothesis $(*)$ in the introduction), the field \mathbb{F}_{p^d} can be computed in deterministic polynomial time. It is the smallest extension of \mathbb{F}_p containing the primitive ℓ th roots of unity.

At Line 10, we compute the \mathbb{F}_{p^d} -algebra $A \otimes \mathbb{F}_{p^d}$ (where the tensor product is taken over \mathbb{F}_p) and the inclusion $i : A \hookrightarrow \mathbb{F}_{p^d}$ sending $t \in A$ to $t \otimes 1$. Suppose $\{b_1, \dots, b_m\}$ is an \mathbb{F}_p -basis of A and $b_i b_j = \sum_{k=1}^m c_{ijk} b_k$ where $c_{ijk} \in \mathbb{F}_p$, then $A \otimes \mathbb{F}_{p^d}$ can be defined as an \mathbb{F}_{p^d} -algebra in the \mathbb{F}_{p^d} -basis $\{b_1 \otimes 1, \dots, b_m \otimes 1\}$ satisfying $(b_i \otimes 1)(b_j \otimes 1) = \sum_{k=1}^m c_{ijk} (b_k \otimes 1)$. It follows from the universal property of tensor products that this definition does not depend on the choice of the basis. See, e.g., [AM69]. In particular, identify A with $\prod_{i=1}^m \mathbb{F}_p$ and then $A \otimes \mathbb{F}_{p^d}$ is simply $\prod_{i=1}^m \mathbb{F}_{p^d}$.

At Line 11, we compute \mathbb{F}_{p^d} -linear automorphism $\sigma \otimes 1$ of $A \otimes \mathbb{F}_{p^d}$ sending $t \otimes u \in A$ to $\sigma(t) \otimes u$. It follows from the universal property of tensor products that such an automorphism exists and is unique. At Line 12, we pick an ℓ th power non-residue γ of \mathbb{F}_{p^d} , which be done in deterministic polynomial time under GRH (or Hypothesis $(*)$ in the introduction). Then at Line 13, we compute $\xi = \gamma^{(p^d-1)/\ell}$, which is a primitive ℓ th root of unity.

At Line 14, we compute a nonzero element $x \in A \otimes \mathbb{F}_{p^d}$ satisfying $(\sigma \otimes 1)(x) = \xi x$. We claim that such an element x exists. To see this, note that as σ has order ℓ , the permutation π of $[m]$ associated with σ has an ℓ -cycle $(i_1 i_2 \cdots i_\ell)$. Then we can choose x to be the element in $A \otimes \mathbb{F}_{p^d} = \prod_{i=1}^m \mathbb{F}_{p^d}$ whose i_j th coordinate is ξ^{-j} for $j \in [\ell]$ and remaining coordinates are zero.

If the element x is a zero divisor of $A \otimes \mathbb{F}_{p^d}$, the preimage of the ideal (x) of $A \otimes \mathbb{F}_{p^d}$ in A under the map i is strictly between $\{0\}$ and A . In this case, we compute a nonzero element a in it (or equivalently, an element a satisfying $i(a) \in (x)$) at Line 19 and return it. Note that a is guaranteed to be a zero divisor of A .

On the other hand, if x is not a zero divisor of $A \otimes \mathbb{F}_{p^d}$, we replace it with a zero divisor $b \neq 0$ in Lines 16–18: suppose $p^d - 1 = k\ell^e$ where k is coprime to ℓ . We compute k at Line 16. As $\gamma \in \mathbb{F}_{p^d}$ is an ℓ th power non-residue, the order of γ^k is $\ell^e = (p^d - 1)/k$. As x is not a zero divisor, we have $x^k \in (A \otimes \mathbb{F}_{p^d})^\times$ and its order divides ℓ^e . Also note that $\sigma \otimes 1$ fixes γ^k (by \mathbb{F}_{p^d} -linearity) and sends x^k to $\xi^k x^k \neq x^k$. So x^k is not a power of γ^k . By Lemma B.1, a zero divisor $b \neq 0$ of $A \otimes \mathbb{F}_{p^d}$ is obtained at Line 17 and we assign its value to x . Then we obtain the zero divisor $a \neq 0$ of A and return it at Line 19 as before. \square

Appendix C

PROOFS OMITTED FROM CHAPTER V

This chapter contains proofs that are omitted from Chapter 5.

Lemma 5.1. *There exists a polynomial-time algorithm that given p and a polynomial $\tilde{f}(X) \in A_0[X]$ satisfying $\tilde{\psi}_0(\tilde{f}) \neq 0$, computes an integer D satisfying $D \equiv 1 \pmod{p}$ and a factorization of $D \cdot \tilde{f}$ into irreducible factors \tilde{f}_i over K_0 . Furthermore all of the factors $\tilde{f}_i(X)$ are in $A_0[X]$.*

Proof. Factorize \tilde{f} into irreducible factors g_1, \dots, g_k over K_0 using polynomial factoring algorithms for number fields [Len83; Lan85]. Note that coefficients of each factor g_i lie in $K_0 = \mathbb{Q}[Y]/(\tilde{h}(Y))$ but not necessarily in $A_0 = \mathbb{Z}[Y]/(\tilde{h}(Y))$. Here a coefficient $\alpha \in K_0$ is represented by a unique polynomial $r_\alpha(Y) \in \mathbb{Q}[Y]$ of degree at most $\deg(\tilde{h}) - 1$ such that $\alpha = r_\alpha(Y) + (\tilde{h}(Y))$. And $\alpha \in A_0$ holds iff the coefficients of $r_\alpha(Y)$ are all integers.

For each factor g_i , use r_α , where α ranges over coefficients of g_i , to compute the smallest $e_i \in \mathbb{Z}$ and $D_i \in \mathbb{N}^+$ coprime to p such that all the coefficients of $p^{e_i}D_i g_i$ are in A_0 . Compute an integer $D \in \mathbb{N}^+$ such that D is a multiple of $\prod_{i=1}^k D_i$ and $D \equiv 1 \pmod{p}$. Compute $\tilde{f}_i := p^{e_i}D_i g_i$ for $i = 2, \dots, k$ and $\tilde{f}_1 = (p^{e_1}D/\prod_{i=2}^k D_i)g_1$. Then the polynomials $\tilde{f}_i(X)$ are all in $A_0[X]$.

It remains to show that the product of \tilde{f}_i equals $D \cdot \tilde{f}$, which reduces to proving $\sum_{i=1}^k e_i = 0$. Note that for all $i \in [k]$, the polynomial $p^{e_i}D_i g_i(X)$ is in $A_0[X]$ but not in $pA_0[X]$, since otherwise we may replace e_i with $e_i - 1$, contradicting the minimality of e_i . The ideal $pA_0[X]$ is a prime ideal of $A_0[X]$, since $A_0[X]/pA_0[X] \cong \mathbb{F}_q[X]$ is an integral domain. Therefore

$$\prod_{i=1}^k p^{e_i}D_i g_i(X) = \left(\prod_{i=1}^k p^{e_i} \right) \cdot \left(\prod_{i=1}^k D_i \right) \cdot \tilde{f}(X)$$

is not in $pA_0[X]$ either. So $\sum_{i=1}^k e_i \leq 0$. Rewrite the equation above as

$$\left(\prod_{i=1}^k p^{-e_i} \right) \cdot \left(\prod_{i=1}^k p^{e_i}D_i g_i(X) \right) = \left(\prod_{i=1}^k D_i \right) \cdot \tilde{f}(X).$$

As $\tilde{\psi}_0(\tilde{f}) \neq 0$, we have $\tilde{f}(X) \notin pA_0[X]$. And the integers D_i are coprime to p and hence not in $pA_0[X]$ either. The equation above then implies $\sum_{i=1}^k e_i \geq 0$. \square

Remark. An alternative way of proving $\sum_{i=1}^k e_i = 0$ is to consider the localization of A_0 at the prime ideal pA_0 and apply Gauss Lemma (see [Lan02, Section IV.2]). We leave the details to the reader.

Lemma 5.4. *The partitions $P(I)$ and the idempotent decompositions $I(P)$ are well defined. And for any idempotent decomposition I of $\bar{\mathcal{O}}_K$, the idempotents $\delta \in I$ correspond one-to-one to the blocks of $P(I)$ via the map*

$$\delta \mapsto B_\delta := \{Hg\mathcal{D}_{\Omega_0} \in H \backslash G / \mathcal{D}_{\Omega_0} : g^{-1}(i_{K,L}(\delta)) \equiv 1 \pmod{\bar{\mathfrak{Q}}_0}\}$$

with the inverse map $B \mapsto \delta_B$.

Proof. We first show that $P(I)$ and $I(P)$ are well defined. For $P(I)$ we note that $g^{-1}(i_{K,L}(\delta)) \pmod{\bar{\mathfrak{Q}}_0}$ depends only on the double coset $Hg\mathcal{D}_{\Omega_0}$ since H fixes $i_{K,L}(\delta) \in i_{K,L}(R_K)$ and \mathcal{D}_{Ω_0} fixes any element modulo $\bar{\mathfrak{Q}}_0$. The relation $g^{-1}(i_{K,L}(\delta)) \equiv g'^{-1}(i_{K,L}(\delta)) \pmod{\bar{\mathfrak{Q}}_0}$ for all $\delta \in I$ is obviously an equivalence relation on $H \backslash G / \mathcal{D}_{\Omega_0}$, and hence defines a partition of $H \backslash G / \mathcal{D}_{\Omega_0}$.

For $I(P)$, we fix $B \subseteq H \backslash G / \mathcal{D}_{\Omega_0}$ and first show that

$$t := \sum_{g\mathcal{D}_{\Omega_0} \in G/\mathcal{D}_{\Omega_0} : Hg\mathcal{D}_{\Omega_0} \in B} {}^g\delta_{\bar{\mathfrak{Q}}_0}$$

is well defined and does not depend on the choices of the representatives g . Note that for $h \in \mathcal{D}_{\Omega_0}$, the primitive idempotents $\delta_{\bar{\mathfrak{Q}}_0}$ and ${}^h\delta_{\bar{\mathfrak{Q}}_0}$ correspond to the same maximal ideal $\bar{\mathfrak{Q}}_0 = {}^h\bar{\mathfrak{Q}}_0$ and hence are equal (see Lemma 3.3). So $\delta_{\bar{\mathfrak{Q}}_0}$ is fixed by \mathcal{D}_{Ω_0} . It follows that t is well defined.

Next we prove $t \in i_{K,L}(R_K)$ so that $\delta_B = i_{K,L}^{-1}(t)$ is well defined. By Lemma 5.3, each double coset $x = Hg\mathcal{D}_{\Omega_0}$ corresponds to a maximal ideal

$$\mathfrak{P}_x := \frac{({}^g\bar{\mathfrak{Q}}_0 \cap \mathcal{O}_K) / p\mathcal{O}_K}{\text{Rad}(\bar{\mathcal{O}}_K)} \cap R_K$$

of R_K . Let δ be the idempotent of R_K satisfying $\delta \equiv 1 \pmod{\mathfrak{P}_x}$ for $x \in B$, and $\delta \equiv 0 \pmod{\mathfrak{P}_x}$ for $x \notin B$ (see Lemma 3.3). It follows that $i_{K,L}(\delta) \equiv 1 \pmod{{}^g\bar{\mathfrak{Q}}_0}$ if $Hg\mathcal{D}_{\Omega_0} \in B$ and $i_{K,L}(\delta) \equiv 0 \pmod{{}^g\bar{\mathfrak{Q}}_0}$ if $Hg\mathcal{D}_{\Omega_0} \notin B$. By definition, we also have $t \equiv 1 \pmod{{}^g\bar{\mathfrak{Q}}_0}$ if $Hg\mathcal{D}_{\Omega_0} \in B$ and $t \equiv 0 \pmod{{}^g\bar{\mathfrak{Q}}_0}$ if $Hg\mathcal{D}_{\Omega_0} \notin B$. So $t = i_{K,L}(\delta)$, as desired. Furthermore, by choosing $B = H \backslash G / \mathcal{D}_{\Omega_0}$ and $t = i_{K,L}(1) = 1$, we see that $\sum_{g \in G/\mathcal{D}_{\Omega_0}} {}^g\delta_{\bar{\mathfrak{Q}}_0} = 1$. It follows that $I(P)$ is a well defined idempotent decomposition of R_K .

For the second claim, we first check that the sets B_δ form a partition of $H \backslash G / \mathcal{D}_{\Omega_0}$ and the map $\delta \mapsto B_\delta$ is injective. To see this, note that if a double coset $Hg\mathcal{D}_{\Omega_0}$ lies in both B_δ and $B_{\delta'}$ for distinct $\delta, \delta' \in I$, then $g^{-1}(i_{K,L}(\delta)) \equiv g^{-1}(i_{K,L}(\delta')) \equiv 1 \pmod{\bar{\Omega}_0}$ by definition. But then $g^{-1}(i_{K,L}(\delta\delta')) \equiv 1 \pmod{\bar{\Omega}_0}$, contradicting the fact that $\delta\delta' = 0$. So the sets B_δ are disjoint and the map $\delta \mapsto B_\delta$ is injective. Furthermore, each $Hg\mathcal{D}_{\Omega_0} \in H \backslash G / \mathcal{D}_{\Omega_0}$ lies in at least one set B_δ since

$$\sum_{\delta \in I} g^{-1}(i_{K,L}(\delta)) \equiv g^{-1}\left(i_{K,L}\left(\sum_{\delta \in I} \delta\right)\right) \equiv 1 \pmod{\bar{\Omega}_0}. \quad (\text{C.1})$$

So the sets B_δ form a partition of $H \backslash G / \mathcal{D}_{\Omega_0}$.

It remains to show that the partition formed by the sets B_δ is $P(I)$. Assume two elements $Hg\mathcal{D}_{\Omega_0}, Hg'\mathcal{D}_{\Omega_0}$ are in the same set B_{δ_0} for some $\delta_0 \in I$, then $g^{-1}(i_{K,L}(\delta_0)) \equiv g'^{-1}(i_{K,L}(\delta_0)) \equiv 1 \pmod{\bar{\Omega}_0}$ by definition. By (C.1), we have $g^{-1}(i_{K,L}(\delta)) \equiv g'^{-1}(i_{K,L}(\delta)) \equiv 0 \pmod{\bar{\Omega}_0}$ for $\delta \in I - \{\delta_0\}$. So $Hg\mathcal{D}_{\Omega_0}, Hg'\mathcal{D}_{\Omega_0}$ are in the same block of $P(I)$.

Conversely, assume two elements $Hg\mathcal{D}_{\Omega_0}, Hg'\mathcal{D}_{\Omega_0}$ are in the same block of $P(I)$. So $g^{-1}(i_{K,L}(\delta)) \equiv g'^{-1}(i_{K,L}(\delta)) \pmod{\bar{\Omega}_0}$ for all $\delta \in I$. And by (C.1), we have $g^{-1}(i_{K,L}(\delta_0)) \equiv g'^{-1}(i_{K,L}(\delta_0)) \equiv 1 \pmod{\bar{\Omega}_0}$ for some $\delta_0 \in I$. Then $Hg\mathcal{D}_{\Omega_0}, Hg'\mathcal{D}_{\Omega_0} \in B_{\delta_0}$. So the partition formed by the sets B_δ is exactly $P(I)$. \square

Lemma 5.5. *The map $I \mapsto P(I)$ is a one-to-one correspondence between the idempotent decompositions of R_K and the partitions of $H \backslash G / \mathcal{D}_{\Omega_0}$, with the inverse map $P \mapsto I(P)$.*

Proof. It suffices to show that the map $B \mapsto \delta_B$ is a one-to-one correspondence between the subsets of $H \backslash G / \mathcal{D}_{\Omega_0}$ and the idempotents of R_K , with the inverse map $\delta \mapsto B_\delta$. Fix $B \subseteq H \backslash G / \mathcal{D}_{\Omega_0}$ and let $\delta = \delta_B$. We verify that $B_\delta = B$. For $h \in G$, we have

$$h^{-1}(i_{K,L}(\delta)) = \sum_{g\mathcal{D}_{\Omega_0} \in G/\mathcal{D}_{\Omega_0} : Hg\mathcal{D}_{\Omega_0} \in B} h^{-1}g\delta_{\bar{\Omega}_0}.$$

Note that the residue of $h^{-1}g\delta_{\bar{\Omega}_0}$ modulo $\bar{\Omega}_0$ equals one if $h\mathcal{D}_{\Omega_0} = g\mathcal{D}_{\Omega_0}$, and zero otherwise. So the residue of $h^{-1}(i_{K,L}(\delta))$ modulo $\bar{\Omega}_0$ equals one if $Hh\mathcal{D}_{\Omega_0} \in B$ and zero otherwise. It follows by definition that $B_\delta = B$.

It then suffices to show that the map $\delta \mapsto B_\delta$ is injective. Consider two distinct idempotents δ, δ' of R_K . As $\delta \neq \delta'$, there exists a maximal ideal ${}^g\bar{\Omega}_0$ of R_L for

which $i_{K,L}(\delta) \not\equiv i_{K,L}(\delta') \pmod{{}^g\bar{\Omega}_0}$ holds. Equivalently, we have $g^{-1}(i_{K,L}(\delta)) \not\equiv g^{-1}(i_{K,L}(\delta')) \pmod{\bar{\Omega}_0}$. Then exactly one of $g^{-1}(i_{K,L}(\delta))$, $g^{-1}(i_{K,L}(\delta'))$ equals one modulo $\bar{\Omega}_0$. So $Hg\mathcal{D}_{\Omega_0}$ is in one of B_δ , $B_{\delta'}$ but not the other. Therefore $B_\delta \neq B_{\delta'}$. \square

Lemma 5.6. *For any maximal ideal \mathfrak{m} of $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$, the group $\langle \sigma_{K,i} \rangle$ generated by $\sigma_{K,i}$ acts transitively on the set of the maximal ideal of $A_{K,i}$ containing \mathfrak{m} .*

Proof. Let $A = \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ and $H = \langle \sigma_{K,i} \rangle$. Equivalently, we want to prove that H acts transitively on the set of the maximal ideals of $A_{K,i}/\mathfrak{m}A_{K,i}$, where the action is induced from that on $A_{K,i}$.

We have a short exact sequence

$$0 \rightarrow \mathfrak{m} \rightarrow A \rightarrow A/\mathfrak{m} \rightarrow 0,$$

which by [AM69, Proposition 2.18] induces an exact sequence

$$\mathfrak{m} \otimes_{\mathbb{F}_q} \mathbb{F}_{q^i} \rightarrow A_{K,i} \rightarrow (A/\mathfrak{m}) \otimes_{\mathbb{F}_q} \mathbb{F}_{q^i} \rightarrow 0.$$

Also note that the image of $\mathfrak{m} \otimes_{\mathbb{F}_q} \mathbb{F}_{q^i}$ in $A_{K,i}$ is $\mathfrak{m}A_{K,i}$. Then we have

$$A_{K,i}/\mathfrak{m}A_{K,i} \cong (A/\mathfrak{m}) \otimes_{\mathbb{F}_q} \mathbb{F}_{q^i}.$$

So we want to prove that H acts transitively on the set of the maximal ideals of $(A/\mathfrak{m}) \otimes_{\mathbb{F}_q} \mathbb{F}_{q^i}$.

Suppose $\mathfrak{m}_1, \dots, \mathfrak{m}_k$ are maximal ideals of $(A/\mathfrak{m}) \otimes_{\mathbb{F}_q} \mathbb{F}_{q^i}$ that form an H -orbit, and $\delta_1, \dots, \delta_k$ are the corresponding primitive idempotents. Define $t := \sum_{i=1}^k \delta_i$ which is a nonzero idempotent fixed by H . It suffices to prove $t = 1$.

Note that we have the exact sequence

$$0 \rightarrow (A/\mathfrak{m})^H \rightarrow A/\mathfrak{m} \xrightarrow{\tau} A/\mathfrak{m},$$

where τ sends $x \in A$ to $x^q - x$. It induces a sequence

$$0 \rightarrow (A/\mathfrak{m})^H \otimes_{\mathbb{F}_q} \mathbb{F}_{q^i} \rightarrow (A/\mathfrak{m}) \otimes_{\mathbb{F}_q} \mathbb{F}_{q^i} \xrightarrow{\tau'} (A/\mathfrak{m}) \otimes_{\mathbb{F}_q} \mathbb{F}_{q^i},$$

where τ' sends $x \in (A/\mathfrak{m}) \otimes_{\mathbb{F}_q} \mathbb{F}_{q^i}$ to $\sigma_{K,i}(x) - x$. This sequence is exact since \mathbb{F}_{q^i} is a flat \mathbb{F}_q -module (see, e.g., [AM69, Proposition 2.19 and Exercise 2.4]). So we have

$$((A/\mathfrak{m}) \otimes_{\mathbb{F}_q} \mathbb{F}_{q^i})^H \cong (A/\mathfrak{m})^H \otimes_{\mathbb{F}_q} \mathbb{F}_{q^i} \cong \mathbb{F}_q \otimes_{\mathbb{F}_q} \mathbb{F}_{q^i} \cong \mathbb{F}_{q^i}$$

and the only nonzero idempotent it contains is 1. It follows that $t = 1$, as desired. \square

Lemma 5.7. *There exists a polynomial-time algorithm `ComputeRings` that given p and a relative number field K over K_0 , computes the following data*

- a p -maximal order \mathcal{O}'_K of K and the inclusion $\mathcal{O}'_K \hookrightarrow K$,
- $\bar{\mathcal{O}}_K$ and the quotient map $\mathcal{O}'_K \rightarrow \bar{\mathcal{O}}_K$,
- $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ and the quotient map $\bar{\mathcal{O}}_K \rightarrow \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$,
- R_K and the inclusion $R_K \hookrightarrow \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$,

where $\bar{\mathcal{O}}_K$, $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$, and R_K are encoded as algebras over \mathbb{F}_p and \mathcal{O}'_K is encoded as an algebra over \mathbb{Z} .

Proof. First use Corollary 4.1 to compute an ordinary number field \tilde{K} isomorphic to K and an isomorphism $\phi : K \rightarrow \tilde{K}$ in some \mathbb{Q} -basis of K . Apply Lemma 3.8 to \tilde{K} and p to compute $\bar{\mathcal{O}}_K$, \mathcal{O}'_K as well as the maps $\mathcal{O}'_K \hookrightarrow \tilde{K}$ and $\mathcal{O}'_K \rightarrow \bar{\mathcal{O}}_K$. Compose $\mathcal{O}'_K \hookrightarrow \tilde{K}$ with ϕ^{-1} to obtain the map $\mathcal{O}'_K \hookrightarrow K$.

Next we compute an \mathbb{F}_p -basis $B = \{x_1, \dots, x_s\}$ of the radical $\text{Rad}(\bar{\mathcal{O}}_K) \subseteq \bar{\mathcal{O}}_K$ using Theorem 5.4. Extend B to an \mathbb{F}_p -basis $B' = \{x_1, \dots, x_s, y_1, \dots, y_t\}$ of $\bar{\mathcal{O}}_K$. Compute $c_{ij}^k \in \mathbb{F}_p$ for $i, j \in [t]$, $k \in [s]$ and $d_{ij}^k \in \mathbb{F}_p$ for $i, j, k \in [t]$ such that

$$y_i y_j = \sum_{k=1}^s c_{ij}^k x_k + \sum_{k=1}^t d_{ij}^k y_k \quad \text{for } i, j \in [t].$$

Then the structure constants of $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ are given by d_{ij}^k in the \mathbb{F}_p -basis $\{y_1 + \text{Rad}(\bar{\mathcal{O}}_K), \dots, y_t + \text{Rad}(\bar{\mathcal{O}}_K)\}$ of $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ since

$$(y_i + \text{Rad}(\bar{\mathcal{O}}_K))(y_j + \text{Rad}(\bar{\mathcal{O}}_K)) = \sum_{k=1}^t d_{ij}^k y_k + \text{Rad}(\bar{\mathcal{O}}_K)$$

holds for $i, j \in [t]$. The map $\bar{\mathcal{O}}_K \rightarrow \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ is given in the basis B' which sends each x_i to zero and each y_i to $y_i + \text{Rad}(\bar{\mathcal{O}}_K)$.

Finally, we compute an \mathbb{F}_p -basis of R_K in $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ by solving the system of \mathbb{F}_p -linear equations given by $x^p = x$. It also gives the inclusion $R_K \hookrightarrow \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$. The structure constants of R_K can be computed from those of $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$. \square

Lemma 5.8. *There exists a polynomial-time algorithm `ComputeRingHoms` that given p , relative number fields K, K' over K_0 , a field embedding $\phi : K \rightarrow K'$*

over K_0 , and the outputs of `ComputeRings` (see Lemma 5.7) on the inputs (K, p) and (K', p) respectively, computes the maps $\bar{\phi} : \bar{\mathcal{O}}_K \rightarrow \bar{\mathcal{O}}_{K'}$, $\hat{\phi} : \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K) \rightarrow \bar{\mathcal{O}}_{K'}/\text{Rad}(\bar{\mathcal{O}}_{K'})$ and $\tilde{\phi} : R_K \rightarrow R_{K'}$.

Proof. To compute the map $\bar{\phi}$, we identify K and K' with ordinary number fields and apply Lemma 3.11: compute isomorphisms $\tau : K \rightarrow \tilde{K}$ and $\tau' : K' \rightarrow \tilde{K}'$ using Corollary 4.1 where \tilde{K} and \tilde{K}' are ordinary number fields. Compute the maps $\mathcal{O}'_K \hookrightarrow \tilde{K}$, $\mathcal{O}'_{K'} \hookrightarrow \tilde{K}'$ by composing $\mathcal{O}'_K \hookrightarrow K$, $\mathcal{O}'_{K'} \hookrightarrow K'$ with τ and τ' respectively. And compute the field embedding $\phi' = \tau' \circ \phi \circ \tau^{-1}$ from \tilde{K} to \tilde{K}' . Now use Lemma 3.11 to obtain the map $\bar{\phi}$.

The map $\hat{\phi} : \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K) \rightarrow \bar{\mathcal{O}}_{K'}/\text{Rad}(\bar{\mathcal{O}}_{K'})$ induced from $\bar{\phi}$ sends $x + \text{Rad}(\bar{\mathcal{O}}_K) \in \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ to $\bar{\phi}(x) + \text{Rad}(\bar{\mathcal{O}}_{K'})$. We can efficiently compute $\hat{\phi}$ from $\bar{\phi}$ since the quotient maps $\bar{\mathcal{O}}_K \rightarrow \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ and $\bar{\mathcal{O}}_{K'} \rightarrow \bar{\mathcal{O}}_{K'}/\text{Rad}(\bar{\mathcal{O}}_{K'})$ are given.

Finally, we restrict $\hat{\phi}$ to $\hat{\phi}|_{R_K} : R_K \rightarrow \bar{\mathcal{O}}_{K'}/\text{Rad}(\bar{\mathcal{O}}_{K'})$ using the given inclusion $R_K \hookrightarrow \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$. Then compute $\tilde{\phi} : R_K \rightarrow R_{K'}$ from $\hat{\phi}|_{R_K}$ by lifting along the given inclusion $R_{K'} \hookrightarrow \bar{\mathcal{O}}_{K'}/\text{Rad}(\bar{\mathcal{O}}_{K'})$. \square

Lemma 5.17. *Under GRH, there exists a subroutine `ComputeAdvice` that given $\mathcal{I} = \{I_K : K \in \mathcal{F}\}$ as in Definition 5.7, either properly refines some idempotent decomposition $I_K \in \mathcal{I}$, or computes e_δ, f_δ for $K \in \mathcal{F}$, $\delta \in I_K$ and an \mathcal{I} -advice. Moreover, the subroutine runs in time polynomial in $\log p$ and the size of \mathcal{F} .*

Proof. See Algorithm 18 for the pseudocode of the subroutine. It enumerates $K \in \mathcal{F}$, $\delta \in I_K$ and computes $e_\delta, f_\delta, s_\delta$ (if $e_\delta > 1$) and t_δ (if $f_\delta > 1$).

Fix $K \in \mathcal{F}$ and $\delta \in I_K$. We compute the ideal J of $\bar{\mathcal{O}}_K$, which is defined to be the preimage of the ideal of $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ generated by $1 - \delta$ (under the natural quotient map). Then J is the product of the maximal ideals \mathfrak{m} of $\bar{\mathcal{O}}_K$ satisfying $\delta \equiv 1 \pmod{\mathfrak{m}/\text{Rad}(\bar{\mathcal{O}}_K)}$. Note that for any such \mathfrak{m} , we have

$$\bar{\mathcal{O}}_K \supseteq \mathfrak{m} \supseteq \mathfrak{m}^2 \supseteq \cdots \supseteq \mathfrak{m}^{e_\delta} = \mathfrak{m}^{e_\delta+1}$$

and $\bar{\mathcal{O}}_K/\mathfrak{m} \cong \mathbb{F}_{q^{f_\delta}}$. So we can compute e_δ as the smallest positive integer i such that $J^i = J^{i+1}$, and compute f_δ as the smallest positive integer i such that the automorphism $x \mapsto x^{q^i}$ fixes $\bar{\mathcal{O}}_K/J$.

Algorithm 18 ComputeAdvice

```

1: for  $K \in \mathcal{F}$  do
2:   for  $\delta \in I_K$  do
3:      $J \leftarrow \{x \in \bar{\mathcal{O}}_K : x + \text{Rad}(\bar{\mathcal{O}}_K) \in (1 - \delta)(\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K))\}$ 
4:     compute  $e_\delta$  as the smallest  $i \in \mathbb{N}^+$  such that  $J^i = J^{i+1}$ 
5:     compute  $f_\delta$  as the smallest  $i \in \mathbb{N}^+$  such that  $x \mapsto x^{q^i}$  fixes  $\bar{\mathcal{O}}_K/J$ 
6:     if  $e_\delta > 1$  then
7:       find  $s_\delta \in J - J^2$ 
8:        $U \leftarrow$  the image of  $\text{Ann}_{\bar{\mathcal{O}}_K}(s_\delta^{e_\delta-1})$  in  $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ 
9:        $U \leftarrow U \cap \delta R_K$ 
10:      compute  $\delta_0 \in U$  satisfying  $(1 - \delta_0)U = \{0\}$ 
11:      if  $\delta_0\delta \notin \{0, \delta\}$  then
12:         $I_K \leftarrow I_K - \{\delta\}$ 
13:         $I_K \leftarrow I_K \cup \{\delta_0\delta, (1 - \delta_0)\delta\}$ 
14:      return
15:      if  $f_\delta > 1$  then
16:        find a primitive  $f_\delta$ th root of unity  $\xi \in \mathbb{F}_{q^{f_\delta}}$ 
17:        find nonzero  $t_\delta \in \delta A_{K, f_\delta}$  satisfying  $\sigma_{K, f_\delta}(t_\delta) = \xi t_\delta$ 
18:         $U \leftarrow t_\delta A_{K, f_\delta} \cap R_K$ 
19:        find  $\delta_0 \in U$  satisfying  $(1 - \delta_0)U = \{0\}$ 
20:        if  $\delta_0\delta \notin \{0, \delta\}$  then
21:           $I_K \leftarrow I_K - \{\delta\}$ 
22:           $I_K \leftarrow I_K \cup \{\delta_0\delta, (1 - \delta_0)\delta\}$ 
23:        return

```

Suppose $e_\delta > 1$. Choose s_δ to be an element in $J^2 - J$. So we have $s_\delta \in \mathfrak{m}$ for all the maximal ideals \mathfrak{m} of $\bar{\mathcal{O}}_K$ satisfying $\delta \equiv 1 \pmod{\mathfrak{m}/\text{Rad}(\bar{\mathcal{O}}_K)}$, and $s_\delta \notin \tilde{\mathfrak{m}}^2$ for some maximal ideal $\tilde{\mathfrak{m}}$ of them.

Next compute the image of $\text{Ann}_{\bar{\mathcal{O}}_K}(s_\delta^{e_\delta-1})$ in $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ and let U be its intersection with δR_K , which is an ideal of R_K . Choose an element δ_0 in U such that $(1 - \delta_0)U = 0$. Then δ_0 is the unique idempotent of R_K that generates U . If $\delta_0\delta \notin \{0, \delta\}$, we use δ_0 to properly refine I_K and return.

As $s_\delta \in \tilde{\mathfrak{m}} - \tilde{\mathfrak{m}}^2$, we have $s_\delta^{e_\delta-1} \in \tilde{\mathfrak{m}}^{e_\delta-1} - \tilde{\mathfrak{m}}^{e_\delta}$ and hence $\text{Ann}_{\bar{\mathcal{O}}_K}(s_\delta^{e_\delta-1}) \subseteq \tilde{\mathfrak{m}}$. So we have $\delta_0 \in U \subseteq \tilde{\mathfrak{m}}/\text{Rad}(\bar{\mathcal{O}}_K)$. But we also have $\delta \equiv 1 \pmod{\tilde{\mathfrak{m}}/\text{Rad}(\bar{\mathcal{O}}_K)}$. It follows that $\delta_0\delta \neq \delta$.

On the other hand, assume $s_\delta \in \tilde{\mathfrak{m}}'^2$ for some maximal ideal $\tilde{\mathfrak{m}}'$ of $\bar{\mathcal{O}}_K$ satisfying $\delta \equiv 1 \pmod{\tilde{\mathfrak{m}}'/\text{Rad}(\bar{\mathcal{O}}_K)}$. We claim $\delta_0\delta \neq 0$, in which case the subroutine properly refines I_K and returns. To see this, note that $s_\delta^{e_\delta-1} \in \tilde{\mathfrak{m}}'^{2(e_\delta-1)} \subseteq \tilde{\mathfrak{m}}'^{e_\delta}$ since $2(e_\delta - 1) \geq e_\delta$. Then $\text{Ann}_{\bar{\mathcal{O}}_K}(s_\delta^{e_\delta-1}) \not\subseteq \tilde{\mathfrak{m}}'$. Let δ' be the idempotent of $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ that generates the image of $\text{Ann}_{\bar{\mathcal{O}}_K}(s_\delta^{e_\delta-1})$ in $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$. Then $\delta' \notin \tilde{\mathfrak{m}}'/\text{Rad}(\bar{\mathcal{O}}_K)$. Note that $\delta_0 = \delta\delta' \notin \tilde{\mathfrak{m}}'/\text{Rad}(\bar{\mathcal{O}}_K)$. So $\delta_0\delta = \delta_0 \neq 0$, as desired.

Now suppose $f_\delta > 1$. We pick a primitive f_δ th root of unity ξ in $\mathbb{F}_{q^{f_\delta}}$ which exists since f_δ divides $|\mathbb{F}_{q^{f_\delta}}^\times| = q^{f_\delta} - 1$.¹ This step can be done efficiently assuming GRH. Choose t_δ to be a nonzero element in $\delta A_{K,f_\delta}$ satisfying $\sigma_{K,f_\delta}(t_\delta) = \xi t_\delta$. We claim that such an element always exists. To see this, note that the quotient map $A_{K,f_\delta} \rightarrow A_{K,f_\delta}/(1-\delta)$ is injective when restricting to $\delta A_{K,f_\delta}$. So it suffices to show that there exists a nonzero element $t \in A_{K,f_\delta}/(1-\delta)$ satisfying $\sigma_{K,f_\delta}(t) + (1-\delta) = \xi t + (1-\delta)$. This follows from the argument used in the proof of Lemma 3.18.

Next compute the ideal $U = t_\delta A_{K,f_\delta} \cap R_K$ of R_K , and choose an element δ_0 in U satisfying $(1-\delta_0)U = 0$. Then δ_0 is the unique idempotent of R_K that generates U . If $\delta_0\delta \notin \{0, \delta\}$, we use δ_0 to properly refine I_K and return.

Assume there exists a maximal ideal \mathfrak{m}_0 of A_{K,f_δ} satisfying $\delta \equiv 1 \pmod{\mathfrak{m}_0}$ and $t_\delta \in \mathfrak{m}_0$. Then $t_\delta A_{K,f_\delta} \subseteq \mathfrak{m}_0$ and hence $U \subseteq \mathfrak{m}_0$. As $\delta \notin \mathfrak{m}_0$, we have $\delta_0\delta \neq \delta$. We claim $\delta_0\delta$ is nonzero, and hence the subroutine properly refines I_K and returns. As $\delta_0 \in t_\delta A_{K,f_\delta} \subseteq \delta A_{K,f_\delta}$. We have $\delta_0\delta = \delta_0$, which generates the ideal U of R_K . So it suffices to prove $U \neq \{0\}$. As $t_\delta \neq 0$, there exists a maximal ideal \mathfrak{m} of A_{K,f_δ} that does not contain t_δ . Let $\mathfrak{m}' = \mathfrak{m} \cap R_K$, which is a maximal ideal of R_K . Let δ' be the primitive idempotent of R_K corresponding to \mathfrak{m}' , i.e., $\delta' \equiv 1 \pmod{\mathfrak{m}'}$ and $\delta' \equiv 0 \pmod{\mathfrak{m}''}$ for all the maximal ideals $\mathfrak{m}'' \neq \mathfrak{m}'$ of R_K . We claim $\delta' \in U$, or equivalently, $\delta' \in t_\delta A_{K,f_\delta}$. For $i \in \mathbb{Z}$, we have $\sigma_{K,f_\delta}^{-i}(t_\delta) = \xi^{-i}t_\delta \notin \mathfrak{m}$ and hence $t_\delta \notin \sigma_{K,f_\delta}^i(\mathfrak{m})$. By Lemma 5.6 and the choice of δ' , the maximal ideals of A_{K,f_δ} not containing δ' are exactly those of the form $\sigma_{K,f_\delta}^i(\mathfrak{m})$, $i \in \mathbb{Z}$. It follows that $\delta' \in t_\delta A_{K,f_\delta}$, as desired.

The claim about the running time is straightforward. □

Lemma 5.18. *The \mathcal{P} -collection $\tilde{\mathcal{C}}$ in Definition 5.8 is well defined.*

Proof. First note that each element $s_{\delta,H}$ (resp. $t_{\delta,H}$) is fixed by H and hence

¹We use the fact that f_δ is coprime to p , which in turn relies on the assumption $p > \deg(f)$.

$g^{-1}s_{\delta,H}, g'^{-1}s_{\delta,H}$ (resp. $g^{-1}t_{\delta,H}, g'^{-1}t_{\delta,H}$) only depend on the cosets Hg and Hg' . Fix $H \in \mathcal{P}$ and $K \in \mathcal{F}$ as in Definition 5.8. Fix $B \in C_H$ and $g, g' \in G$ such that $Hg\mathcal{D}_{\Omega_0}, Hg'\mathcal{D}_{\Omega_0} \in B$. Let δ be the unique idempotent in I_K such that $\tilde{\tau}_H(\delta) = \delta_B$. Then $e_\delta = e(B)$ and $f_\delta = f(B)$.

Suppose $e(\delta) > 1$. By Definition 5.7, we have $s_{\delta,H} \in \mathfrak{m} - \mathfrak{m}^2$ for all the maximal ideals \mathfrak{m} of $\bar{\mathcal{O}}_{LH}$ satisfying $\tilde{\tau}_H(\delta) \equiv 1 \pmod{\mathfrak{m}/\text{Rad}(\bar{\mathcal{O}}_{LH})}$. Consider the maximal ideal

$$\mathfrak{m}_{g,H} = ({}^g\Omega_0 \cap \mathcal{O}_{LH})/p\mathcal{O}_{LH}$$

of $\bar{\mathcal{O}}_{LH}$. As $\tilde{\tau}_H(\delta) = \delta_B$ and $Hg\mathcal{D}_{\Omega_0} \in B$, we have

$$\tilde{\tau}_H(\delta) \equiv 1 \pmod{\mathfrak{m}_{g,H}/\text{Rad}(\bar{\mathcal{O}}_{LH})}.$$

Therefore $s_{\delta,H} \in \mathfrak{m}_{g,H} - \mathfrak{m}_{g,H}^2$. So $s_{\delta,H} + \mathfrak{m}_{g,H}^2$ is a nonzero element in $\mathfrak{m}_{g,H}/\mathfrak{m}_{g,H}^2$. Let \mathfrak{m}_g be the maximal ideal ${}^g\Omega_0/p\mathcal{O}_L$ of $\bar{\mathcal{O}}_L$, and let $k = e(\Omega_0)/e_\delta$. Using the natural inclusion

$$\mathfrak{m}_{g,H}/\mathfrak{m}_{g,H}^2 \hookrightarrow \mathfrak{m}_g^k/\mathfrak{m}_g^{k+1},$$

we see $s_{\delta,H} + \mathfrak{m}_g^{k+1}$ is a nonzero element in $\mathfrak{m}_g^k/\mathfrak{m}_g^{k+1}$. Let $\mathfrak{m}_e := \Omega_0/p\mathcal{O}_L$, so that ${}^g\mathfrak{m}_e = \mathfrak{m}_g$. Then $g^{-1}s_{\delta,H} + \mathfrak{m}_e^{k+1}$ is a nonzero element in $\mathfrak{m}_e^k/\mathfrak{m}_e^{k+1}$. The same argument shows that $g'^{-1}s_{\delta,H} + \mathfrak{m}_e^k$ is a nonzero element in $\mathfrak{m}_e^k/\mathfrak{m}_e^{k+1}$ as well. As $\mathfrak{m}_e^k/\mathfrak{m}_e^{k+1}$ is an one-dimensional vector space over $\bar{\mathcal{O}}_L/\mathfrak{m}_e \cong \kappa_{\Omega_0}$, we see that there exists a unique scalar $c \in \kappa_{\Omega_0}^\times$ satisfying

$$g^{-1}s_{\delta,H} + \mathfrak{m}_e^{k+1} = c \cdot (g'^{-1}s_{\delta,H} + \mathfrak{m}_e^{k+1}).$$

Note $\mathfrak{m}_e^{k+1} = (\Omega_0/p\mathcal{O}_L)^{e(\Omega_0)/e_\delta+1}$. We see that the second condition in Definition 5.8 is well defined.

Now suppose $f(\delta) > 1$. By Definition 5.7, we have $t_{\delta,H} \notin \mathfrak{m}$ for all the maximal ideals \mathfrak{m} of A_{LH,f_δ} satisfying $\tilde{\tau}_H(\delta) \equiv 1 \pmod{\mathfrak{m}}$. As $\tilde{\tau}_H(\delta) = \delta_B$, $Hg\mathcal{D}_{\Omega_0} \in B$ and $\frac{\Omega_0/p\mathcal{O}_L}{\text{Rad}(\bar{\mathcal{O}}_L)} \subseteq \mathfrak{m}_0$, we have

$$\tilde{\tau}_H(\delta) \equiv 1 \pmod{{}^g\mathfrak{m}_0}.$$

So $t_{\delta,H} \notin {}^g\mathfrak{m}_0$. Then $g^{-1}t_{\delta,H} + \mathfrak{m}_0$ is a nonzero element in $A_{L,f_\delta}/\mathfrak{m}_0$. The same argument shows that $g'^{-1}t_{\delta,H} + \mathfrak{m}_0$ is a nonzero element in $A_{L,f_\delta}/\mathfrak{m}_0$ as well. It follows that there exists a unique scalar $c \in (A_{L,f_\delta}/\mathfrak{m}_0)^\times$ satisfying

$$g^{-1}t_{\delta,H} + \mathfrak{m}_0 = c \cdot (g'^{-1}t_{\delta,H} + \mathfrak{m}_0). \quad (\text{C.2})$$

We also check that Definition 5.8 is independent of the choice of \mathfrak{m}_0 : Let \mathfrak{m}'_0 be another maximal ideal of A_{L,f_δ} containing $\frac{\Omega_0/p\mathcal{O}_L}{\text{Rad}(\mathcal{O}_L)}$. By Lemma 5.6, we have $\mathfrak{m}'_0 = \sigma_{L,f_\delta}^i(\mathfrak{m}_0)$ for some $i \in \mathbb{Z}$. Let $\sigma = \sigma_{L,f_\delta}^i$. Then (C.2) is equivalent to

$$\sigma(g^{-1}t_{\delta,H}) + \mathfrak{m}'_0 = \sigma(c) \cdot (\sigma(g'^{-1}t_{\delta,H}) + \mathfrak{m}'_0) \quad (\text{C.3})$$

where $\sigma(c) \in (A_{L,f_\delta}/\mathfrak{m}'_0)^\times$. Also note that

$$\sigma(g^{-1}t_{\delta,H}) = g^{-1}(\sigma(t_{\delta,H})) = g^{-1}(\xi^i t_{\delta,H}) = \xi^i g^{-1}t_{\delta,H}.$$

and similarly $\sigma(g'^{-1}t_{\delta,H}) = \xi^i g'^{-1}t_{\delta,H}$. Substituting them in (C.3) and canceling $\xi^i + \mathfrak{m}'_0$ on both sides, we obtain

$$g^{-1}t_{\delta,H} + \mathfrak{m}'_0 = \sigma(c) \cdot (g'^{-1}t_{\delta,H} + \mathfrak{m}'_0).$$

Note that $\sigma(c)$ and c have the same order. We see that choosing \mathfrak{m}'_0 instead of \mathfrak{m}_0 does not affect the definition.

Finally, it is easy to see that the conditions in Definition 5.8 are equivalence relations on $H \backslash G$. So they do define a partition C_H on $H \backslash G$. \square

Lemma 5.24. *Under GRH, there exists a subroutine `SurjectivityTest` that updates I_K in time polynomial in $\log p$ and the size of \mathcal{F} so that the partitions $C_H \in \mathcal{C}$ are refined. Moreover, at least one partition C_H is properly refined unless for all $H \in \mathcal{P}$, the map $\pi_H : H \backslash G \rightarrow H \backslash G / \mathcal{D}_{\Omega_0}$ sending $Hg \in H \backslash G$ to $Hg\mathcal{D}_{\Omega_0}$ maps each block of \tilde{C}_H surjectively to a block of C_H .*

To prove Lemma 5.24, we first prove the following lemma, generalizing Lemma B.1:

Lemma C.1. *There exists an algorithm `SplitByExp` that, given a semisimple \mathbb{F}_p -algebra A , $m \in \mathbb{N}^+$, and nonzero elements $x, y \in A$ satisfying the following conditions*

- x and y generate the same ideal of A
- Let n_x (resp. n_y) be the smallest positive integer such that x^{n_x} (resp. y^{n_y}) is an idempotent. Then n_y divides n_x and all the prime factors of n_x divide m

returns an element $z = x^k - y \in A$ satisfying $zA \subsetneq xA$ in time polynomial in m and $\log |A|$, where $k \in \mathbb{N}$.

Proof. We find $k \in \mathbb{N}$ such that $x^k - y$ satisfies the requirement. Let $I = \text{Ann}_A(x)$. By replacing A , x , and y with A/I , $x + I$ and $y + I$ respectively, we reduce to the case $x, y \in A^\times$, and the goal is to find $k \in \mathbb{N}$ such that $z = x^k - y$ is a zero divisor. In addition, we find the smallest $d \in \mathbb{N}^+$ such that the ideal J generated by $\{x^{p^d} - x : x \in A\}$ is a proper ideal of A . By replacing A with A/J , we may assume $J = \{0\}$. Then A is a finite product of copies of \mathbb{F}_{p^d} .

Enumerate the prime factors ℓ of m . For each ℓ , compute $e_\ell \in \mathbb{N}$ and $f_\ell \in \mathbb{N}^+$ such that $p^d - 1 = \ell^{e_\ell} f_\ell$ and f_ℓ is coprime to ℓ . Let $n_{x,\ell}$ and $n_{y,\ell}$ be the order of x^{f_ℓ} and y^{f_ℓ} respectively. Then $n_{x,\ell}, n_{y,\ell}$ are powers of ℓ and $n_{y,\ell} | n_{x,\ell}$. Use the algorithm in Lemma B.1 (applied to x^{f_ℓ} and y^{f_ℓ}) to compute $k_\ell \in \mathbb{N}$ such that $x^{k_\ell f_\ell} - y^{f_\ell}$ is a zero divisor. If $x^{k_\ell f_\ell} - y^{f_\ell} \neq 0$, we use Lemma 3.17 to find an idempotent $\gamma \notin \{0, 1\}$ of A and solve the problem recursively on the quotient ring $A/(1 - \gamma)$. So assume $x^{k_\ell f_\ell} = y^{f_\ell}$. Then the order of x^{k_ℓ}/y divides f_ℓ and hence is coprime to ℓ .

Compute k_ℓ and e_ℓ for all the prime factors ℓ of m as above. Use the extended Euclidean algorithm to find $k \in \mathbb{N}$ satisfying $k \equiv k_\ell \pmod{\ell^{e_\ell}}$ for all ℓ . Then k is the desired integer.

We claim $x^k = y$. To see this, note that for each ℓ , we have $x^k/y = (x^{k_\ell}/y) \cdot x^{t\ell^{e_\ell}}$ for some $t \in \mathbb{Z}$. As the orders of x^{k_ℓ}/y and $x^{t\ell^{e_\ell}}$ are both coprime to ℓ , so is the order of x^k/y . Therefore the order of x^k/y is coprime to m . But the orders of x^k and y are only divisible by prime factors of m . So $x^k/y = 1$, as desired. \square

The pseudocode of the subroutine `SurjectivityTest` is given in Algorithm 19. It enumerates $K \in F$ and $\delta \in I_K$. For each K and δ , a set S of ideals of A_{K,f_δ} is computed. And for each $I \in S$, we find $\delta_0 \in I \cap R_K$ satisfying $(1 - \delta_0)(I \cap R_K) = \{0\}$, which is the unique idempotent of R_K that generates the ideal $I \cap R_K$ of R_K .² If $\delta_0 \delta \notin \{0, \delta\}$, we use δ_0 to refine I_K and return.

Fix $K \in F$ and $\delta \in I_K$. The corresponding set S is computed as follows: first assume $f_\delta > 1$. We compute the largest factor r of $q^{f_\delta} - 1$ coprime to f_δ , so that all the prime factors of $(q^{f_\delta} - 1)/r$ divide f_δ . Compute an element $\gamma \in \mathbb{F}_{q^{f_\delta}}^\times$ of order $(q^{f_\delta} - 1)/r$, which can be done efficiently assuming GRH.³ By the second

²Here R_K is regarded as a subring of A_{K,f_δ} via the inclusions $R_K \hookrightarrow \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ and $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K) \hookrightarrow A_{K,f_\delta}$.

³For example, we can achieve this by computing an ℓ th power non-residue γ_ℓ for each prime factor ℓ of f_δ . By raising γ_ℓ to its r_ℓ th power, where r_ℓ is the largest factor of $q^{f_\delta} - 1$ coprime to ℓ , we may assume the order of γ_ℓ is $(q^{f_\delta} - 1)/r_\ell$. Then let γ be the product of all γ_ℓ .

Algorithm 19 SurjectivityTest

```

1: for  $K \in \mathcal{F}$  do
2:   for  $\delta \in I_K$  do
3:      $S \leftarrow \emptyset$ 
4:     if  $f_\delta > 1$  then
5:        $r \leftarrow$  the largest factor of  $q^{f_\delta} - 1$  coprime to  $f_\delta$ 
6:       compute  $\gamma \in \mathbb{F}_{q^{f_\delta}}^\times$  of order  $(q^{f_\delta} - 1)/r$ 
7:       call SplitByExp on  $(A_{K,f_\delta}, f_\delta, \delta\gamma, \delta t_\delta^r)$  to obtain  $x \in A_{K,f_\delta}$ 
8:        $S \leftarrow S \cup \{xA_{K,f_\delta}\}$ 
9:     if  $e_\delta > 1$  then
10:       $J \leftarrow$  the preimage of  $\delta(\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K))$  in  $\bar{\mathcal{O}}_K$ 
11:      find  $\delta' \in \text{Ann}_{\bar{\mathcal{O}}_K}(J^{e_\delta})$  satisfying  $(1 - \delta')\text{Ann}_{\bar{\mathcal{O}}_K}(J^{e_\delta}) = \{0\}$ 
12:      lift  $\delta's_\delta \in \bar{\mathcal{O}}_K$  to  $\tilde{s} \in \mathcal{O}'_K$ 
13:      compute the image  $s$  of  $\tilde{s}^{e_\delta}/p$  in  $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ 
14:       $r' \leftarrow$  the largest factor of  $q^{f_\delta} - 1$  coprime to  $e_\delta$ 
15:      compute  $\mu \in \mathbb{F}_{q^{f_\delta}}^\times$  of order  $(q^{f_\delta} - 1)/r'$ 
16:      call SplitByExp on  $(A_{K,f_\delta}, e_\delta, \delta\mu, s^{r'})$  to obtain  $y \in A_{K,f_\delta}$ 
17:       $S \leftarrow S \cup \{yA_{K,f_\delta}\}$ 
18:      if  $f_\delta > 1$  then
19:        for  $i \leftarrow 0$  to  $f_\delta - 1$  do
20:           $S \leftarrow S \cup \{yA_{K,f_\delta} + \sigma_{K,f_\delta}^i(x)A_{K,f_\delta}\}$ 
21:      for  $I \in S$  do
22:        find  $\delta_0 \in I \cap R_K$  satisfying  $(1 - \delta_0)(I \cap R_K) = \{0\}$ 
23:        if  $\delta_0\delta \notin \{0, \delta\}$  then
24:           $I_K \leftarrow I_K - \{\delta\}$ 
25:           $I_K \leftarrow I_K \cup \{\delta_0\delta, (1 - \delta_0)\delta\}$ 
26:      return

```

condition in Definition 5.7, the element δt_δ^r generates the ideal $\delta A_{K,f_\delta}$ of A_{K,f_δ} , and so does $\delta\gamma$. We call the subroutine SplitByExp in Lemma C.1 on the input $(A_{K,f_\delta}, f_\delta, \delta\gamma, \delta t_\delta^r)$ to obtain $x \in A_{K,f_\delta}$, and add the ideal xA_{K,f_δ} to S .

Next assume $e_\delta > 1$. Compute the preimage J of $(1 - \delta)(\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K))$ under the quotient map $\bar{\mathcal{O}}_K \rightarrow \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$. Then J is the product of the maximal ideals \mathfrak{m} of $\bar{\mathcal{O}}_K$ satisfying $\delta \equiv 1 \pmod{\mathfrak{m}/\text{Rad}(\bar{\mathcal{O}}_K)}$. Find $\delta' \in \text{Ann}_{\bar{\mathcal{O}}_K}(J^{e_\delta})$ satisfying $(1 - \delta')\text{Ann}_{\bar{\mathcal{O}}_K}(J^{e_\delta}) = \{0\}$, so that δ' is the unique idempotent of $\bar{\mathcal{O}}_K$ generating

$\text{Ann}_{\bar{\mathcal{O}}_K}(J^{e_\delta})$. Lift $\delta' s_\delta \in \bar{\mathcal{O}}_K$ to $\tilde{s} \in \mathcal{O}'_K$.

We claim $\tilde{s}^{e_\delta} \in p\mathcal{O}_K$: this is equivalent to $(\delta' s_\delta)^{e_\delta} = \delta' s_\delta^{e_\delta} = 0$. By the first condition in Definition 5.7, we have $s_\delta^{e_\delta} \in \mathfrak{m}^{e_\delta}$ for all the maximal ideals \mathfrak{m} of $\bar{\mathcal{O}}_K$ satisfying $\delta \equiv 1 \pmod{\mathfrak{m}/\text{Rad}(\bar{\mathcal{O}}_K)}$. And by the definition of J , it holds that $\delta' \in \mathfrak{m}^k$ for all the maximal ideals \mathfrak{m} of $\bar{\mathcal{O}}_K$ satisfying $\delta \in \mathfrak{m}/\text{Rad}(\bar{\mathcal{O}}_K)$ and $k \in \mathbb{N}$. It follows that $\delta' s_\delta^{e_\delta} = 0$ and hence $\tilde{s}^{e_\delta} \in p\mathcal{O}_K$.

Compute the image s of $\tilde{s}^{e_\delta}/p \in \mathcal{O}_K$ in $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$. This is done by first computing $\tilde{s}^{e_\delta} + \mathcal{O}_K \in \bar{\mathcal{O}}_K$ using Lemma 3.9 and then computing s using the quotient map $\bar{\mathcal{O}}_K \rightarrow \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$. Next compute the largest factor r' of $q^{f_\delta} - 1$ coprime to e_δ , so that all the prime factors of $(q^{f_\delta} - 1)/r'$ divide e_δ . Compute an element $\mu \in \mathbb{F}_{q^{f_\delta}}^\times$ of order $(q^{f_\delta} - 1)/r'$, which can be done efficiently assuming GRH. By the first condition in Definition 5.7, the element $s^{r'}$ generates the ideal $\delta A_{K,f_\delta}$ of A_{K,f_δ} , and so does $\delta\mu$.⁴ We call the subroutine `SplitByExp` on the input $(A_{K,f_\delta}, e_\delta, \delta\mu, s^{r'})$ to obtain $y \in A_{K,f_\delta}$, and add the ideal yA_{K,f_δ} to S . In addition, if $f_\delta > 1$, we enumerate $i = 0, 1, \dots, f_\delta - 1$ and for each i , we add the ideal of A_{K,f_δ} generated by y and $\sigma_{K,f_\delta}^i(x)$ to S , where $x \in A_{K,f_\delta}$ is computed in the case $f_\delta > 1$ above.

Now we prove Lemma 5.24.

Proof of Lemma 5.24. Assume for some $H \in \mathcal{P}$, $B \in C_H$, and $\tilde{B} \in \tilde{C}_H$, the map $\pi_H : Hh \mapsto Hh\mathcal{D}_{\Omega_0}$ maps \tilde{B} to a proper subset of B . Let K be the field in \mathcal{F} isomorphic to L^H over K_0 . Let δ be the idempotent in I_K satisfying $\tilde{\tau}_H(\delta) = \delta_B$ (see Definition 5.4). We show that in the corresponding iteration of the loop in Lines 3–27, we compute a set S that contains an ideal I of A_{K,f_δ} such that the unique idempotent $\delta_0 \in R_K$ generating $I \cap R_K$ satisfies $\delta_0\delta \notin \{0, \delta\}$. Consequently, some partition in \mathcal{C} is properly refined.

Choose $g, g' \in G$ such that $Hg\mathcal{D}_{\Omega_0} \in B - \pi_H(\tilde{B})$ and $Hg' \in \tilde{B}$. Let \mathfrak{m}_0 be an arbitrary maximal ideal of A_{L,f_δ} containing $\frac{\Omega_0/p\mathcal{O}_L}{\text{Rad}(\mathcal{O}_L)}$. Fix $\sigma \in \mathcal{D}_{\Omega_0}$ whose image in $\text{Gal}(\kappa_{\Omega_0}/\bar{\mathcal{O}}_{K_0})$ is the Frobenius automorphism $x \mapsto x^q$ over \mathbb{F}_q .

We necessarily have $f_\delta > 1$ or $e_\delta > 1$. First assume $f_\delta > 1$. Let $\gamma \in \mathbb{F}_{q^{f_\delta}}^\times$ be of order $(q^{f_\delta} - 1)/r$, where r is the largest factor of $q^{f_\delta} - 1$ coprime to f_δ . Consider an element $x \in A_{L^H,f_\delta}$ of the form $x = (\delta_B\gamma)^k - \delta_B t_{\delta,H}^r = \delta_B(\gamma^k - t_{\delta,H}^r)$ such that $xA_{L^H,f_\delta} \subsetneq \delta_B A_{L^H,f_\delta}$, where $k \in \mathbb{N}$. Let δ_0 be the unique idempotent of

⁴We let $A_{K,f_\delta} = \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ if $f_\delta = 1$.

R_{L^H} generating $x A_{L^H, f_\delta} \cap R_{L^H}$. The assumption $x A_{L^H, f_\delta} \subsetneq \delta_B A_{L^H, f_\delta}$ implies $\delta_0 \delta_B = \delta_0$ and $\delta_0 \neq \delta_B$. If $\delta_0 \neq 0$, by identifying K with L^H via the isomorphism $\tau_H : K \rightarrow L^H$, we see the ideal added to S at Line 8 is used in Lines 24–26 to properly refine I_K . So assume $\delta_0 = 0$, or equivalently $x A_{L^H, f_\delta} \cap R_{L^H} = \{0\}$.

Consider arbitrary $h \in G$, and let δ_1 be the primitive idempotent of R_{L^H} corresponding to the maximal ideal ${}^h \tilde{\mathcal{Q}}_0 \cap R_{L^H}$. Then a maximal ideal \mathfrak{m} of R_L satisfies $\delta_1 \equiv 1 \pmod{\mathfrak{m}}$ iff $\mathfrak{m} = {}^{h'} \tilde{\mathcal{Q}}_0$ for some $h' \in Hh$. This follows from Lemma 5.3 and the fact that \mathcal{D}_{Ω_0} fixes $\tilde{\mathcal{Q}}_0$ setwisely. So a maximal ideal \mathfrak{m}' of A_{L, f_δ} satisfies $\delta_1 \equiv 1 \pmod{\mathfrak{m}'}$ iff $\mathfrak{m}' \supseteq {}^{h'} \tilde{\mathcal{Q}}_0$ for some $h' \in Hh$. As $x A_{L^H, f_\delta} \cap R_{L^H} = \{0\}$, we have $\delta_1 \notin x A_{L^H, f_\delta}$. So for some $h' \in Hh$ and a maximal ideal $\mathfrak{m}' \supseteq {}^{h'} \tilde{\mathcal{Q}}_0$ of A_{L, f_δ} , we have $x \in \mathfrak{m}' \cap A_{L^H, f_\delta}$, and hence $h^{-1}x = h'^{-1}x \in {}^{h'^{-1}}\mathfrak{m}' \supseteq \tilde{\mathcal{Q}}_0$. By Lemma 5.6, we have ${}^{h'^{-1}}\mathfrak{m}' = \sigma_{L, f_\delta}^{-i_0}(\mathfrak{m}_0)$ for some $i_0 \in \mathbb{Z}$. Therefore

$$\sigma_{L, f_\delta}^{i_0}({}^{h^{-1}}x) \in \mathfrak{m}_0. \quad (\text{C.4})$$

Suppose the element h above satisfies $Hh\mathcal{D}_{\Omega_0} \in B$. Then $\sigma_{L, f_\delta}^{i_0}({}^{h^{-1}}\delta_B) = {}^{h^{-1}}\delta_B \equiv 1 \pmod{\mathfrak{m}_0}$. As $x = \delta_B(\gamma^k - t_{\delta, H}^r)$, (C.4) implies

$$\xi^{i_0 r} {}^{h^{-1}}(t_{\delta, H}^r) = {}^{h^{-1}}(\sigma_{L, f_\delta}^{i_0}(t_{\delta, H}^r)) = \sigma_{L, f_\delta}^{i_0}({}^{h^{-1}}(t_{\delta, H}^r)) \equiv \gamma^k \pmod{\mathfrak{m}_0},$$

where ξ is the primitive f_δ th root of unity satisfying $\sigma_{K, f_\delta}(t_\delta) = \xi \cdot t_\delta$ as in Definition 5.7. Choosing h to be g and g' respectively and using the fact r is coprime to f_δ , we see that there exists a unique integer $i \in \{0, \dots, f_\delta - 1\}$ satisfying

$$\xi^{i r} {}^{g^{-1}}(t_{\delta, H}^r) \equiv {}^{g'^{-1}}(t_{\delta, H}^r) \pmod{\mathfrak{m}_0}. \quad (\text{C.5})$$

As

$$\xi^{i r} {}^{g^{-1}}(t_{\delta, H}^r) \equiv {}^{g^{-1}}(\sigma_{L, f_\delta}^i(t_{\delta, H}^r)) \equiv \sigma_{L, f_\delta}^i({}^{g^{-1}}(t_{\delta, H}^r)) \equiv \sigma^{i g^{-1}}(t_{\delta, H}^r) \pmod{\mathfrak{m}_0},$$

and r is coprime to f_δ , we see that i is the unique integer in $\{0, \dots, f_\delta - 1\}$ such that the order of the element c in $(A_{L, f_\delta}/\mathfrak{m}_0)^\times$ satisfying

$$\sigma^{i g^{-1}} t_{\delta, H} + \mathfrak{m}_0 = c \cdot ({}^{g'^{-1}} t_{\delta, H} + \mathfrak{m}_0)$$

is coprime to f_δ . So for all $\omega \in \mathcal{I}_{\Omega_0}$, the third condition in Definition 5.8 is satisfied by $Hg\sigma^{-i}\omega^{-1}$ and Hg' , and is not satisfied by $Hg\sigma^{-i'}\omega^{-1}$ and Hg' for $i' \in \{0, \dots, f_\delta - 1\} - \{i\}$. In particular, if $e_\delta = 1$, then $Hg\sigma^{-i}$ and Hg' are in the same block \tilde{B} by Definition 5.8, contradicting the assumption $Hg\mathcal{D}_{\Omega_0} \notin \pi_H(\tilde{B})$. So the subroutine properly refines I_K if $f_\delta > 1$ and $e_\delta = 1$.

Next assume $e_\delta > 1$. Consider the ideal yA_{K,f_δ} of A_{K,f_δ} added to S at Line 18, and let δ_0 be the idempotent of R_K generating $yA_{K,f_\delta} \cap R_K$. Note $yA_{K,f_\delta} \subsetneq \delta A_{K,f_\delta}$. So $\delta_0\delta = \delta_0 \neq \delta$. If $\delta \neq 0$, we properly refine I_K using δ_0 in Lines 24–26. So assume $\delta_0 = 0$, or equivalently $yA_{K,f_\delta} \cap R_K = \{0\}$. Using the isomorphism $\tau_H : K \rightarrow L^H$, we regard y as an element of A_{L^H,f_δ} . So the assumption becomes $yA_{L^H,f_\delta} \cap R_{L^H} = \{0\}$.

Let c be the unique element in $\kappa_{\Omega_0}^\times$ satisfying

$$\sigma^i g^{-1} s_{\delta,H} + I = c \cdot (g'^{-1} s_{\delta,H} + I), \quad (\text{C.6})$$

where $I = (\mathfrak{Q}_0/p\mathcal{O}_L)^{e(\mathfrak{Q}_0)/e_\delta+1}$ (see Definition 5.8), and i is the unique integer in $\{0, \dots, f_\delta - 1\}$ satisfying (C.5) above (if $f_\delta = 1$, we let σ^i be the identity). Then the element \tilde{s} computed at Line 13 (regarded as an element of \mathcal{O}_L) satisfies

$$\sigma^i g^{-1} \tilde{s} + \mathfrak{Q}_0^{e(\mathfrak{Q}_0)/e_\delta+1} = c \cdot (g'^{-1} \tilde{s} + \mathfrak{Q}_0^{e(\mathfrak{Q}_0)/e_\delta+1})$$

and hence

$$\sigma^i g^{-1} (\tilde{s}^{e_\delta}) + \mathfrak{Q}_0^{e(\mathfrak{Q}_0)+1} = c^{e_\delta} \cdot (g'^{-1} (\tilde{s}^{e_\delta}) + \mathfrak{Q}_0^{e(\mathfrak{Q}_0)+1}).$$

We have $p \in \mathfrak{Q}_0^{e(\mathfrak{Q}_0)} - \mathfrak{Q}_0^{e(\mathfrak{Q}_0)+1}$ and it is fixed by G . So the element s computed at Line 14 (regarded as an element of $\bar{\mathcal{O}}_L/\text{Rad}(\bar{\mathcal{O}}_L)$) satisfies

$$\sigma^i g^{-1} s + \mathfrak{m} = c^{e_\delta} \cdot (g'^{-1} s + \mathfrak{m}), \quad (\text{C.7})$$

where $\mathfrak{m} = \frac{\mathfrak{Q}_0/p\mathcal{O}_L}{\text{Rad}(\mathcal{O}_L)}$.

Fix a generator ω of \mathcal{L}_{Ω_0} . The proof of Lemma 5.23 shows that

$$\omega \sigma^i g^{-1} s_{\delta,H} + I = c' (\sigma^i g^{-1} s_{\delta,H} + I) \quad (\text{C.8})$$

for some primitive e_δ th root of unity $c' \in \kappa_{\Omega_0}^\times$.

If $f_\delta = 1$, we have $A_{L^H,f_\delta} = \bar{\mathcal{O}}_{L^H}/\text{Rad}(\bar{\mathcal{O}}_{L^H})$, and its maximal ideals correspond one-to-one to those of R_{L^H} . So $yA_{L^H,f_\delta} \cap R_{L^H} = \{0\}$ implies $y = 0$. Note that y is of the form $\delta_B \mu^\ell - s^{r'}$ where $\ell \in \mathbb{N}$, r' is the largest factor of $q^{f_\delta} - 1$ coprime to e_δ , and μ is an element in $\mathbb{F}_{q^{f_\delta}}^\times$ of order $(q^{f_\delta} - 1)/r'$. We have $g^{-1} \delta_B, g'^{-1} \delta_B \equiv 1 \pmod{\mathfrak{m}}$ since $Hg\mathcal{D}_{\Omega_0}, Hg'\mathcal{D}_{\Omega_0} \in B$. As $g^{-1} y = g'^{-1} y = 0$, we have $g^{-1} (s^{r'}) \equiv g'^{-1} (s^{r'}) \pmod{\mathfrak{m}}$. Combining it with (C.7), we see $c^{r'}$ is an e_δ th root of unity. On the other hand, we know c' is a primitive e_δ th root of unity, and so is $c'^{r'}$ since r' is coprime to e_δ . Therefore there exists $j \in \{0, \dots, e_\delta - 1\}$ such that $(c'^{r'})^j c^{r'} = 1$. Then the

order of $c'^j c$ divides r' , and hence is coprime to e_δ . On the other hand, by (C.6) and (C.8), we have

$$\omega^j \sigma^i g^{-1} s_{\delta,H} + I = c'^j c \cdot (g'^{-1} s_{\delta,H} + I).$$

So by Definition 5.8 and the fact $Hg' \in \tilde{B}$, we have $Hg\sigma^{-i}\omega^{-j} \in \tilde{B}$. But this is a contradiction to the assumption $Hg\mathcal{D}_{\Omega_0} \notin \pi_H(\tilde{B})$.

Next consider the case $f_\delta > 1$ (and $e_\delta > 1$). Let $x, y \in A_{L^H, f_\delta}$ be as above. We claim that there exists $i' \in \{0, \dots, f_\delta - 1\}$ such that

$$yA_{L^H, f_\delta} + \sigma_{L^H, f_\delta}^{i'}(x)A_{L^H, f_\delta} \subsetneq \delta_B A_{L^H, f_\delta}. \quad (\text{C.9})$$

To see this, choose a maximal ideal \mathfrak{m}_1 of A_{L^H, f_δ} containing y but not δ_B , which exists since $yA_{L^H, f_\delta} \subsetneq \delta_B A_{L^H, f_\delta}$. Let $\mathfrak{m}'_1 = \mathfrak{m}_1 \cap R_{L^H}$. As $xA_{L^H, f_\delta} \cap R_{L^H} = \{0\}$, there exists a maximal ideal $\mathfrak{m}_2 \supseteq \mathfrak{m}'_1$ of A_{L^H, f_δ} containing x . By Lemma 5.6, there exists $i' \in \mathbb{Z}$ such that $\sigma_{L^H, f_\delta}^{i'}(\mathfrak{m}_2) = \mathfrak{m}_1$ and hence $\sigma_{L^H, f_\delta}^{i'}(x) \in \mathfrak{m}_1$. As $\sigma_{L^H, f_\delta}^{f_\delta}$ fixes $\delta_B A_{L^H, f_\delta}$, we may assume $i' \in \{0, \dots, f_\delta - 1\}$. As \mathfrak{m}_1 contains both y and $\sigma_{L^H, f_\delta}^{i'}(x)$, but not δ_B , the claim follows.

Let $I = yA_{L^H, f_\delta} + \sigma_{L^H, f_\delta}^{i'}(x)A_{L^H, f_\delta} \subsetneq \delta_B A_{L^H, f_\delta}$. Let δ_0 be the idempotent of R_K generating $I \cap R_K$. As $I \subsetneq \delta_B A_{L^H, f_\delta}$, we have $\delta_0 \delta = \delta_0 \neq \delta$. If $\delta_0 \neq 0$, we see it is used in Lines 24–26 to properly refine I_K . So assume $\delta_0 = 0$, or equivalently $I \cap R_K = \{0\}$. Let $x' = \sigma_{L^H, f_\delta}^{i'}(x)$. Then there exists $i_1, i_2 \in \mathbb{Z}$ such that

$$\sigma_{L, f_\delta}^{i_1}(g^{-1}y), \sigma_{L, f_\delta}^{i_1}(g^{-1}x'), \sigma_{L, f_\delta}^{i_2}(g'^{-1}y), \sigma_{L, f_\delta}^{i_2}(g'^{-1}x') \in \mathfrak{m}_0.$$

As $y = \delta_B \mu^\ell - s^{r'}$ and $x' = \sigma_{L^H, f_\delta}^{i'}(\gamma^k - t_{\delta, H}^r)$, we have

$$\sigma_{L, f_\delta}^{i_1}(g^{-1}(s^{r'})) \equiv \sigma_{L, f_\delta}^{i_2}(g'^{-1}(s^{r'})) \pmod{\mathfrak{m}_0} \quad (\text{C.10})$$

and

$$\sigma_{L, f_\delta}^{i_1}(g^{-1}(t_{\delta, H}^r)) \equiv \sigma_{L, f_\delta}^{i_2}(g'^{-1}(t_{\delta, H}^r)) \pmod{\mathfrak{m}_0}. \quad (\text{C.11})$$

As $\sigma_{L, f_\delta}(t_{\delta, H}) = \xi \cdot t_{\delta, H}$ and G commutes with σ_{L, f_δ} , (C.11) implies

$$\xi^{(i_1 - i_2)r} g^{-1}(t_{\delta, H}^r) \equiv g'^{-1}(t_{\delta, H}^r) \pmod{\mathfrak{m}_0}.$$

On the other hand, we know i is the unique integer in $\{0, \dots, f_\delta - 1\}$ satisfying (C.5). So $i_1 - i_2 \equiv i \pmod{f_\delta}$. Let $s' = \sigma_{L, f_\delta}^{i_2}(s)$. Then by (C.10), Lemma 5.22 and the fact that G commutes with σ_{L, f_δ} , we have

$$\sigma^i g^{-1}(s^{r'}) \equiv g'^{-1}(s^{r'}) \pmod{\mathfrak{m}_0}.$$

On the other hand, as σ_{L,f_δ} fixes $\mathfrak{m} = \frac{\Omega_0/p\mathcal{O}_L}{\text{Rad}(\mathcal{O}_L)}$ setwisely, (C.7) implies

$$\sigma^i g^{-1} s' + \mathfrak{m} = \sigma_{L,f_\delta}^{i_2}(c^{e_\delta}) \cdot (g'^{-1} s' + \mathfrak{m}).$$

It follows that $\sigma_{L,f_\delta}^{i_2}(c)$ is an e_δ th root of unity. So c is also an e_δ th root of unity, as in the case $e_\delta > 1$, $f_\delta = 1$. The same proof in the case $e_\delta > 1$, $f_\delta = 1$ then shows that there exists $j \in \{0, \dots, e_\delta - 1\}$ such that $Hg\sigma^{-i}\omega^{-j} \in \tilde{B}$, which contradicts the assumption $Hg\mathcal{D}_{\Omega_0} \notin \pi_H(\tilde{B})$. \square

Lemma 5.25. *Under GRH, there exists a subroutine RingHomTest that updates I_K in time polynomial in $\log p$ and the size of \mathcal{F} so that the partitions $C_H \in \mathcal{C}$ are refined. Moreover, at least one partition C_H is properly refined unless $\tilde{\mathcal{C}}$ is compatible and invariant.*

We need the following notation: suppose K, K' are extensions of K_0 and $\phi : K' \hookrightarrow K$ is an embedding of K' in K over K_0 . Recall that ϕ induces a homomorphism of \mathbb{F}_q -algebras $\hat{\phi} : \bar{\mathcal{O}}_{K'}/\text{Rad}(\bar{\mathcal{O}}_{K'}) \rightarrow \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$. Also suppose ψ is an embedding of \mathbb{F}_{q^i} in \mathbb{F}_{q^j} over \mathbb{F}_q where $i, j \in \mathbb{N}^+$. Then $\hat{\phi}$ and ψ determine a homomorphism of \mathbb{F}_q -algebras $A_{K',i} \rightarrow A_{K,j}$ sending $a \otimes b \in A_{K',i}$ to $\hat{\phi}(a) \otimes \psi(b) \in A_{K,j}$ for $a \in \bar{\mathcal{O}}_{K'}/\text{Rad}(\bar{\mathcal{O}}_{K'})$ and $b \in \mathbb{F}_{q^i}$. We denote this map by $\hat{\phi} \otimes_{\mathbb{F}_q} \psi$.

The pseudocode of the subroutine RingHomTest is given in Algorithm 20. It enumerates $(K, K') \in \mathcal{F}^2$, embeddings $\phi : K' \hookrightarrow K$ over K_0 , and $(\delta, \delta') \in I_K \times I_{K'}$ such that $\tilde{\phi}(\delta')\delta = \delta$. For each $(K, K', \phi, \delta, \delta')$, a set S of ideals of A_{K,f_δ} is computed. And for each $I \in S$, we find $\delta_0 \in I \cap R_K$ satisfying $(1 - \delta_0)(I \cap R_K) = \{0\}$, which is the unique idempotent of R_K that generates the ideal $I \cap R_K$ of R_K . If $\delta_0 \notin \{0, \delta\}$, we use δ_0 to refine I_K and return.

Fix $(K, K', \phi, \delta, \delta')$. The corresponding set S is computed as follows: Note we have $f_{\delta'} | f_\delta$ and $e_{\delta'} | e_\delta$. First assume $f_\delta > 1$. Compute the largest factor r of $q^{f_\delta} - 1$ coprime to f_δ . Then compute an element $\gamma \in \mathbb{F}_{q^{f_\delta}}^\times$ of order $(q^{f_\delta} - 1)/r$, which can be done efficiently assuming GRH. Call the subroutine SplitByExp in Lemma C.1 on $(A_{K,f_\delta}, f_\delta, \delta\gamma, \delta t_\delta^r)$ to obtain $x \in A_{K,f_\delta}$. Also perform the following computation if $f_{\delta'} > 1$: compute an embedding $\psi : \mathbb{F}_{q^{f_{\delta'}}} \rightarrow \mathbb{F}_{q^{f_\delta}}$ over \mathbb{F}_q deterministically in polynomial time using Lenstra's algorithm [Len91]. Compute $t = (\hat{\phi} \otimes_{\mathbb{F}_q} \psi)(t_{\delta'}) \in A_{K,f_\delta}$. By Definition 5.7 and the fact $\tilde{\phi}(\delta')\delta = \delta$, we have $\delta t^r A_{K,f_\delta} = \delta A_{K,f_\delta}$. Call the subroutine SplitByExp on $(A_{K,f_\delta}, f_\delta, \delta\gamma, \delta t^r)$ to obtain $x' \in A_{K,f_\delta}$. Then add the ideal $x' A_{K,f_\delta} + \sigma_{K,f_\delta}^i(x) A_{K,f_\delta}$ to S for all $i \in \{0, 1, \dots, f_\delta - 1\}$.

Algorithm 20 RingHomTest

```

1: for  $(K, K') \in \mathcal{F}^2$  do
2:   for embedding  $\phi : K' \hookrightarrow K$  over  $K_0$  do
3:     for  $(\delta, \delta') \in I_K \times I_{K'}$  satisfying  $\tilde{\phi}(\delta')\delta = \delta$  do
4:        $S \leftarrow \emptyset$ 
5:       if  $f_\delta > 1$  then
6:          $r \leftarrow$  the largest factor of  $q^{f_\delta} - 1$  coprime to  $f_\delta$ 
7:         compute  $\gamma \in \mathbb{F}_{q^{f_\delta}}^\times$  of order  $(q^{f_\delta} - 1)/r$ 
8:         call SplitByExp on  $(A_{K, f_\delta}, f_\delta, \delta\gamma, \delta t_\delta^r)$  to obtain  $x \in A_{K, f_\delta}$ 
9:         if  $f_{\delta'} > 1$  then
10:          compute an embedding  $\psi : \mathbb{F}_{q^{f_{\delta'}}} \rightarrow \mathbb{F}_{q^{f_\delta}}$  over  $\mathbb{F}_q$ 
11:           $t \leftarrow (\hat{\phi} \otimes_{\mathbb{F}_q} \psi)(t_{\delta'}) \in A_{K, f_\delta}$ 
12:          call SplitByExp on  $(A_{K, f_\delta}, f_\delta, \delta\gamma, \delta t^r)$  to obtain  $x' \in A_{K, f_\delta}$ 
13:          for  $i \leftarrow 0$  to  $f_\delta - 1$  do
14:             $S \leftarrow S \cup \{x' A_{K, f_\delta} + \sigma_{K, f_\delta}^i(x) A_{K, f_\delta}\}$ 
15:          if  $e_{\delta'} > 1$  then
16:             $J \leftarrow$  the preimage of  $(1 - \delta)(\bar{\mathcal{O}}_K / \text{Rad}(\bar{\mathcal{O}}_K))$  in  $\bar{\mathcal{O}}_K$ 
17:            compute  $u \in \text{Ann}_{\bar{\mathcal{O}}_K}(J^{e_\delta})$  such that  $u\bar{\phi}(s_{\delta'}) - s_\delta^{e_\delta/e_{\delta'}} \in J^{e_\delta/e_{\delta'}+1}$ 
18:             $\bar{u} \leftarrow u + \text{Rad}(\bar{\mathcal{O}}_K) \in \bar{\mathcal{O}}_K / \text{Rad}(\bar{\mathcal{O}}_K)$ 
19:             $r' \leftarrow$  the largest factor of  $q^{f_\delta} - 1$  coprime to  $e_\delta$ 
20:            compute  $\mu \in \mathbb{F}_{q^{f_\delta}}^\times$  of order  $(q^{f_\delta} - 1)/r'$ 
21:            call SplitByExp on  $(A_{K, f_\delta}, e_\delta, \delta\mu, \bar{u}^{r'})$  to obtain  $y \in A_{K, f_\delta}$ 
22:             $S \leftarrow S \cup \{y A_{K, f_\delta}\}$ 
23:          if  $f_\delta > 1$  then
24:            for  $i \leftarrow 0$  to  $f_\delta - 1$  do
25:               $S \leftarrow S \cup \{y A_{K, f_\delta} + \sigma_{K, f_\delta}^i(x) A_{K, f_\delta}\}$ 
26:          for  $I \in S$  do
27:            find  $\delta_0 \in I \cap R_K$  satisfying  $(1 - \delta_0)(I \cap R_K) = \{0\}$ 
28:            if  $\delta_0\delta \notin \{0, \delta\}$  then
29:               $I_K \leftarrow I_K - \{\delta\}$ 
30:               $I_K \leftarrow I_K \cup \{\delta_0\delta, (1 - \delta_0)\delta\}$ 
31:          return

```

If $e_{\delta'} > 1$, we perform the following computation: first compute the preimage J of $(1 - \delta)(\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K))$ under the quotient map $\bar{\mathcal{O}}_K \rightarrow \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$. Then J is the product of the maximal ideals \mathfrak{m} of $\bar{\mathcal{O}}_K$ satisfying $\delta \equiv 1 \pmod{\mathfrak{m}/\text{Rad}(\bar{\mathcal{O}}_K)}$. Compute $u \in \text{Ann}_{\bar{\mathcal{O}}_K}(J^{e_\delta})$ such that $u\bar{\phi}(s_{\delta'}) - s_\delta^{e_\delta/e_{\delta'}} \in J^{e_\delta/e_{\delta'}+1}$. We claim such u exists: by the Chinese remainder theorem, it suffices to show, for each maximal ideal \mathfrak{m} of $\bar{\mathcal{O}}_K$ containing J , that

$$u\bar{\phi}(s_{\delta'}) \equiv s_\delta^{e_\delta/e_{\delta'}} \pmod{\mathfrak{m}^{e_\delta/e_{\delta'}+1}}$$

has a solution in $\bar{\mathcal{O}}_K$. Fix such \mathfrak{m} . We have $s_\delta \in \mathfrak{m} - \mathfrak{m}^2$ by Lemma 5.17 and hence $s_\delta^{e_\delta/e_{\delta'}} \in \mathfrak{m}^{e_\delta/e_{\delta'}} - \mathfrak{m}^{e_\delta/e_{\delta'}+1}$. Let $\mathfrak{m}' = \bar{\phi}^{-1}(\mathfrak{m})$. By Lemma 5.17 and the fact $\bar{\phi}(\delta')\delta = \delta$, we have $s_{\delta'} \in \mathfrak{m}' - \mathfrak{m}'^2$ and hence $\bar{\phi}(s_{\delta'}) \in \mathfrak{m}^{e_\delta/e_{\delta'}} - \mathfrak{m}^{e_\delta/e_{\delta'}+1}$. The claim follows by noting $\mathfrak{m}^{e_\delta/e_{\delta'}}/\mathfrak{m}^{e_\delta/e_{\delta'}+1}$ is an one-dimensional vector space over $\bar{\mathcal{O}}_K/\mathfrak{m}$. Next compute

$$\bar{u} := u + \text{Rad}(\bar{\mathcal{O}}_K) \in \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K).$$

Compute the largest factor r' of $q^{f_\delta} - 1$ coprime to e_δ , so that all the prime factors of $(q^{f_\delta} - 1)/r'$ divide e_δ . And compute an element $\mu \in \mathbb{F}_{q^{f_\delta}}^\times$ of order $(q^{f_\delta} - 1)/r'$, which can be done efficiently assuming GRH. Note that $\bar{u}^{r'} A_{K,f_\delta} = \delta A_{K,f_\delta}$.⁵ Call the subroutine `SplitByExp` on the input $(A_{K,f_\delta}, e_\delta, \delta\mu, \bar{u}^{r'})$ to obtain $y \in A_{K,f_\delta}$, and add the ideal yA_{K,f_δ} to S . In addition, if $f_\delta > 1$, we enumerate $i = 0, 1, \dots, f_\delta - 1$, and for each i we add the ideal of A_{K,f_δ} generated by y and $\sigma_{K,f_\delta}^i(x)$ to S , where $x \in A_{K,f_\delta}$ is computed in the case $f_\delta > 1$ above.

Now we prove Lemma 5.25.

Proof of Lemma 5.25. Assume the algorithm does not properly refine any I_K . We prove that $\tilde{\mathcal{C}}$ is compatible and invariant. Fix $H, H' \in \mathcal{P}$ and a map $\phi^* : H \backslash G \rightarrow H' \backslash G$ that is either a projection $\pi_{H,H'}$ (with $H \subseteq H'$) or a conjugation $c_{H,h}$ (with $H' = hHh^{-1}$). Consider $g, g' \in G$ for which $Hg, Hg' \in H \backslash G$ are in the same block of $\tilde{\mathcal{C}}_H$. We want to show that $\phi^*(Hg), \phi^*(Hg') \in H' \backslash G$ are in the same block of $\tilde{\mathcal{C}}_{H'}$.

Let B be the block of C_H containing both $Hg\mathcal{D}_{\Omega_0}$ and $Hg'\mathcal{D}_{\Omega_0}$. Let $\bar{\phi}^* : H \backslash G/\mathcal{D}_{\Omega_0} \rightarrow H' \backslash G/\mathcal{D}_{\Omega_0}$ be the map $\pi_{H,H'}^{\mathcal{D}_{\Omega_0}}$ if $\phi^* = \pi_{H,H'}$, or $c_{H,h}^{\mathcal{D}_{\Omega_0}}$ if $\phi^* = c_{H,h}$. As \mathcal{C} is compatible and invariant, there exists $B' \in C_{H'}$ containing both $\bar{\phi}^*(Hg\mathcal{D}_{\Omega_0})$ and $\bar{\phi}^*(Hg'\mathcal{D}_{\Omega_0})$. Let K (resp. K') be the field in \mathcal{F} isomorphic to L^H (resp. $L^{H'}$)

⁵Again, we let $A_{K,f_\delta} = \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)$ if $f_\delta = 1$.

over K_0 . Let δ (resp. δ') be the idempotent in I_K (resp. $I_{K'}$) satisfying $\tilde{\tau}_H(\delta) = \delta_B$ (resp. $\tilde{\tau}_{H'}(\delta') = \delta_{B'}$). Let \mathfrak{m}_0 be an arbitrary maximal ideal of A_{L,f_δ} containing $\frac{\Omega_0/p\mathcal{O}_L}{\text{Rad}(\bar{\mathcal{O}}_L)}$. Fix an embedding $\psi : \mathbb{F}_{q^{f_{\delta'}}} \rightarrow \mathbb{F}_{q^{f_\delta}}$ over \mathbb{F}_q . Let $\phi : L^{H'} \hookrightarrow L^H$ be the natural inclusion if $\phi^* = \pi_{H,H'}$, or the map $x \mapsto h^{-1}x$ if $\phi^* = c_{H,h}$. Finally, let $s = \bar{\phi}(s_{\delta',H'})$ if $e_{\delta'} > 1$, and let $t = (\hat{\phi} \otimes_{\mathbb{F}_q} \psi)(t_{\delta',H'})$ if $f_{\delta'} > 1$.

We claim that the following two conditions are satisfied:

1. If $e_{\delta'} > 1$, the order of the unique element c in $\kappa_{\Omega_0}^\times$ satisfying

$$g^{-1}s + I' = c \cdot (g'^{-1}s + I')$$

is coprime to $e_{\delta'}$, where $I' = (\Omega_0/p\mathcal{O}_L)^{e(\Omega_0)/e_{\delta'}+1}$.

2. If $f_{\delta'} > 1$, the order of the unique element c in $(A_{L,f_\delta}/\mathfrak{m}_0)^\times$ satisfying

$$g^{-1}t + \mathfrak{m}_0 = c \cdot (g'^{-1}t + \mathfrak{m}_0)$$

is coprime to $f_{\delta'}$.

To see this claim implies that $\phi^*(Hg)$ and $\phi^*(Hg')$ are in the same block of $\tilde{C}_{H'}$, pick $\bar{g}, \bar{g}' \in G$ such that $H'\bar{g} = H'\phi^*(Hg)$ and $H'\bar{g}' = H'\phi^*(Hg')$. Then $c \in \kappa_{\Omega_0}^\times$ in the first condition is also the unique element satisfying $\bar{g}^{-1}s_{\delta',H'} + I' = c \cdot (\bar{g}'^{-1}s_{\delta',H'} + I')$. And $c \in (A_{L,f_\delta}/\mathfrak{m}_0)^\times$ in the second condition is also the unique element satisfying $\bar{g}^{-1}t_{\delta',H'} + \mathfrak{m}'_0 = c \cdot (\bar{g}'^{-1}t_{\delta',H'} + \mathfrak{m}'_0)$, where $\mathfrak{m}'_0 \supseteq \frac{\Omega_0/p\mathcal{O}_L}{\text{Rad}(\bar{\mathcal{O}}_L)}$ is the preimage of \mathfrak{m}_0 under $\text{id} \otimes_{\mathbb{F}_q} \psi : A_{L,f_{\delta'}} \rightarrow A_{L,f_\delta}$, and id is the identity map on $\bar{\mathcal{O}}_L/\text{Rad}(\bar{\mathcal{O}}_L)$. It follows by Definition 5.8 that $\phi^*(Hg)$ and $\phi^*(Hg')$ are in the same block assuming if these two conditions are satisfied.

The rest of the proof focuses on verifying the above two conditions. First assume $f_{\delta'} > 1$. Suppose $x = (\delta_B\gamma)^k - \delta_B t_{\delta,H}^r$ and $x' = (\delta_B\gamma)^{k'} - \delta_B t^r$ satisfy $x A_{L^H, f_\delta}, x' A_{L^H, f_\delta} \subsetneq \delta_B A_{L^H, f_\delta}$, where $k, k' \in \mathbb{N}$. Then there exists $i \in \{0, \dots, f_\delta - 1\}$ such that

$$I_i := x' A_{L^H, f_\delta} + \sigma_{L^H, f_\delta}^i(x) A_{L^H, f_\delta} \subsetneq \delta_B A_{L^H, f_\delta}.$$

This follows from the same argument in the proof of Lemma 5.24 that shows the existence of $i' \in \{0, \dots, f_\delta - 1\}$ satisfying (C.9). We may also assume $I_i \cap R_{L^H} = \{0\}$: otherwise, by identifying K with L^H using the isomorphism τ_H , we see the subroutine finds an idempotent δ_0 of R_K at Line 23 satisfying $\delta_0\delta \notin \{0, \delta\}$, and properly refines I_K .

By Lemma 5.6 and the assumption $I_i \cap R_{L^H} = \{0\}$, we know there exist $i_1, i_2 \in \mathbb{Z}$ such that

$$\sigma_{L, f_\delta}^{i_1}(g^{-1}x'), \sigma_{L, f_\delta}^{i_1+i}(g^{-1}x), \sigma_{L, f_\delta}^{i_2}(g'^{-1}x'), \sigma_{L, f_\delta}^{i_2+i}(g'^{-1}x) \in \mathfrak{m}_0. \quad (\text{C.12})$$

By Definition 5.7, there exist primitive f_δ th roots of unity $\xi, \xi' \in \mathbb{F}_{q^{f_\delta}}$ satisfying $\sigma_{L, f_\delta}(t_{\delta, H}) = \xi \cdot t_{\delta, H}$ and $\sigma_{L, f_\delta}(t) = \xi' \cdot t$. As $x = \delta_B \gamma^k - \delta_B t_{\delta, H}^r$ and $x' = \delta_B \gamma^{k'} - \delta_B t^r$, (C.12) implies

$$g^{-1}(t_{\delta, H}^r) \equiv \xi^{(i_2-i_1)r} g'^{-1}(t_{\delta, H}^r) \pmod{\mathfrak{m}_0} \quad (\text{C.13})$$

and

$$g^{-1}(t^r) \equiv \xi^{(i_2-i_1)r} g'^{-1}(t^r) \pmod{\mathfrak{m}_0}. \quad (\text{C.14})$$

On the other hand, as $Hg\mathcal{D}_{\Omega_0}, Hg'\mathcal{D}_{\Omega_0} \in B$, we know from Definition 5.8 that the order of the unique element $c \in (A_{L, f_\delta}/\mathfrak{m}_0)^\times$ satisfying $g^{-1}t_{\delta, H} + \mathfrak{m}_0 = c \cdot (g'^{-1}t_{\delta, H} + \mathfrak{m}_0)$ is coprime to f_δ . As r is coprime to f_δ , we see from (C.13) that $i_2 - i_1$ is divisible by f_δ . Then (C.14) becomes $g^{-1}(t^r) \equiv g'^{-1}(t^r) \pmod{\mathfrak{m}_0}$. So the order of the unique element $c \in (A_{L, f_\delta}/\mathfrak{m}_0)^\times$ satisfying $g^{-1}t + \mathfrak{m}_0 = c \cdot (g'^{-1}t + \mathfrak{m}_0)$ is coprime to f_δ , as desired.

Next assume $e_{\delta'} > 1$. Let \bar{u} be the element computed at Line 18 and regard it as an element of $\bar{\mathcal{O}}_{L^H}/\text{Rad}(\bar{\mathcal{O}}_{L^H})$ by identifying K with L^H using the isomorphism τ_H . Suppose $y = (\delta_B \mu)^\ell - \bar{u}^{r'}$ satisfies $yA_{L^H, f_\delta} \subsetneq \delta_B A_{L^H, f_\delta}$, where $\ell \in \mathbb{N}$. We may assume $yA_{L^H, f_\delta} \cap R_{L^H} = \{0\}$, since otherwise the idempotent δ_0 generating $yA_{L^H, f_\delta} \cap R_{L^H}$ satisfies $\delta_0 \delta \notin \{0, \delta\}$ and is used to properly refine I_K .

If $e_{\delta'} > 1$ and $f_\delta = 1$, the ring A_{L^H, f_δ} is just $\bar{\mathcal{O}}_{L^H}/\text{Rad}(\bar{\mathcal{O}}_{L^H})$, and we have $y = 0$ in this case. So $g^{-1}y = g'^{-1}y = 0$, which implies

$$g^{-1}(\bar{u}^{r'}) \equiv g'^{-1}(\bar{u}^{r'}) \pmod{\frac{\bar{\Omega}_0/p\mathcal{O}_L}{\text{Rad}(\bar{\mathcal{O}}_L)}}. \quad (\text{C.15})$$

Let $c_1, c_2 \in \kappa_{\bar{\Omega}_0}^\times$ be the residues of $g^{-1}\bar{u}$ and $g'^{-1}\bar{u}$ modulo $\frac{\bar{\Omega}_0/p\mathcal{O}_L}{\text{Rad}(\bar{\mathcal{O}}_{L^H})}$ respectively. Then $(c_2/c_1)^{r'} = 1$. So the order of c_2/c_1 divides r' , which is coprime to e_δ . By Definition 5.8, the order of the unique element $c \in \kappa_{\bar{\Omega}_0}^\times$ satisfying

$$g^{-1}s_{\delta, H} + I = c \cdot (g'^{-1}s_{\delta, H} + I)$$

is coprime to e_δ (and hence to $e_{\delta'}$), where $I = (\bar{\Omega}_0/p\mathcal{O}_L)^{e(\bar{\Omega}_0)/e_\delta+1}$. Then we have

$$g^{-1}s_{\delta, H}^{e_\delta/e_{\delta'}} + I' = c^{e_\delta/e_{\delta'}} \cdot \left(g'^{-1}s_{\delta, H}^{e_\delta/e_{\delta'}} + I' \right) \quad (\text{C.16})$$

and the order of $c^{e_\delta/e_{\delta'}}$ is coprime to $e_{\delta'}$. By the definition of \bar{u} , we may rewrite (C.16) as

$$c_1(g^{-1}s) + I' = c_2c^{e_\delta/e_{\delta'}} \cdot (g'^{-1}s + I').$$

As the order of c_2/c_1 and that of $c^{e_\delta/e_{\delta'}}$ are coprime to $e_{\delta'}$, the second condition above is satisfied.

Finally, assume $e_{\delta'} > 1$ and $f_\delta > 1$. Then there exists $j \in \{0, \dots, f_\delta - 1\}$ such that

$$I'_j := yA_{L^H, f_\delta} + \sigma_{L^H, f_\delta}^i(x)A_{L^H, f_\delta} \subsetneq \delta_B A_{L^H, f_\delta},$$

where $x = \delta_B \gamma^k - \delta_B t_{\delta, H}^r$ is as above. Again we may assume $I'_j \cap R_{L^H} = \{0\}$ since otherwise I_K is properly refined. Then there exist $i_1, i_2 \in \mathbb{Z}$ such that

$$\sigma_{L, f_\delta}^{i_1}(g^{-1}y), \sigma_{L, f_\delta}^{i_1+j}(g^{-1}x), \sigma_{L, f_\delta}^{i_2}(g'^{-1}y), \sigma_{L, f_\delta}^{i_2+j}(g'^{-1}x) \in \mathfrak{m}_0.$$

As $x = \delta_B \gamma^k - \delta_B t_{\delta, H}^r$ and $\sigma_{L, f_\delta}(t_{\delta, H}) = \xi \cdot t_{\delta, H}$, again we conclude that $i_2 - i_1$ is divisible by f_δ . As the order of σ_{L^H, f_δ} on $\delta_B A_{L^H, f_\delta}$ is f_δ , we may assume $i_1 = i_2$. As $y = (\delta_B \mu)^\ell - \bar{u}^{r'}$, we have

$$\sigma_{L, f_\delta}^{i_1}(g^{-1}(\bar{u}^{r'})) \equiv \sigma_{L, f_\delta}^{i_1}(g'^{-1}(\bar{u}^{r'})) \pmod{\mathfrak{m}_0}.$$

As σ_{L, f_δ} fixes every maximal ideal of $\bar{\mathcal{O}}_L/\text{Rad}(\bar{\mathcal{O}}_L)$ setwisely, and

$$\mathfrak{m}_0 \cap (\bar{\mathcal{O}}_L/\text{Rad}(\bar{\mathcal{O}}_L)) = \frac{\mathfrak{Q}_0/p\mathcal{O}_L}{\text{Rad}(\bar{\mathcal{O}}_L)},$$

we see (C.15) still holds. The rest of the proof is the same as in the case $e_{\delta'} > 1$, $f_\delta = 1$. □

Lemma 5.26. *Let G be a finite group acting transitively on a set S . Let \mathcal{D} be a subgroup of G and let k be the number of \mathcal{D} -orbits in S . Suppose $k > 1$. Let $\ell \in \mathbb{N}^+$ be the least prime factor of k . Let $\mathcal{P} = \mathcal{P}_m$ be the system of stabilizers of depth m for some $m \geq \ell$ (with respect to the action of G on S). Then for any $x \in S$ and any \mathcal{P} -scheme of double cosets \mathcal{C} with respect to \mathcal{D} that is homogeneous on G_x , there exists no antisymmetric $(\mathcal{C}, \mathcal{D})$ -separated \mathcal{P} -scheme.*

Proof. Assume to the contrary that there exist $x \in S$, a \mathcal{P} -scheme of double cosets $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ with respect to \mathcal{D} that is homogeneous on G_x , and an antisymmetric $(\mathcal{C}, \mathcal{D})$ -separated \mathcal{P} -scheme $\tilde{\mathcal{C}} = \{\tilde{C}_H : H \in \mathcal{P}\}$. As G acts transitively on S , we know \mathcal{C} is homogeneous on G_x for all $x \in S$.

Fix $x_0 \in S$ and consider the bijection $\lambda_{x_0} : S \rightarrow G_{x_0} \backslash G$ sending ${}^g x_0$ to $G_{x_0} g^{-1}$. It is an equivalence between the action of G on S and that on $G_{x_0} \backslash G$ by inverse right translation. Let B_0 be a block of $\tilde{C}_{G_{x_0}}$ and define $T := \lambda_{x_0}^{-1}(B_0) \subseteq S$. As \mathcal{C} is homogeneous on G_{x_0} and $\tilde{\mathcal{C}}$ is $(\mathcal{C}, \mathcal{D})$ -separated, we know T is a complete set of representatives of the \mathcal{D} -orbits in S , and hence $|B_0| = |T| = k$.

The group G acts diagonally on $S^{(\ell)}$. And $\text{Sym}(\ell)$ acts on $S^{(\ell)}$ by permuting the coordinates. As the two actions commute, we know $\text{Sym}(\ell)$ permutes the G -orbits in $S^{(\ell)}$. Fix $z \in T^{(\ell)}$ and let H_z be the subgroup of $\text{Sym}(\ell)$ fixing Gz setwisely. Using the bijection $\lambda_z : Gz \rightarrow G_z \backslash G$, the action of H_z on Gz induces an action on $G_z \backslash G$. In the proof of Lemma 2.18, we showed that the latter action induces a semiregular action on the set of the blocks of \tilde{C}_{G_z} .

Let $U_z := T^{(\ell)} \cap Gz$. Suppose $z = (z_1, \dots, z_\ell)$. For $g \in G$, the element ${}^g z$ is in U_z iff $\lambda_{x_0}({}^g z_i) \in \lambda_{x_0}(T) = B_0$ for all $i \in [\ell]$. Fix $i \in [\ell]$ and choose $g_i \in G$ satisfying ${}^{g_i} x_0 = z_i$. Then $c_{x_0, g_i} : G_{x_0} \backslash G \rightarrow G_{z_i} \backslash G$ sends B_0 to a block $B_i \in \tilde{C}_{G_{z_i}}$. Also note that $c_{x_0, g_i} \circ \lambda_{x_0} = \lambda_{z_i}$. So $\lambda_{x_0}({}^g z_i) \in B_0$ is equivalent to $\lambda_{z_i}({}^g z_i) \in B_i$. As

$$\lambda_{z_i}({}^g z_i) = G_{z_i} g^{-1} = \pi_{G_z, G_{z_i}}(G_z g^{-1}) = \pi_{G_z, G_{z_i}} \circ \lambda_z({}^g z),$$

we see that $\lambda_z(U_z)$ consists of the elements $x \in G_z \backslash G$ satisfying $\pi_{G_z, G_{z_i}}(x) \in B_i$ for all $i \in [\ell]$. By compatibility of $\tilde{\mathcal{C}}$, the set $\lambda_z(U_z)$ is a disjoint union of blocks of \tilde{C}_{G_z} . Moreover, by regularity of $\tilde{\mathcal{C}}$, the cardinality of these blocks are all divisible by $|B_0| = k$.

Note that the action of H_z on Gz fixes the set U_z setwisely. So the semiregular action of H_z on the set of the blocks of \tilde{C}_{G_z} restricts to a semiregular action on the subset of the blocks in $\lambda_z(U_z)$. By the previous paragraph, we know $|U_z|$ is a multiple of $k|H_z|$.

The set $T^{(\ell)}$ is a disjoint union of subsets of the form U_z where $z \in T^{(\ell)}$. The group $\text{Sym}(\ell)$ permutes these subsets. By the orbit-stabilizer theorem, each $\text{Sym}(\ell)$ -orbit O is a disjoint union of $|\text{Sym}(\ell)|/|H_z|$ subsets of the same cardinality $|U_z|$, where z is an arbitrary element in O . So

$$|O| = \frac{|\text{Sym}(\ell)|}{|H_z|} \cdot |U_z|$$

which is a multiple of $k\ell!$ by the previous paragraph. It follows that $|T^{(\ell)}| = k(k-1) \cdots (k-\ell+1)$ is a multiple of $k\ell!$. But this is impossible since none of the factors $k-1, \dots, k-\ell+1$ are divisible by the prime number ℓ . \square

Appendix D

LIST OF ALGORITHMS

Name	Reference	Page
ComputeQuotientRing	Lemma 3.8	56
ComputeResidue	Lemma 3.9	57
ComputeEmbeddings	Lemma 3.10	58
ComputeRingHom	Lemma 3.11	58
ExtractFactors	Algorithm 1, Theorem 3.7	60
ComputePscheme	Algorithm 2, Theorem 3.8	63
CompatibilityAndInvarianceTest	Algorithm 3, Lemma 3.13	65
FreeModuleTest	Lemma 3.16	67
SplitByZeroDivisor	Lemma 3.17	67
RegularityTest	Algorithm 4, Lemma 3.14	68
StrongAntisymmetryTest	Algorithm 5, Lemma 3.15	70
PschemeAlgorithm	Algorithm 6, Theorem 3.9	75
Automorphism	Algorithm 17, Lemma 3.18	241

Table D.1: Algorithms and subroutines in the \mathcal{P} -scheme algorithm

Name	Reference	Page
AdjoinRoot	Lemma 4.8	86
SplittingField	Algorithm 7, Lemma 4.9	87
Stabilizers	Algorithm 8, Lemma 4.10	88
Tower	Theorem 4.1	90
GeneralAction	Algorithm 9, Theorem 4.2	92
SubgroupSystem	Algorithm 16, Lemma 8.1	205

Table D.2: Algorithms for constructing number fields

Name	Reference	Page
ComputeRelEmbeddings	Lemma 4.7	86
ComputeRings	Lemma 5.7	116
ComputeRingHoms	Lemma 5.8	116
ExtractFactorsV2	Algorithm 10, Theorem 5.6	119
ComputeDoubleCosetPscheme ¹	Algorithm 11, Theorem 5.7	123
CompatibilityAndInvarianceTestV2	Lemma 5.12	124
RegularityTestV2	Lemma 5.13	125
StrongAntisymmetryTestV2	Lemma 5.14	125
RamificationIndexTest	Algorithm 12, Lemma 5.15	126
InertiaDegreeTest	Algorithm 13, Lemma 5.16	128
ComputeOrdinaryPscheme	Algorithm 14, Theorem 5.8	137
GeneralizedPschemeAlgorithm	Algorithm 15, Theorem 5.9	139
ComputeAdvice	Algorithm 18, Lemma 5.17	250
SplitByExp	Lemma C.1	253
SurjectivityTest	Algorithm 19, Lemma 5.24	255
RingHomTest	Algorithm 20, Lemma 5.25	261

Table D.3: Algorithms and subroutines in the generalized \mathcal{P} -scheme algorithm

¹The subroutine `ComputeDoubleCosetPscheme` is not actually used in the generalized \mathcal{P} -scheme algorithm, but only serves as a preliminary version of `ComputeOrdinaryPscheme`.

BIBLIOGRAPHY

- [AMM77] L. Adleman, K. Manders, and G. Miller. “On taking roots in finite fields”. In: *Proceedings of the 18th Annual Symposium on Foundations of Computer Science*. 1977, pp. 175–178.
- [Aro13] M. Arora. “Extensibility of association schemes and GRH-based deterministic polynomial factoring”. PhD thesis. Universitäts-und Landesbibliothek Bonn, 2013.
- [Aro+14] M. Arora, G. Ivanyos, M. Karpinski, and N. Saxena. “Deterministic polynomial factoring and association schemes”. In: *LMS Journal of Computation and Mathematics* 17.01 (2014), pp. 123–140.
- [Asc00] M. Aschbacher. *Finite Group Theory*. Cambridge University Press, 2000.
- [AG84] M. Aschbacher and R. Guralnick. “Some applications of the first cohomology group”. In: *Journal of Algebra* 90.2 (1984), pp. 446–460.
- [AM69] M. F. Atiyah and I. G. Macdonald. *Introduction to Commutative Algebra*. Addison-Wesley, 1969.
- [BCP82] L. Babai, P. J. Cameron, and P. P. Pálffy. “On the orders of primitive groups with restricted nonabelian composition factors”. In: *Journal of Algebra* 79.1 (1982), pp. 161–168.
- [Bai04] R. A. Bailey. *Association Schemes: Designed Experiments, Algebra and Combinatorics*. Vol. 84. Cambridge University Press, 2004.
- [BI84] E. Bannai and T. Ito. *Algebraic Combinatorics*. Benjamin/Cummings, 1984.
- [Ben05] C. Benbenisty. “On actions of primitive groups”. PhD thesis. Hebrew University, 2005.
- [Ber67] E. R. Berlekamp. “Factoring polynomials over finite fields”. In: *Bell System Technical Journal* 46.8 (1967), pp. 1853–1859.
- [Ber70] E. R. Berlekamp. “Factoring polynomials over large finite fields”. In: *Mathematics of Computation* 24.111 (1970), pp. 713–735.
- [Bha+17] V. Bhargava, G. Ivanyos, R. Mittal, and N. Saxena. “Irreducibility and r -th root finding over finite fields”. In: *arXiv preprint arXiv:1702.00558* (2017).
- [BKS15] J. Bourgain, S. Konyagin, and I. Shparlinski. “Character sums and deterministic polynomial root finding in finite fields”. In: *Mathematics of Computation* 84.296 (2015), pp. 2969–2977.

- [BCN89] A. E. Brouwer, A. M. Cohen, and A. Neumaier. *Distance-Regular Graphs*. Springer-Verlag, 1989.
- [Bur07] T. C. Burness. “On base sizes for actions of finite classical groups”. In: *Journal of the London Mathematical Society* 75.3 (2007), pp. 545–562.
- [BGS11] T. C. Burness, R. M. Guralnick, and J. Saxl. “On base sizes for symmetric groups”. In: *Bulletin of the London Mathematical Society* 43.2 (2011), pp. 386–391.
- [BLS09] T. C. Burness, M. W. Liebeck, and A. Shalev. “Base sizes for simple groups and a conjecture of Cameron”. In: *Proceedings of the London Mathematical Society* 98.1 (2009), pp. 116–162.
- [BOW10] T. C. Burness, E. A. O’Brien, and R. A. Wilson. “Base sizes for sporadic simple groups”. In: *Israel Journal of Mathematics* 177.1 (2010), pp. 307–333.
- [BS15] T. Burness and Á. Seress. “On Pyber’s base size conjecture”. In: *Transactions of the American Mathematical Society* 367.8 (2015), pp. 5633–5651.
- [Cam92] P. J. Cameron. “Some open problems on permutation groups”. In: *Groups, Combinatorics and Geometry* 165 (1992), pp. 340–351.
- [CGS78] P. J. Cameron, J.-M. Goethals, and J. J. Seidel. “The Krein condition, spherical designs, Norton algebras and permutation groups”. In: *Indagationes Mathematicae (Proceedings)*. Vol. 81. 1. 1978, pp. 196–206.
- [CK93] P. J. Cameron and W. M. Kantor. “Random permutations: some group-theoretic aspects”. In: *Combinatorics, Probability and Computing* 2.03 (1993), pp. 257–262.
- [CZ81] D. G. Cantor and H. Zassenhaus. “A new algorithm for factoring polynomials over finite fields”. In: *Mathematics of Computation* 36.154 (1981), pp. 587–592.
- [CH00] Q. Cheng and M. A. Huang. “Factoring polynomials over finite fields and stable colorings of tournaments”. In: *Proceedings of the 4th Algorithmic Number Theory Symposium*. 2000, pp. 233–245.
- [Chi89] A. L. Chistov. “The complexity of constructing the ring of integers of a global field”. In: *Doklady Akademii Nauk SSSR*. Vol. 306. 5. 1989, pp. 1063–1067.
- [Coh93] H. Cohen. *A Course in Computational Algebraic Number Theory*. Springer, 1993.
- [Con+85] J. H. Conway, R. T. Curtis, S. P. Norton, and R. A. Parker. *ATLAS of Finite Groups*. Oxford University Press, 1985.

- [Coo78] B. N. Cooperstein. “Minimal degree for a permutation representation of a classical group”. In: *Israel Journal of Mathematics* 30.3 (1978), pp. 213–235.
- [DM96] J. D. Dixon and B. Mortimer. *Permutation Groups*. Vol. 163. Springer, 1996.
- [DHM16] H. Duyan, Z. Halasi, and A. Maróti. “A proof of Pyber’s base size conjecture”. In: *arXiv preprint arXiv:1611.09487* (2016).
- [ER63] P. Erdős and A. Rényi. “Asymmetric graphs”. In: *Acta Mathematica Hungarica* 14.3-4 (1963), pp. 295–315.
- [Evd92] S. A. Evdokimov. “Factorization of solvable polynomials over finite fields and the generalized Riemann hypothesis”. In: *Journal of Soviet Mathematics* 59.3 (1992), pp. 842–849.
- [Evd94] S. A. Evdokimov. “Factorization of polynomials over finite fields in subexponential time under GRH”. In: *Proceedings of the 1st Algorithmic Number Theory Symposium*. 1994, pp. 209–219.
- [Faw13] J. B. Fawcett. “The base size of a primitive diagonal group”. In: *Journal of Algebra* 375 (2013), pp. 302–321.
- [FT63] W. Feit and J. Thompson. “Solvability of groups of odd order”. In: *Pacific Journal of Mathematics* 13.3 (1963), pp. 775–787.
- [FR85] K. Friedl and L. Rónyai. “Polynomial time solutions of some problems of computational algebra”. In: *Proceedings of the 17th Annual ACM Symposium on Theory of Computing*. 1985, pp. 153–162.
- [Gao01] S. Gao. “On the deterministic complexity of factoring polynomials”. In: *Journal of Symbolic Computation* 31.1 (2001), pp. 19–36.
- [Gat87] J. von zur Gathen. “Factoring polynomials and primitive elements for special primes”. In: *Theoretical Computer Science* 52.1 (1987), pp. 77–89.
- [GG13] J. von zur Gathen and J. Gerhard. *Modern Computer Algebra*. Cambridge University Press, 2013.
- [GP01] J. von zur Gathen and D. Panario. “Factoring polynomials over finite fields: A survey”. In: *Journal of Symbolic Computation* 31.1-2 (2001), pp. 3–17.
- [GS92] J. von zur Gathen and V. Shoup. “Computing Frobenius maps and factoring polynomials”. In: *Computational Complexity* 2.3 (1992), pp. 187–224.
- [GM98] D. Gluck and K. Magaard. “Base sizes and regular orbits for coprime affine permutation groups”. In: *Journal of the London Mathematical Society* 58.3 (1998), pp. 603–618.

- [GSS98] D. Gluck, Á. Seress, and A. Shalev. “Bases for primitive permutation groups and a conjecture of Babai”. In: *Journal of Algebra* 199.2 (1998), pp. 367–378.
- [GLS94] D. Gorenstein, R. Lyons, and R. Solomon. *The Classification of the Finite Simple Groups*. American Mathematical Society, 1994.
- [Gua09] Y. Guan. “Factoring polynomials and Grobner bases”. PhD thesis. Clemson University, 2009.
- [Har77] R. Hartshorne. *Algebraic Geometry*. Springer, 1977.
- [Hig67] D. G. Higman. “Intersection matrices for finite permutation groups”. In: *Journal of Algebra* 6.1 (1967), pp. 22–42.
- [Hua84] M. A. Huang. “Factorization of polynomials over finite fields and factorization of primes in algebraic number fields”. In: *Proceedings of the 16th Annual ACM Symposium on Theory of Computing*. 1984, pp. 175–182.
- [Hua91a] M. A. Huang. “Factorization of polynomials over finite fields and decomposition of primes in algebraic number fields”. In: *Journal of Algorithms* 12.3 (1991), pp. 482–489.
- [Hua91b] M. A. Huang. “Generalized Riemann hypothesis and factoring polynomials over finite fields”. In: *Journal of Algorithms* 12.3 (1991), pp. 464–481.
- [ILF97] V. V. Ishkhanov, B. B. Lur’e, and D. K. Faddeev. *The Embedding Problem in Galois Theory*. American Mathematical Society, 1997.
- [Iva+12] G. Ivanyos, M. Karpinski, L. Rónyai, and N. Saxena. “Trading GRH for algebra: algorithms for factoring polynomials and related structures”. In: *Mathematics of Computation* 81.277 (2012), pp. 493–531.
- [IKS09] G. Ivanyos, M. Karpinski, and N. Saxena. “Schemes for deterministic polynomial factoring”. In: *Proceedings of the International Symposium on Symbolic and Algebraic Computation*. 2009, pp. 191–198.
- [KS98] E. Kaltofen and V. Shoup. “Subquadratic-time factoring of polynomials over finite fields”. In: *Mathematics of Computation* 67.223 (1998), pp. 1179–1197.
- [KU11] K. S. Kedlaya and C. Umans. “Fast polynomial factorization and modular composition”. In: *SIAM Journal on Computing* 40.6 (2011), pp. 1767–1802.
- [KL90] P. B. Kleidman and M. W. Liebeck. *The Subgroup Structure of the Finite Classical Groups*. Vol. 129. Cambridge University Press, 1990.
- [KM00] J. Klüners and G. Malle. “Explicit Galois realization of transitive groups of degree up to 15”. In: *Journal of Symbolic Computation* 30.6 (2000), pp. 675–716.

- [Knu98] D. Knuth. *The Art of Computer Programming, Volume Two, Seminumerical Algorithms*. Addison-Wesley, 1998.
- [KP00] L. G. Kovács and C. E. Praeger. “On minimal faithful permutation representations of finite groups”. In: *Bulletin of the Australian Mathematical Society* 62.02 (2000), pp. 311–317.
- [LMO79] J. C. Lagarias, H. L. Montgomery, and A. M. Odlyzko. “A bound for the least prime ideal in the Chebotarev density theorem”. In: *Inventiones Mathematicae* 54.3 (1979), pp. 271–296.
- [Lan84] S. Landau. “Polynomial time algorithms for Galois groups”. In: *Proceedings of the International Symposium on Symbolic and Algebraic Computation*. 1984, pp. 225–236.
- [Lan85] S. Landau. “Factoring polynomials over algebraic number fields”. In: *SIAM Journal on Computing* 14.1 (1985), pp. 184–195.
- [LM85] S. Landau and G. L. Miller. “Solvability by radicals is in polynomial time”. In: *Journal of Computer and System Sciences* 30.2 (1985), pp. 179–208.
- [Lan02] S. Lang. *Algebra*. Springer, 2002.
- [Len90] H. W. Lenstra Jr. “Algorithms for finite fields”. In: *Number Theory and Cryptography*. 1990, pp. 76–85.
- [Len91] H. W. Lenstra Jr. “Finding isomorphisms between finite fields”. In: *Mathematics of Computation* (1991), pp. 329–347.
- [Len92] H. W. Lenstra Jr. “Algorithms in algebraic number theory”. In: *Bulletin of the American Mathematical Society* 26.2 (1992), pp. 211–244.
- [Len83] A. K. Lenstra. “Factoring polynomials over algebraic number fields”. In: *Computer Algebra* (1983), pp. 245–254.
- [LLL82] A. K. Lenstra, H. W. Lenstra Jr., and L. Lovász. “Factoring polynomials with rational coefficients”. In: *Mathematische Annalen* 261.4 (1982), pp. 515–534.
- [LPS88] M. W. Liebeck, C. E. Praeger, and J. Saxl. “On the O’Nan-Scott theorem for finite primitive permutation groups”. In: *Journal of the Australian Mathematical Society (Series A)* 44.03 (1988), pp. 389–396.
- [LS99] M. W. Liebeck and A. Shalev. “Simple groups, permutation groups, and probability”. In: *Journal of the American Mathematical Society* 12.2 (1999), pp. 497–520.
- [LS02] M. W. Liebeck and A. Shalev. “Bases of primitive linear groups”. In: *Journal of Algebra* 252.1 (2002), pp. 95–113.
- [LS14] M. W. Liebeck and A. Shalev. “Bases of primitive linear groups II”. In: *Journal of Algebra* 403 (2014), pp. 223–228.

- [LS03] A. Lubotzky and D. Segal. *Subgroup Growth*. Birkhäuser, 2003.
- [Luk82] E. M. Luks. “Isomorphism of graphs of bounded valence can be tested in polynomial time”. In: *Journal of Computer and System Sciences* 25.1 (1982), pp. 42–65.
- [Luk93] E. M. Luks. “Permutation groups and polynomial-time computation”. In: *Groups and Computation*. Vol. 11. DIMACS series in Discrete Mathematics and Theoretical Computer Science. 1993, p. 139.
- [Mar77] D. A. Marcus. *Number Fields*. Springer, 1977.
- [Mil83] G. L. Miller. “Isomorphism of k -contractible graphs. A generalization of bounded valence and bounded genus”. In: *Information and Control* 56.1 (1983), pp. 1–20.
- [Mum99] D. Mumford. *The Red Book of Varieties and Schemes*. Springer, 1999.
- [Neu99] J. Neukirch. *Algebraic Number Theory*. Springer-Verlag, 1999.
- [Neu63] B. H. Neumann. “Twisted wreath products of groups”. In: *Archiv der Mathematik* 14.1 (1963), pp. 1–6.
- [Pál82] P. P. Pálffy. “A polynomial bound for the orders of primitive solvable groups”. In: *Journal of Algebra* 77.1 (1982), pp. 127–137.
- [Pil90] J. Pila. “Frobenius maps of abelian varieties and finding roots of unity in finite fields”. In: *Mathematics of Computation* 55.192 (1990), pp. 745–763.
- [Pra90] C. E. Praeger. “The inclusion problem for finite primitive permutation groups”. In: *Proceedings of the London Mathematical Society* 3.1 (1990), pp. 68–88.
- [PLN97] C. E. Praeger, C. H. Li, and A. C. Niemeyer. “Finite transitive permutation groups and finite vertex-transitive graphs”. In: *Graph Symmetry*. Springer, 1997, pp. 277–318.
- [Pyb93] L. Pyber. “Asymptotic results for permutation groups”. In: *Groups and Computation*. Vol. 11. DIMACS series in Discrete Mathematics and Theoretical Computer Science. 1993, pp. 197–219.
- [PS97] L. Pyber and A. Shalev. “Asymptotic results for primitive permutation groups”. In: *Journal of Algebra* 188.1 (1997), pp. 103–124.
- [RB72] K. B. Reid and E. Brown. “Doubly regular tournaments are equivalent to skew Hadamard matrices”. In: *Journal of Combinatorial Theory, Series A* 12.3 (1972), pp. 332–338.
- [Rón88] L. Rónyai. “Factoring polynomials over finite fields”. In: *Journal of Algorithms* 9.3 (1988), pp. 391–400.
- [Rón89] L. Rónyai. “Factoring polynomials modulo special primes”. In: *Combinatorica* 9.2 (1989), pp. 199–206.

- [Rón90] L. Rónyai. “Computing the structure of finite algebras”. In: *Journal of Symbolic Computation* 9.3 (1990), pp. 355–373.
- [Rón92] L. Rónyai. “Galois groups and factoring polynomials over finite fields”. In: *SIAM Journal on Discrete Mathematics* 5.3 (1992), pp. 345–365.
- [Sch85] R. Schoof. “Elliptic curves over finite fields and the computation of square roots mod p ”. In: *Mathematics of Computation* 44.170 (1985), pp. 483–494.
- [SS98] K. See and S. Y. Song. “Association schemes of small order”. In: *Journal of Statistical Planning and Inference* 73.1 (1998), pp. 225–271.
- [Ser96] Á. Seress. “The minimal base size of primitive solvable permutation groups”. In: *Journal of the London Mathematical Society* 53.2 (1996), pp. 243–255.
- [Ser03] Á. Seress. *Permutation Group Algorithms*. Vol. 152. Cambridge University Press, 2003.
- [Sha54] I. R. Shafarevich. “Construction of fields of algebraic numbers with given solvable Galois group”. In: *Izvestiya Rossiiskoi Akademii Nauk, Seriya Matematicheskaya* 18.6 (1954), pp. 525–578.
- [Sho90] V. Shoup. “On the deterministic complexity of factoring polynomials over finite fields”. In: *Information Processing Letters* 33.5 (1990), pp. 261–267.
- [Sho91] V. Shoup. “Smoothness and factoring polynomials over finite fields”. In: *Information Processing Letters* 38.1 (1991), pp. 39–42.
- [Sup76] D. A. Suprunenko. *Matrix Groups*. 45. American Mathematical Society, 1976.
- [Uma08] C. Umans. “Fast polynomial factorization and modular composition in small characteristic”. In: *Proceedings of the 40th Annual ACM Symposium on Theory of Computing*. 2008, pp. 481–490.
- [Wae91] B. L. van der Waerden. *Algebra, Vols. I, II*. Springer, 1991.
- [WR76] P. J. Weinberger and L. P. Rothschild. “Factoring polynomials over algebraic number fields”. In: *ACM Transactions on Mathematical Software (TOMS)* 2.4 (1976), pp. 335–350.
- [Wie64] H. Wielandt. *Finite Permutation Groups*. Academic Press, 1964.
- [Yun76] D. Y. Y. Yun. “On square-free decomposition algorithms”. In: *Proceedings of the 3rd ACM Symposium on Symbolic and Algebraic Computation*. 1976, pp. 26–35.

LIST OF NOTATIONS

- \mathbb{N}^+ . set of positive integers.
 $[k]$. set $\{1, 2, \dots, k\}$.
 $A - B$. set difference $\{x : x \in A \text{ and } x \notin B\}$.
 $|S|$. cardinality of S .
 \log . logarithmic function with base 2.
 $\coprod S_i$. disjoint union of sets S_i .
 0_S . coarsest partition of a set S .
 ∞_S . finest partition of a set S .
 $S^{(k)}$. set of k -tuples of distinct elements from S .
 $f \circ g$. composition of functions f and g , from right to left.
 $\text{char}(K)$. characteristic of a field K .
 $\text{gcd}(f, g)$. greatest common divisor of polynomials f and g .
 e . identity element of a group.
 gH . left coset $\{gh : h \in H\}$.
 Hg . right coset $\{hg : h \in H\}$.
 G/H . left coset space $\{gH : g \in G\}$.
 $H \backslash G$. right coset space $\{Hg : g \in G\}$.
 HgK . double coset $\{hgh' : h \in H, h' \in K\}$.
 $H \backslash G / K$. double coset space $\{HgK : g \in G\}$.
 $[G : H]$. index of a subgroup H in G .
 $\langle H_1, \dots, H_k \rangle$. join of subgroups H_1, \dots, H_k .
 $\langle g_1, \dots, g_k \rangle$. subgroup generated by g_1, \dots, g_k .
 $H \trianglelefteq G$. H is a normal subgroup of G .
 $N_G(H)$. normalizer of H in G .
 $Z(G)$. center of G .

- $(a_1 a_2 \cdots a_n)$. permutation sending a_i to a_{i+1} for $1 \leq i < n$ and a_n to a_1 .
- $\text{Sym}(S), \text{Sym}(n)$. symmetric group.
- $\text{Alt}(S), \text{Alt}(n)$. alternating group.
- $\text{Aut}(G)$. automorphism group of a group G .
- $\text{Inn}(G)$. inner automorphism group of a group G .
- $\text{Out}(G)$. outer automorphism group of a group G .
- ${}^g x$. alias for the element $\varphi(g, x)$ where φ is a group action.
- ${}^g T$. set $\{{}^g x : x \in T\}$.
- Gx . G -orbit $\{{}^g x : g \in G\}$ of an element x .
- G_x . stabilizer of an element x .
- G_T . pointwise stabilizer of a set T .
- $G_{\{T\}}$. setwise stabilizer of a set T .
- G_{x_1, \dots, x_k} . pointwise stabilizer of $\{x_1, \dots, x_k\}$.
- S^G . set of fixed points of G in a set S .
- A^G . subgroup (resp. subring, subfield) of G -invariant elements of the abelian group (resp. ring, field) A .
- \mathcal{P}_m . system of stabilizers of depth m .
- λ_x . the map from a G -orbit S containing x to $G_x \backslash G$ sending ${}^g x$ to $G_x g^{-1}$.
- $\pi_{H, H'}$. projection from $H \backslash G$ to $H' \backslash G$.
- $c_{H, g}$. conjugation from $H \backslash G$ to $gHg^{-1} \backslash G$.
- $d(G), d'(G)$. See Definition 2.8.
- $b(G)$. minimal base size of a permutation group G .
- π_i^k . projection from $S^{(k)}$ to $S^{(k-1)}$ omitting the k th coordinate.
- π_T^k . projection from $S^{(k)}$ to $S^{(k-1)}$ omitting the coordinates with indices in T .
- c_g^k . permutation of $S^{(k)}$ sending x to ${}^g x$.
- $\Pi(\mathcal{C})$. m -scheme constructed from a \mathcal{P} -scheme \mathcal{C} (see Definition 2.12).
- $\mathcal{C}(\Pi)$. \mathcal{P} -scheme constructed from an m -scheme Π (see Definition 2.13).
- 1_S . block $\{(x, x) : x \in S\}$ of an association scheme on S .

- $\mathcal{C}_{g',g''}^g$. See Definition 2.15.
- $\Pi(P)$. 3-collection constructed from a partition P (see Definition 2.16).
- $P(\Pi)$. partition constructed from a 3-collection Π (see Definition 2.16).
- $(x), xR, Rx$. ideal of a ring R generated by x .
- \mathcal{O}_K . ring of integers of a number field K .
- $\text{Aut}(K/K_0)$. automorphism group of a field extension K/K_0 .
- $\text{Gal}(K/K_0)$. Galois group of a Galois extension K/K_0 .
- $\text{Gal}(f/K_0)$. Galois group of K/K_0 where K is the splitting field of f over K_0 .
- $\bar{\mathcal{O}}_K$. quotient ring $\mathcal{O}_K/p\mathcal{O}_K$.
- $i_{K,L}$. inclusion $\bar{\mathcal{O}}_K \hookrightarrow \bar{\mathcal{O}}_L$ in Chapter 3, or $R_K \hookrightarrow R_L$ in Chapter 5.
- $P(I)$. See Definition 3.2 and Definition 5.4.
- $I(P)$. See Definition 3.2 and Definition 5.4.
- δ_B . See Definition 3.2 and Definition 5.4.
- B_δ . See Lemma 3.5 and Lemma 5.4.
- \mathcal{O}'_K . a p -maximal order of K .
- $\bar{\phi}$. ring homomorphism $\bar{\mathcal{O}}_K \rightarrow \bar{\mathcal{O}}_{K'}$ induced from an embedding $\phi : K \hookrightarrow K'$.
- \mathcal{P}^\sharp . poset of subfields corresponding to \mathcal{P} via Galois correspondence.
- τ_H . fixed isomorphism $K \rightarrow L^H$ (which is K_0 -linear in Chapter 5).
- \cong_{K_0} . isomorphism over a field K_0 .
- $c(\mathcal{P})$. complexity of a subgroup system \mathcal{P} .
- $\|\alpha\|$. greatest absolute value of $i(\alpha)$ where i ranges over embeddings $\mathbb{Q}(\alpha) \hookrightarrow \mathbb{C}$.
- \mathcal{P}_+ . subgroup system $\{U : H \subseteq U \subseteq N_G(H), H \in \mathcal{P}\}$.
- $\tilde{h}, \psi_0, A_0, K_0$. See Chapter 5.
- $\kappa_{\mathfrak{P}}$. residue field of \mathfrak{P} .
- $e(\mathfrak{P})$. ramification index of \mathfrak{P} over $p\mathcal{O}_{K_0}$.
- $f(\mathfrak{P})$. inertia degree of \mathfrak{P} over $p\mathcal{O}_{K_0}$.
- $\mathcal{D}_{\mathfrak{P}}$. decomposition group of \mathfrak{P} over K_0 .

- $\mathcal{I}_{\mathfrak{P}}$. inertia group of \mathfrak{P} over K_0 .
- $\mathcal{W}_{\mathfrak{P}}$. wild inertia group of \mathfrak{P} over K_0 .
- $e(Hg\mathcal{D})$. ramification index of a double coset $Hg\mathcal{D}$.
- $f(Hg\mathcal{D})$. inertia degree of a double coset $Hg\mathcal{D}$.
- $\text{Rad}(A)$. radical of a ring A .
- $\text{Rad}(g)$. radical of a polynomial g .
- R_K . ring $\{x \in \bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K) : x^p = x\}$.
- $\pi_{H,H'}^{\mathcal{D}}$. projection from $H \backslash G/\mathcal{D}$ to $H' \backslash G/\mathcal{D}$.
- $c_{H,g}^{\mathcal{D}}$. conjugation from $H \backslash G/\mathcal{D}$ to $gHg^{-1} \backslash G/\mathcal{D}$.
- $A \otimes_{\mathbb{F}_q} B$. tensor product of A and B over \mathbb{F}_q .
- $A_{K,i}$. ring $(\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K)) \otimes_{\mathbb{F}_q} \mathbb{F}_{q^i}$.
- $\sigma_{K,i}$. automorphism of $A_{K,i}$ sending $a \otimes b$ to $a^q \otimes b$.
- $\text{Ann}_R(S)$. annihilator of S in R .
- $\hat{\phi}$. ring homomorphism $\bar{\mathcal{O}}_K/\text{Rad}(\bar{\mathcal{O}}_K) \rightarrow \bar{\mathcal{O}}_{K'}/\text{Rad}(\bar{\mathcal{O}}_{K'})$ induced from $\phi : K \hookrightarrow K'$.
- $\tilde{\phi}$. ring homomorphism $R_K \rightarrow R_{K'}$ induced from $\phi : K \hookrightarrow K'$.
- e_δ, f_δ . See Section 5.7.
- $\mathcal{P}|_H$. restriction of a subgroup system \mathcal{P} to H .
- $\mathcal{C}|_H$. restriction of a \mathcal{P} -collection \mathcal{C} to H .
- $\Pi|_{x_1, \dots, x_k}$. See Definition 6.3.
- \mathcal{P}_{cl} . closure of a subgroup system \mathcal{P} .
- $\Pi||_T$. restriction of an m -collection Π to a subset T .
- $\text{AGL}(V)$. general affine group on V .
- $G \wr G'$. wreath product of groups G and G' .
- $d_{\text{Sym}}(n)$. alias for $d(G)$ where $G = \text{Sym}(S)$ acts naturally on S , $|S| = n$.
- $\text{GL}(V), \text{GL}_n(q)$. general linear group.
- $\Gamma\text{L}(V), \Gamma\text{L}_n(q)$. general semilinear group.
- $\text{PGL}(V), \text{PGL}_n(q)$. projective linear group.

$\mathrm{P}\Gamma\mathrm{L}(V)$, $\mathrm{P}\Gamma\mathrm{L}_n(q)$. projective semilinear group.

$d_{\mathrm{GL}}(n, q)$, $d_{\mathrm{TL}}(n, q)$, $d_{\mathrm{PGL}}(n, q)$, $d_{\mathrm{P}\Gamma\mathrm{L}}(n, q)$. See Definition 7.2.

$m(n)$, $m'(n)$. See Definition 7.3.

$\mathcal{P}_{G,N}$. See Definition 8.1.

$\mathrm{soc}(G)$. socle of a finite group G .

$\mathrm{Hol}(G)$. holomorph of a group G .

$\mathrm{Map}(S, T)$. set of all maps from the set S to the set T .

$T \mathrm{twr}_{\varphi} P$. twisted wreath product with respect to the data (T, P, φ) .

$U(k, q)$. fully deleted permutation module for $\mathrm{Sym}(k)$ over \mathbb{F}_q .

$G_1 \otimes \cdots \otimes G_k$. tensor product of the linear groups G_i .

$g_1 \otimes \cdots \otimes g_k$. image of (g_1, \dots, g_k) in $G_1 \otimes \cdots \otimes G_k$.

INDEX

A

action, 12
affine scheme, 122, 235
affine type, 209, 215
algebraic integer, 52
algebraic number, 52
almost simple group, 209
almost simple type, 209, 212
annihilator, 115
antisymmetry
 of a \mathcal{P} -scheme, 19
 of a \mathcal{P} -scheme of double cosets,
 113
 of an m -scheme, 23
association scheme, 30
automorphism group
 of a field extension, 52
 of a group, 11

B

base, 21
basis of a free module, 51
block, 10

C

$(\mathcal{C}, \mathcal{D})$ -separatedness, 131
center of a group, 11
CFSG, 201
Chinese remainder theorem, 50
classical group, 201
classification of finite simple groups,
 see CFSG
closure, 99, 162
compatibility
 of a \mathcal{P} -collection, 18, 235

 of a \mathcal{P} -collection of double
 cosets, 112
 of an m -collection, 22
complete
 factorization, 1
 idempotent decomposition, 51
completely reducible, 44
complexity of a subgroup system, 81
composition factor, 201
composition series, 201
conjugate
 subfield, 52
 subgroup, 11
conjugation, 17, 112
coprime ideals, 50
cyclotomic scheme, 39

D

decomposition group, 107
Dedekind domain, 53
depth, 15
diagonal action, 12
diagonal type, 210, 223
direct product
 of \mathcal{P} -collections, 175
 of m -schemes, 176
discreteness
 of a \mathcal{P} -scheme, 19
 of a \mathcal{P} -scheme of double cosets,
 113
 of an m -scheme, 23
discriminant of a polynomial, 96
double coset, 11
doubly regular tournament, 39

E

equivalent actions, 16
 exceptional group of Lie type, 201

F

faithful action, 12
 finitely generated module, 51
 fixed
 subfield, 13
 subring, 13
 flat module, 114, 247
 free module, 51
 Frobenius automorphism, 107
 fully deleted permutation module,
 215
 fundamental theorem of Galois
 theory, 52

G

Galois closure, 52
 Galois extension, 52
 Galois group, 52
 Galois theory, 52
 general
 affine group, 173
 linear group, 181
 semilinear group, 181
 generalized \mathcal{P} -scheme algorithm, 139
 generalized Riemann hypothesis, *see*

GRH

G -invariant, 13
 G -module, 13
 G -orbit, *see* orbit
 great common divisor, 11
 GRH, 3
 group action, *see* action
 G -set, 12

H

having a singleton, 143
 holomorph, 210
 homogeneity
 of a \mathcal{P} -scheme, 19
 of a \mathcal{P} -scheme of double cosets,
 113
 of an m -scheme, 23

I

\mathcal{I} -advice, 130
 ideal, 49
 idempotent, 45, 50
 idempotent decomposition, 45, 51
 imprimitive
 m -scheme, 171
 permutation group, 89
 imprimitive wreath product action,
 179
 index of a subgroup, 11
 induction of a \mathcal{P} -scheme, 154
 inertia degree
 of a double coset, 109
 of a prime ideal, 106
 inertia group, 107
 inner automorphism, 11
 integral, 52
 invariance
 of a \mathcal{P} -collection, 18, 235
 of a \mathcal{P} -collection of double
 cosets, 112
 of an m -collection, 22
 inverse right translation, 16
 of \mathcal{P} -schemes, 148
 irreducible lifted polynomial, 5, 102
 irreducible linear group, 215

J

join, 11
 join-closed, 166
 Jordan-Hölder theorem, 201

K

k -transitive, 12
 $(k + 1/2)$ -transitive, 12

L

left coset, 11
 left translation, 16
 Lie rank, 206
 lifted polynomial, 5, 102
 linear group, 182
 locally constant

- inertia degrees, 109
- ramification indices, 109

 lying over, 53, 106

M

matching, 29
 maximal ideal, 49
 maximal subgroup, 11
 m -collection, 22
 minimal base size, 21
 monic polynomial, 11
 m -scheme, 22

N

natural action of linear groups, 182,
 194
 nilpotent, 109
 nilradical, *see* radical of a ring
 non-standard action, 213
 normal subgroup, 11
 normalization, 123
 normalizer, 11

number field, 52

O

O’Nan-Scott theorem, 212
 Odd Order Theorem, 174
 orbit, 12

- \mathcal{P} -scheme, 33
- m -scheme, 33

 orbit-stabilizer theorem, 17
 orbital graph, 173
 orthogonal idempotents, 45, 50
 outer automorphism group, 12

P

Paley tournament, 39
 partially ordered set, *see* poset
 partition, 10
 \mathcal{P} -collection, 18

- of double cosets, 112

 permutation group, 12
 permutation isomorphic actions, 16
 permutation representation, 12
 p -maximal order, 56
 pointwise stabilizer, 12
 poset, 15
 power

- non-residue, 3
- residue, 3

 primary tensor, 217
 prime ideal, 49
 primitive

- m -scheme, 171
- association scheme, 172
- idempotent, 50
- linear group, 215
- permutation group, 89

 primitive element, 52

- primitive element theorem, 52, 84
 - primitive wreath product action, 211
 - product type, 211, 226
 - projection, 17, 112
 - projective
 - linear group, 182
 - semilinear group, 182
 - space, 182
 - proper
 - factorization, 1, 142
 - ideal, 49
 - idempotent decomposition, 51
 - refinement, 10
 - \mathcal{P} -scheme, 18, 235
 - of double cosets, 112
 - \mathcal{P} -scheme algorithm, 75
 - Pyber's base size conjecture, 213
- Q**
- quasisimple group, 216
- R**
- radical
 - of a polynomial, 110
 - of a ring, 109
 - ramification index
 - of a double coset, 109
 - of a prime ideal, 106
 - rank of a module, 51
 - refinement, 10
 - regular action, 12
 - regularity
 - of a \mathcal{P} -collection, 18, 235
 - of a \mathcal{P} -collection of double cosets, 113
 - of a prime, 122
 - of an m -collection, 23
 - relative number field, 79
 - residue field, 106
 - restriction
 - of a \mathcal{P} -collection, 149
 - of a partition, 10
 - of a subgroup system, 148
 - of an m -collection to a subset, 168
 - right coset, 11
 - ring of integers, 52
- S**
- scalar linear transformation, 182
 - schemes conjecture, 147, 159
 - for permutation groups, 147, 160
 - Schreier conjecture, 212
 - self-reduction of discreteness, 186
 - semilinear transformation, 181
 - semiregular action, 12
 - semisimple
 - algebra, 50
 - ring, 50
 - set of fixed points, 12
 - set of imprimitivity, 89
 - setwise stabilizer, 12
 - simple group, 201
 - size, 55
 - socle, 209
 - splitting field, 52
 - splitting of prime ideals, 53, 106
 - complete splitting, 53
 - sporadic simple group, 201
 - square-free, 44
 - square-free factorization, 44, 115
 - stabilizer, 12
 - standard action, 181, 189, 213
 - strong antisymmetry

- of a \mathcal{P} -scheme, 20
- of a \mathcal{P} -scheme of double cosets, 113
- of an m -scheme, 23
- structure constants, 3, 55
- subfield system, 61, 80
- subgroup system, 15
- subquotient, 11
- symmetry
 - of a \mathcal{P} -scheme, 19
 - of an m -scheme, 23
- system of stabilizers, 15

T

- tensor product of linear groups, 215

- transitive action, 12
- twisted wreath type, 212, 228

U

- unique factorization domain, 1, 11

V

- valency, 30

W

- wild inertia group, 108

- wreath product

- of \mathcal{P} -schemes, 178

- of m -schemes, 179

- of groups, 177