

Deterministic polynomial factoring over finite fields: a uniform approach via \mathcal{P} -schemes

Zeyu Guo

Department of Computer Science and Engineering, IIT Kanpur¹

Abstract

We introduce a family of combinatorial objects called \mathcal{P} -schemes, where \mathcal{P} is a collection of subgroups of a finite group G . A \mathcal{P} -scheme is a collection of partitions of 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 as well as m -schemes [IKS09].

We apply the theory of \mathcal{P} -schemes to deterministic polynomial factoring over finite fields: suppose $\tilde{f}(X) \in \mathbb{Z}[X]$ and a prime number p are given, such that $f(X) := \tilde{f}(X) \bmod p$ factorizes into $n = \deg(\tilde{f})$ distinct linear factors over the finite field \mathbb{F}_p . We show that, assuming the generalized Riemann hypothesis (GRH), $f(X)$ can be completely factorized in deterministic polynomial time if the Galois group G of $\tilde{f}(X)$ is an almost simple primitive permutation group on the set of roots of $\tilde{f}(X)$, and the socle of G is a subgroup of $\text{Sym}(k)$ for k up to $2^{O(\sqrt{\log n})}$. This is the first deterministic polynomial-time factoring algorithm for primitive Galois groups of superpolynomial order.

We prove our result by developing a generic factoring algorithm and analyzing it using \mathcal{P} -schemes. We also show that the main results achieved by known GRH-based deterministic polynomial factoring algorithms can be derived from our generic algorithm in a uniform way.

Finally, we investigate the *schemes conjecture* in [IKS09], and formulate analogous conjectures associated with various families of permutation groups. We show that these conjectures form a hierarchy of relaxations of the original schemes conjecture, and their positive resolutions would imply deterministic polynomial-time factoring algorithms for various families of Galois groups under GRH.

Keywords: polynomial factoring, permutation group, finite field, algebraic combinatorics

1. Introduction

We are interested in the problem of deterministic univariate polynomial factoring over finite fields: given a univariate monic polynomial $f(X) \in \mathbb{F}_q[X]$ over a finite field \mathbb{F}_q , our goal is to deterministically compute the *complete factorization* of f over \mathbb{F}_q , i.e., the factorization $f(X) = \prod_{i=1}^k f_i(X)$, where each $f_i(X)$ is a monic irreducible factor of $f(X)$ over \mathbb{F}_q . We are also interested in the more moderate goal of deterministically computing a *proper factorization* of $f(X)$, i.e., factoring $f(X)$ into more than one (possibly reducible) factors over \mathbb{F}_q .

¹Part of this work was done while the author was at Department of Computing and Mathematical Sciences, Caltech.

Email address: zguo@cse.iitk.ac.in (Zeyu Guo)

This manuscript has been published in the Journal of Symbolic Computation. DOI: 10.1016/j.jsc.2019.02.011

© 2019. This manuscript version is made available under the CC-BY-NC-ND 4.0 license. <http://creativecommons.org/licenses/by-nc-nd/4.0/>

1.1. Previous work

As a fundamental problems in computer algebra, univariate polynomial factoring over finite fields has been extensively studied over the years. A polynomial-time factoring algorithm is required to factorize a degree- n monic polynomial $f(X) \in \mathbb{F}_q[X]$ in time polynomial in $O(n \log q)$. If randomness is allowed, such algorithms are well known [Ber70, CZ81, vzGS92, KS98, Uma08, KU11]. On the other hand, despite much effort, finding a *deterministic* polynomial-time algorithm remains a long-standing open problem. Berlekamp [Ber67] gave the first deterministic algorithm for the general case, whose running time is polynomial in n and q (instead of n and $\log q$). All currently known (unconditional) deterministic algorithms take exponential time [Ber67, Ber70, Sho90, BKS15]. In particular, the papers [Sho90, BKS15] achieve the running time $O((n \log q)^c p^{1/2})$, where $p = \text{char}(\mathbb{F}_q)$ and $c > 0$ is a constant. The $p^{1/2}$ -dependence on the characteristic p remains the best known [BKS15]. Faster algorithms are known in some special cases [vzG87, Rón89, Sho91, Sch85, Pil90, IKRS12].

A lot more is known if one accepts the generalized Riemann hypothesis (GRH). Assuming GRH, the work [AMM77] gave a deterministic polynomial-time algorithm factorizing polynomials of the form $X^n - a \in \mathbb{F}_p[X]$, where a has an n -th root in \mathbb{F}_p . Several GRH-based deterministic algorithms were proposed since then. In one line of research [Hua91a, Hua91b, Evd92, Rón92], the finite field over which $f(X)$ is defined is assumed to be a prime field \mathbb{F}_p , and a *lifted polynomial* of $f(X)$ is assumed to be given, i.e., a monic polynomial $\tilde{f}(X) \in \mathbb{Z}[X]$ satisfying $\tilde{f}(X) \bmod p = f(X)$. In particular, Huang [Hua91a, Hua91b] proved that $f(X) \in \mathbb{F}_p[X]$ can be deterministically factorized in polynomial time under GRH if the Galois group G of $\tilde{f}(X)$ is abelian. This was generalized in [Evd92] to the case that G is solvable. For a general Galois group G , the work [Rón92] proposed a deterministic algorithm under GRH that runs in time polynomial in $|G|$ and the size of the input. In general, however, the cardinality of G may be as large as $n!$, 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 $\text{poly}(n^n, \log q)$ under GRH. Building on Rónyai's work, Evdokimov [Evd94] showed that the problem can be solved in quasipolynomial time by presenting a deterministic $\text{poly}(n^{\log n}, \log q)$ -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 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, AIKS14]. 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 of $f(X)$, 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* called the *schemes conjecture* was proposed in [IKS09], whose affirmative resolution would imply a polynomial-time factoring algorithm under GRH.

1.2. Our results

Our main result extends the Galois-theoretic approach [Hua91a, Hua91b, Evd92, Rón92]. For simplicity, in this paper we always make the following assumption about the input polynomial $f(X)$:

Assumption 1. $f(X)$ is a monic polynomial defined over a prime field \mathbb{F}_p that factorizes completely into distinct linear factors over \mathbb{F}_p .

This is considered to be the most difficult case in the literature, and there exists a standard deterministic polynomial-time reduction that reduces the general factoring problem to this special case [Ber70, Yun76].

Lifted polynomial. Like the results in [Hua91a, Hua91b, Evd92, Rón92], our algorithm uses a *lifted polynomial* of $f(X)$, which is defined to be a monic polynomial $\tilde{f}(X) \in \mathbb{Z}[X]$ satisfying $\tilde{f}(X) \bmod p = f(X)$. Note that given $f(X)$, we can always choose $\tilde{f}(X)$ efficiently.

Main result. We give a deterministic polynomial factoring algorithm for the case that the Galois group of $\tilde{f}(X)$ is an *almost simple primitive permutation group* on the set of roots of $\tilde{f}(X)$. Recall that a finite group G is *almost simple* if it has a non-abelian simple subgroup H such that $H \leq G \leq \text{Aut}(H)$ [DM96]. The subgroup H is unique and is called the *socle* of G . And a permutation group G on a set S is *primitive* if there is no G -invariant partition of S other than the finest partition and the coarsest partition [DM96], where a partition P is G -invariant if $P = {}^g P := \{gU : U \in P\}$ for $g \in G$.

We state our main result as follows.

Theorem 1.1. *Under GRH, there exists a deterministic algorithm that, given $f(X) \in \mathbb{F}_p[X]$ of degree n satisfying Assumption 1 and a lifted polynomial $\tilde{f}(X) \in \mathbb{Z}[X]$ such that the Galois group G of $\tilde{f}(X)$ is an almost simple primitive permutation group on the set of roots of $\tilde{f}(X)$, and the socle of G is (isomorphic to) a subgroup of $\text{Sym}(k)$ for some $k \in \mathbb{N}^+$, computes the complete factorization of $f(X)$ over \mathbb{F}_p in time polynomial in $k^{\log k}$ and the size of the input. In particular, the algorithm runs in polynomial time for $k = 2^{O(\sqrt{\log n})}$.*

Theorem 1.1 gives the first deterministic factoring algorithm that runs in polynomial time for primitive Galois groups of superpolynomial order. For example, consider the special case that $G = \text{Sym}(k)$ and $k = 2^{\Theta(\sqrt{\log n})}$, and suppose the action of G on the set of n roots of $\tilde{f}(X)$ is equivalent to its action on the set of t -subsets of $[k]$ (induced from the natural action of G on $[k]$), where $t = \Theta(\sqrt{\log n})$ and $n = \binom{k}{t}$. In this case, our algorithm runs in polynomial time even though $|G| = k!$ is doubly exponential in $\sqrt{\log n}$. This should be compared with Evdokimov’s algorithm [Evd94] and Rónyai’s algorithm [Rón92]. The former takes quasipolynomial time, and the latter takes time polynomial in $|G|$ and the size of the input.

Theorem 1.1 is proved using a “generic” factoring algorithm developed in this paper, which also provides a general framework for GRH-based deterministic polynomial factoring. In particular, we show that the main results achieved by known GRH-based deterministic factoring algorithms, including those in the Galois-theoretic approach [Hua91a, Hua91b, Evd92, Rón92] and those in the combinatorial approach [Rón88, Evd94, IKS09], can be derived from our generic factoring algorithm in a uniform way.

Remark. While we only address the special case of almost simple primitive permutation groups in this paper, we remark that it is typically considered to be the most difficult case for many problems on permutation groups. Indeed, a common framework for proofs in permutation group theory is first reducing the general problem to the almost simple primitive case (via the O’Nan-Scott Theorem [LPS88]), and then further analyzing this special case (which often involves the classification of finite simple groups). See [Asc08, Page 2] for an explanation of this framework. Our work is motivated by such reductions. The extension of our result to general permutation groups is the subject of future publications, and the analysis presented in this paper may serve as an important ingredient.

1.3. Our techniques

To prove our results, we introduce a family of combinatorial objects called \mathcal{P} -schemes. Then we develop a generic factoring algorithm that can be analyzed using \mathcal{P} -schemes.

\mathcal{P} -schemes. Given a finite group G and a collection \mathcal{P} of subgroups of G , a \mathcal{P} -scheme is a set of partitions $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ satisfying a list of axioms, where for $H \in \mathcal{P}$, C_H is a partition of the right coset space $H \backslash G = \{Hg : g \in G\}$. The formal definition is given in Definition 3.5.

\mathcal{P} -schemes are naturally related to the classical notion of *association schemes* [BI84] in algebraic combinatorics and the notion of *m-schemes* studied in [IKS09, Aro13, AIKS14]. It was shown in [IKS09]

that an m -scheme, when $m = 3$, gives rise to an association scheme and vice versa. Our notion of \mathcal{P} -schemes further generalize m -schemes. Such connection will be further discussed in Subsection 3.2.

Our definition of \mathcal{P} -schemes can be seen as both combinatorial and group-theoretic: it involves a finite group G , which is absent in the definition of m -schemes and that of association schemes. This allows us to apply tools from finite group theory and permutation group theory to study properties of \mathcal{P} -schemes. On the other hand, \mathcal{P} -schemes are even more general than m -schemes: an m -scheme corresponds to a \mathcal{P} -scheme with $G = \text{Sym}(n)$ and a special collection \mathcal{P} of subgroups. In general, one can choose various G and \mathcal{P} and study the corresponding \mathcal{P} -schemes. Such generality is crucially used in our approach to deterministic polynomial factoring, and it may also find applications to other problems.

A generic factoring algorithm. We design a generic factoring algorithm that connects \mathcal{P} -schemes with deterministic polynomial factoring, which we explain now.

Denote by L the splitting field of $\tilde{f}(X)$ over \mathbb{Q} and by G the Galois group $\text{Gal}(L/\mathbb{Q})$. One step of our algorithm is a subroutine that uses $\tilde{f}(X)$ to construct a collection \mathcal{F} of subfields of L . This step is the generic part of our algorithm, as there are multiple choices of \mathcal{F} . Once \mathcal{F} is given, we associate with \mathcal{F} a collection \mathcal{P} of subgroups of G , defined by

$$\mathcal{P} = \{H \leq G : L^H \cong K \text{ for some } K \in \mathcal{F}\}.$$

We will discuss \mathcal{P} -schemes with respect to this collection \mathcal{P} .

The main theorem on the generic factoring algorithm depends on two properties of \mathcal{P} -schemes called *strong antisymmetry* and *discreteness*, which are defined in Subsection 3.1. The theorem states as follows:

Theorem 1.2. *Let \mathcal{I} be a family of instances of the input $(f(X), \tilde{f}(X))$ where $f(X)$ satisfies Assumption 1. Suppose there exists a deterministic algorithm that given an instance $s = (f(X), \tilde{f}(X)) \in \mathcal{I}$, constructs in time $T(s)$ a collection \mathcal{F} of number fields that are isomorphic to subfields of the splitting field L of $\tilde{f}(X)$ over \mathbb{Q} , such that*

- $\mathbb{Q}[X]/(\tilde{f}_i(X)) \in \mathcal{F}$ for all monic irreducible factors $\tilde{f}_i(X)$ of $\tilde{f}(X)$ over \mathbb{Q} , and
- all strongly antisymmetric \mathcal{P} -schemes are discrete on $\text{Gal}(L/\mathbb{Q}(\alpha)) \in \mathcal{P}$ for all roots α of $\tilde{f}(X)$ in L , where \mathcal{P} is the collection of subgroups of $G = \text{Gal}(L/\mathbb{Q})$ associated with \mathcal{F} .

Then under GRH, there exists a deterministic algorithm that given $s = (f(X), \tilde{f}(X)) \in \mathcal{I}$, outputs the complete factorization of $f(X)$ over \mathbb{F}_p in time polynomial in $T(s)$ and the size of s .

For the more moderate goal of computing a proper factorization of $f(X)$, we have a similar theorem that requires a weaker property of \mathcal{P} -schemes. See Theorem 4.24.

The hypothetical algorithm in Theorem 1.2 is generally not difficult to implement. The real challenge is to prove that the corresponding collection \mathcal{P} satisfies the second condition of the theorem. Thus Theorem 1.2 reduces the problem of deterministic polynomial factoring to a combinatorial problem about \mathcal{P} -schemes. By choosing various G and \mathcal{P} and verifying the condition, we are able to recover the main results of known factoring algorithms [Hua91a, Hua91b, Rón88, Rón92, Evd92, Evd94, IKS09] in a uniform way.

Remark. For simplicity, we always assume $f(X)$ is defined over a prime field \mathbb{F}_p and completely factorizes into distinct linear factors over \mathbb{F}_p . There exists a standard reduction to this special case [Ber70, Yun76]. Alternatively, we can directly generalize Theorem 1.2 so that it holds for a general polynomial. Such a generalized algorithm was developed in the author's Ph.D. thesis [Guo17, Chapter 5]. Finally, note that the assumption that $f(X)$ and $\tilde{f}(X)$ are monic is justified by substitution of variables.²

²Suppose $f(X) \in \mathbb{F}_p[X]$ and $\tilde{f}(X) \in \mathbb{Z}[X]$ are (possibly non-monic) polynomials of the same degree n satisfying $\tilde{f}(X) \pmod{p} = f(X)$. Let c be the leading coefficient of $\tilde{f}(X)$ and let $\bar{c} = c \pmod{p} \in \mathbb{F}_p$. Then $f'(X) := \bar{c}^{n-1} \cdot f(X/\bar{c}) \in \mathbb{F}_p[X]$ and $\tilde{f}'(X) := c^{n-1} \cdot \tilde{f}(X/c) \in \mathbb{Z}[X]$ are monic. This gives a reduction to the monic case.

Proving Theorem 1.1. Finally, we derive Theorem 1.1 from Theorem 1.2. This is achieved by first analyzing a special kind of actions of symmetric groups called *standard actions*. Then we apply Liebeck and Shalev’s result [LS99] that almost simple primitive permutation groups in non-standard actions have bounded minimal base size. A technical lemma called the *self-reduction lemma* also plays a key role in our proof.

1.4. Overview of the generic factoring algorithm

We give an overview of the algorithm in Theorem 1.2. For simplicity, assume $\tilde{f}(X)$ is irreducible over \mathbb{Q} . Let $F = \mathbb{Q}(\alpha)$, where α is an arbitrary root of $\tilde{f}(X)$ in its splitting field. For a number field K , denote by \mathcal{O}_K the ring of integers of K .

Reducing to the problem of computing an idempotent decomposition. It is well known that computing a (proper) factorization of $f(X)$ is equivalent to finding a nonzero zero divisor of the ring $\mathbb{F}_p[X]/(f(X))$ [Rón88, Evd94, IKS09]. We focus on finding zero divisors that are *idempotents*, i.e., those elements x satisfying $x^2 = x$. More specifically, we reduce the problem of factoring $f(X)$ to the problem of finding a set I of (nonzero) mutually orthogonal idempotents of the ring $\tilde{\mathcal{O}}_F := \mathcal{O}_F/p\mathcal{O}_F$ satisfying $\sum_{x \in I} x = 1$. We call such a set I an *idempotent decomposition* of $\tilde{\mathcal{O}}_F$.

Computing idempotent decompositions. Let L be the splitting field of $\tilde{f}(X)$ over \mathbb{Q} , and let $G = \text{Gal}(L/\mathbb{Q})$. By [Rón92], a nontrivial idempotent decomposition of $\tilde{\mathcal{O}}_F$ can be found efficiently if an efficiently computable nontrivial automorphism of $\tilde{\mathcal{O}}_F$ is given. When F is Galois over \mathbb{Q} , the Galois group G naturally provides nontrivial automorphisms of $\tilde{\mathcal{O}}_F$, which can be efficiently computed thanks to known efficient factoring algorithms for polynomials over number fields [Len83, Lan85]. Using this idea, Rónyai [Rón92] gave a polynomial-time factoring algorithm in the case that F is Galois over \mathbb{Q} .

When F is not Galois over \mathbb{Q} (i.e., $F \neq L$), not every automorphism in $G = \text{Gal}(L/\mathbb{Q})$ induces an automorphism of F (or $\tilde{\mathcal{O}}_F$). One of our key observations is that F may still admit a nontrivial automorphism group, from which we can compute a partial factorization of $f(X)$. Indeed, let H be the subgroup of G fixing F . Then the automorphism group of F is isomorphic to $N_G(H)/H$, where $N_G(H) := \{g \in G : gHg^{-1} = H\}$ is the *normalizer* of H in G . The corresponding fixed subfield $F' = F^{N_G(H)/H}$ is the smallest subfield of F such that F/F' is Galois. See Figure 1 for an illustration.




Figure 1: The tower of fields and Galois groups.

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 $\tilde{\mathcal{O}}_K := \mathcal{O}_K/p\mathcal{O}_K$ instead of $\tilde{\mathcal{O}}_F$, where \mathcal{O}_K denotes the ring of integers of K . For example, suppose G is the full symmetric group $\text{Sym}(S)$ permuting the set S of roots of $\tilde{f}(X)$ in L . As $F = \mathbb{Q}(\alpha)$, we know $H = \text{Gal}(L/F)$ is the stabilizer G_α . Let β be a root of $\tilde{f}(X)$ different from α .

Then the automorphism group of $K = F(\beta) = \mathbb{Q}(\alpha, \beta)$ is the nontrivial group $N_G(G_{\alpha, \beta})/G_{\alpha, \beta} \cong \text{Sym}(2)$, which contains the transposition $(\alpha \beta) \neq e$.

Motivated by the above observation, we design the algorithm so that it computes not only an idempotent decomposition of \bar{O}_F , but also idempotent decompositions of the rings \bar{O}_K simultaneously, where K ranges over a collection \mathcal{F} of subfields of L . Moreover, using ideas in [Rón92, Evd94, IKS09], our algorithm guarantees that these idempotent decompositions satisfy a list of constraints.

Relating idempotent decompositions to \mathcal{P} -schemes. It can be shown that for each $K \in \mathcal{F}$, the idempotent decomposition of \bar{O}_K corresponds to a partition of the coset space $H \backslash G$, where $H = \text{Gal}(L/K)$. These partitions altogether form a \mathcal{P} -scheme, and the constraints satisfied by the idempotent decompositions are captured by various properties of the \mathcal{P} -scheme. This allows us to analyze the algorithm in terms of \mathcal{P} -schemes, which eventually leads to Theorem 1.2.

1.5. Schemes conjectures for permutation groups

The paper [IKS09] proposed a combinatorial conjecture on m -schemes, called the *schemes conjecture*, whose affirmative resolution would imply a deterministic polynomial-time factoring algorithm under GRH. Proving this conjecture appears to be difficult. However, one can observe that an m -scheme corresponds to a \mathcal{P} -scheme in a special setting where the group G is a full symmetric group. This observation suggests that one should first formulate and attack the analogous conjectures for “less complex” groups.

For each family \mathcal{G} of permutation groups, we formulate an analogous conjecture, called the *schemes conjecture for \mathcal{G}* . Its statement depends on a quantity $d(G)$ of a permutation group G . We define $d(G)$ as well as the related quantity $d'(G)$ in Subsection 3.1. The schemes conjecture for \mathcal{G} then simply states:

Schemes conjecture for \mathcal{G} . $d(G)$ is bounded by an absolute constant $c_{\mathcal{G}}$ for $G \in \mathcal{G}$.

The following theorem connects $d(G)$ with deterministic polynomial factoring:

Theorem 1.3. *Under GRH, there exists a deterministic algorithm that, given a polynomial $f(X)$ of degree n satisfying Assumption 1 and a lifted polynomial $\tilde{f}(X) \in \mathbb{Z}[X]$ with the Galois group $G := \text{Gal}(\tilde{f}/\mathbb{Q})$, computes the complete factorization (resp. a proper factorization) of $f(X)$ over \mathbb{F}_p in time polynomial in $n^{d(G)}$ (resp. $n^{d'(G)}$) and the size of the input.*

In particular, the complete factorization of $f(X)$ can be computed in deterministic polynomial time from $\tilde{f}(X)$, provided that the Galois group of $\tilde{f}(X)$ is in \mathcal{G} and the schemes conjecture for \mathcal{G} is true.

Moreover, we show that these conjectures form a hierarchy of relaxations of the original schemes conjecture in [IKS09]. Specifically, for two families $\mathcal{G}, \mathcal{G}'$ of permutation groups such that every $G \in \mathcal{G}$ is (permutation isomorphic to) a subgroup of some $G' \in \mathcal{G}'$, 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 relaxation of the original schemes conjecture. It is thus natural to first tackle the easier conjectures in this hierarchy, and progress in this approach may shed some light on the problem of deterministic polynomial factoring over finite fields.

The connection with bases. The quantity $d(G)$ is related to the notion of *bases* in permutation group theory. A *base* of a permutation group G on a set S is a subset $T \subseteq S$ whose pointwise stabilizer G_T is trivial. Bases have been studied numerously in permutation group theory and also play a vital role in computational group theory [Bab90].

Denote by $b(G)$ the *minimal base size* of G . We show that $d(G)$ is always bounded by $b(G)$ for any nontrivial permutation group G . It follows that if a family \mathcal{G} of permutation groups has bounded minimal base size, then the schemes conjecture for \mathcal{G} is true. Such examples include the family of primitive solvable groups [Ser96], and more generally, the family of primitive permutation groups not involving $\text{Alt}(k)$ for a constant $k \in \mathbb{N}^+$ [GSS98, LS99].

Outline of the paper. Notations and preliminaries are given in Section 2. Section 3 is devoted to the basic theory of \mathcal{P} -schemes: in Subsection 3.1, we define \mathcal{P} -schemes and its various properties. In Subsection 3.2, we review the definition of m -schemes in [IKS09], and discuss its connection with \mathcal{P} -schemes. In Section 4, we describe the generic factoring algorithm, and prove Theorem 1.2 as well as Theorem 1.3. Theorem 1.1 is proved in Section 5. Finally, we discuss schemes conjectures for permutation groups in Section 6.

2. Preliminaries

For $k \in \mathbb{N}^+$, denote by $[k]$ the set $\{1, 2, \dots, k\}$. Write $A - B$ for the set difference $\{x : x \in A \text{ and } x \notin B\}$ of two sets A and B . The cardinality of a set S is denoted by $|S|$. A *partition* of a set S is a set P of nonempty subsets of S satisfying $S = \coprod_{B \in P} B$, where \coprod denotes the disjoint union. Each $B \in P$ is called a *block* of P . Denote by 0_S the coarsest partition of S , and by ∞_S the finest partition of S . For $T \subseteq S$ and a partition P of S , we have the partition $P|_T := \{B \cap T : B \in P\} - \{\emptyset\}$ of T , called the *restriction* of P to T . For a set S and $k \in \mathbb{N}^+$, define $S^{(k)} := \{(x_1, \dots, x_k) \in S^k : x_i \neq x_j \text{ for } i \neq j\}$.

Write e for the identity element of a group. The *normalizer* of a subgroup H in a group G is $N_G(H) := \{g \in G : gHg^{-1} = H\}$. For a 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 $S = [n]$. For $a, b \in S$, denote by $(a \ b)$ the transposition in $\text{Sym}(S)$ swapping a and b . We let $(a \ b) = e$ if $a = b$.

Group actions. Let G be a group and S be a set. A (*left*) *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*. Write ${}^g x$ for $\varphi(g, x)$ when φ is clear from the context. For $T \subseteq S$, write ${}^g T$ for the set $\{{}^g x : x \in T\}$. Write S^G for the subset of elements in S fixed by G . When S is a ring or a field and G respects the operations in S , the subset S^G is also a subring (resp. subfield) of S .

Let S be a G -set. The *orbit* or G -*orbit* of $x \in S$ 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 $T \subseteq S$, the *pointwise stabilizer* of T is $G_T := \{g \in G : {}^g x = x \text{ for all } x \in T\}$. For $T = \{x_1, \dots, x_k\} \subseteq S$, we also write G_{x_1, \dots, x_k} for G_T .

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 $x \in S$. An 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 diagonal action of G on $S^{(k)}$ is transitive.

An action of G on a finite set S gives a group homomorphism $\rho : G \rightarrow \text{Sym}(S)$. The image $\rho(G)$ is called a *permutation group* on S . When ρ is injective and clear from the context, we just say G is a permutation group on S . Denote by $b(G)$ the *minimal base size* of G , i.e., the minimum cardinality of a set $T \subseteq S$ satisfying $G_T = \{e\}$. It is well known that $|G| \leq |S|^{b(G)}$ [DM96].

All permutation groups in this paper are finite, and act on finite sets.

Definition 2.1 (left and inverse right translation). *Let H and K be subgroups of G . We say K acts on $H \setminus G$ by inverse right translation if ${}^g Hh = Hhg^{-1}$ for $Hh \in H \setminus G$ and $g \in K$.*

Now further assume $K \leq N_G(H)$. We say K acts on $H \setminus G$ by left translation if ${}^g Hh = Hgh$ for $Hh \in H \setminus G$ and $g \in K$. It induces a semiregular action of $K/(K \cap H)$ on $H \setminus G$, and we also say $K/(K \cap H)$ acts on $H \setminus G$ by left translation.

Equivalent actions and permutation isomorphic actions. Let G be a group and let S, T be G -sets. We say the action of G on S and that on T are *equivalent* if there exists a bijective map $\lambda : S \rightarrow T$ satisfying $\lambda({}^g x) = {}^g(\lambda(x))$ for $x \in S$ 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

bijjective map $\lambda : S \rightarrow T$ satisfying $\lambda(gx) = \phi^{(g)}(\lambda(x))$ for $x \in S$ and $g \in G$. When the actions are clear, we often simply say G is permutation isomorphic to H .

It is well known that any transitive group action is equivalent to the action on a right coset space by inverse right translation:

Lemma 2.2. *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.*

3. Introducing \mathcal{P} -schemes

In this section, we introduce \mathcal{P} -schemes and discuss their connection with m -schemes [IKS09].

3.1. Basic definitions

Let G be a finite group. We define the following two kinds of maps between right coset spaces:

- (projection) for $H \leq H' \leq G$, 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 $H \leq G$ 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 3.1. *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 $c_{gHg^{-1},g'} \circ c_{H,g} = c_{H,g'g}$.*
- *(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 actions of G on $H \backslash G$, $H' \backslash G$ and $gHg^{-1} \backslash G$ by inverse right translation.*

The proof is straightforward and left to the reader. We also define the notion of *subgroup systems*:

Definition 3.2 (subgroup system). *A set \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 $g \in G$.*

The following kind of subgroup systems play a key role in this paper.

Definition 3.3 (system of stabilizers). *Suppose G is a finite group acting on a finite set S . For $m \in \mathbb{N}$, define the set of pointwise stabilizers*

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

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).

Definition 3.4 (\mathcal{P} -collection). *Let \mathcal{P} be a subgroup system over a finite group G . A \mathcal{P} -collection is a family $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ indexed by \mathcal{P} , where C_H is a partition of $H \backslash G$ for $H \in \mathcal{P}$.³*

³Throughout this paper, the subscript H of C_H always indicates which right coset space it partitions.

We are now ready to define \mathcal{P} -schemes, the central object of this paper.

Definition 3.5 (\mathcal{P} -scheme). A \mathcal{P} -collection $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ is a \mathcal{P} -scheme if it satisfies the following axioms:

- (compatibility) for $H, H' \in \mathcal{P}$ with $H \leq H'$ 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 bijectively to a block of $C_{gHg^{-1}}$.
- (regularity) for $H, H' \in \mathcal{P}$ with $H \leq H'$, 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 set B' .

Example (trivial \mathcal{P} -schemes). For $H \in \mathcal{P}$, let C_H be the coarsest partition $0_{H \backslash G}$ of $H \backslash G$. Then $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ is a \mathcal{P} -scheme. Similarly, \mathcal{C} is a \mathcal{P} -scheme if we choose C_H to be the finest partition $\infty_{H \backslash G}$ of $H \backslash G$ for $H \in \mathcal{P}$.

Example (orbit \mathcal{P} -scheme). Suppose K is a subgroup of G . For $H \in \mathcal{P}$, let C_H be the partition of $H \backslash G$ into K -orbits with respect to the action of K on $H \backslash G$ by inverse right translation, i.e.,

$$C_H = \{\{Hgh^{-1} : h \in K\} : g \in G\}.$$

It can be shown that $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ is a \mathcal{P} -scheme [Guo17, Theorem 2.2], called the *orbit \mathcal{P} -scheme* associated with K . This notion generalizes the notion of *orbit schemes* in [IKS09]. For more discussion of orbit \mathcal{P} -schemes, see [Guo17, Section 2.4].

In a \mathcal{P} -scheme, the partition of $H \backslash G$ determines that of $H' \backslash G$ whenever $H \leq H'$:

Lemma 3.6. Let $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ be a \mathcal{P} -scheme. For $H, H' \in \mathcal{P}$ with $H \leq H'$, we have $C_{H'} = \{\pi_{H, H'}(B) : B \in C_H\}$.

Proof. Consider $B \in C_H$ and $B' \in C_{H'}$ such that $\pi_{H, H'}(B) \cap B' \neq \emptyset$. By compatibility of \mathcal{C} , we have $\pi_{H, H'}(B) \subseteq B'$. Assume to the contrary that $\pi_{H, H'}(B) \neq B'$. Choose $y \in \pi_{H, H'}(B)$ and $y' \in B' - \pi_{H, H'}(B)$. 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 of \mathcal{C} . So $\pi_{H, H'}(B) = B'$. \square

Next, we define the following properties of \mathcal{P} -schemes:

Definition 3.7 (homogeneity and discreteness). A \mathcal{P} -scheme $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ is homogeneous on a subgroup $H \in \mathcal{P}$ if C_H is the coarsest partition $0_{H \backslash G}$, and otherwise inhomogeneous on H . It is discrete on H if C_H is the finest partition $\infty_{H \backslash G}$, and otherwise non-discrete on H .

Definition 3.8 (antisymmetry). A \mathcal{P} -scheme $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ is antisymmetric if for $H \in \mathcal{P}$ and $g \in N_G(H) - H$, the map $c_{H, g} : H \backslash G \rightarrow gHg^{-1} \backslash G = H \backslash G$ sends every block of C_H to a different block of C_H , or equivalently, the semiregular action of $N_G(H)/H$ on $H \backslash G$ by left translation induces a semiregular action of $N_G(H)/H$ on C_H .

Antisymmetry of \mathcal{P} -schemes is crucial in our factoring algorithm. The following lemma states that it implies discreteness on H for all $H \in \mathcal{P}$ if the trivial subgroup is in \mathcal{P} .

Lemma 3.9. Suppose \mathcal{P} is a subgroup system over a finite group G , and $\{e\} \in \mathcal{P}$. Then 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 $\{e\} \setminus G$ by left translation, we have $C_{\{e\}} = \infty_{\{e\} \setminus G}$ by antisymmetry. Now consider an arbitrary subgroup $H \in \mathcal{P}$. By Lemma 3.6, we have $C_H = \{\pi_{\{e\}, H}(B) : B \in \infty_{\{e\} \setminus G}\} = \infty_{H \setminus G}$. So \mathcal{C} is discrete on H . \square

Finally, we introduce *strong antisymmetry* of \mathcal{P} -schemes, strengthening antisymmetry in Definition 3.8.

Definition 3.10. 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

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

the composition $\sigma_k \circ \dots \circ \sigma_1$ is the identity map on $B_0 = B_k$. In other words, \mathcal{C} is strongly antisymmetric if no nontrivial permutation of any block in any partition C_H can 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 $(\pi_{H_i, H_{i-1}}|_{B_i})^{-1}$.

Lemma 3.11. A strongly antisymmetric \mathcal{P} -scheme is antisymmetric.

Proof. 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 $\sigma = c_{H, g}|_B : B \rightarrow B$. It sends $x \in B$ to ${}^g H x$ 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, σ is not the identity map. By definition, \mathcal{C} is not strongly antisymmetric. \square

The functions $d(\cdot)$ and $d'(\cdot)$. We introduce two functions $d(\cdot)$ and $d'(\cdot)$ below. They are directly related to the running time of the algorithm in Theorem 1.3.

Definition 3.12. Let G be a permutation group on a finite set S . For $m \in \mathbb{N}^+$, let \mathcal{P}_m be the system of stabilizers of depth m over G . Define $d(G), d'(G) \in \mathbb{N}^+$ as follows.

- Let $d(G)$ 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 $|S| > 1$, let $d'(G)$ 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 $d'(G) = 1$.

Lemma 3.13. $1 \leq d'(G) \leq d(G) \leq \max\{b(G), 1\} \leq \max\{|S| - 1, 1\}$.

Proof. The last inequality holds since $b(G) \leq b(\text{Sym}(S)) = |S| - 1$. For $m \geq \max\{b(G), 1\}$, the system of stabilizers \mathcal{P}_m contains the trivial subgroup. So all strongly antisymmetric \mathcal{P}_m -schemes are discrete on G_x for all $x \in S$ by Lemma 3.9. Therefore $d(G) \leq \max\{b(G), 1\}$. The other inequalities are trivial. \square

Also note $d(G) = d(H)$ and $d'(G) = d'(H)$ if G and H are permutation isomorphic. This is because a \mathcal{P} -scheme for G can be turned into a \mathcal{P} -scheme for H and vice versa, by using the isomorphism between G and H and the bijection between the sets they act on.

Relation to polynomial factoring. To motivate the definitions, we now briefly explain how \mathcal{P} -schemes are related to factoring. More details are given in Section 4.

As mentioned in the introduction, we factorize $f(X)$ by computing idempotent decompositions of rings $\bar{\mathcal{O}}_K$, where K ranges over a collection of number fields. Each K corresponds to a group $H \in \mathcal{P}$ via the Galois correspondence $K \mapsto H = \text{Gal}(L/K)$. We will show that idempotent decompositions of $\bar{\mathcal{O}}_K$ correspond one-to-one to partitions of $H \setminus G$ (see Lemma 4.17). Thus the idempotent decompositions we

compute are encoded by a \mathcal{P} -collection $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$. In order to successfully factorize $f(X)$, we want the partitions C_H to be as fine as possible.

Our algorithm proceeds by repeatedly testing compatibility, invariance, and regularity, and strong antisymmetry of \mathcal{C} . The partitions C_H are refined whenever any of these tests fails. Therefore, after \mathcal{C} stabilizes, it is guaranteed to be a strongly antisymmetric \mathcal{P} -scheme \mathcal{C} .

Finally, the factorization is complete (resp. proper) iff \mathcal{C} is discrete (resp. inhomogeneous) on H for certain subgroups $H \in \mathcal{P}$. So the algorithm always computes the complete factorization (resp. a proper factorization) if all strongly antisymmetric \mathcal{P} -schemes are discrete (resp. proper) on H , which is guaranteed by the second condition of Theorem 1.2 (resp. Theorem 4.24).

3.2. The connection with m -schemes

We review the definition of m -schemes in [IKS09] and discuss its connection with \mathcal{P} -schemes.

The definition of m -schemes. Let S be a finite set and let $m \in \mathbb{N}^+$. Define an m -collection on S to be a collection of partitions $\Pi = \{P_1, \dots, P_m\}$ where P_k is a partition of $S^{(k)}$ for $k \in [m]$.

For $k \in [m]$, $\text{Sym}(k)$ acts on $S^{(k)}$ 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 $\pi_T^k : S^{(k)} \rightarrow S^{(k-|T|)}$ 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 3.14 (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 bijectively 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 set B' .

Furthermore, we say Π is 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 a different block. And Π is homogeneous if $P_1 = 0_S$.

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

Definition 3.15. *An m -scheme $\Pi = \{P_1, \dots, P_m\}$ on S is discrete if $P_1 = \infty_S$. It is 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 $i \in [m]$, $T \subsetneq [i]$, and $B \in P_i$.*

Example. Let $S = [3]$. Let $P_1 = 0_S = \{\{3\}\}$ and $P_2 = \{\{(1, 2), (2, 3), (3, 1)\}, \{(1, 3), (3, 2), (2, 1)\}\}$. Then $\Pi = \{P_1, P_2\}$ is an antisymmetric homogeneous 2-scheme on S . Let B be the block in P_2 containing $(1, 2)$. Then both π_1^2 and π_2^2 maps $B \in P_2$ bijectively to $\{3\} \in P_1$. Note $\pi_1^2|_B$ sends $(1, 2)$ to 2 and $\pi_2^2|_B$ sends $(1, 2)$ to 1. So $\pi_1^2|_B \circ (\pi_2^2|_B)^{-1}$ sends 1 to 2 and hence nontrivially permutes $\{3\}$. Therefore, Π is not strongly antisymmetric.

Strong antisymmetry vs. matchings. The papers [IKS09, AIKS14] defined *matchings* of m -schemes and studied m -schemes that have no matching. We will not use this notion, but consider *strong antisymmetry* of m -schemes (and \mathcal{P} -schemes) instead. For completeness, we discuss matchings in Appendix A. In particular, Lemma A.2 states that a strongly antisymmetric m -scheme never has a matching.

The difference between strong antisymmetry and (the nonexistence) of matchings is as follows: a matching [IKS09, AIKS14] corresponds to a nontrivial permutation of the form $\pi_{T'}^i|_B \circ (\pi_T^i|_B)^{-1}$, which is a composition of two maps. In our definition of strong antisymmetry, we consider permutations that are compositions of *finitely many* maps of the form $c_g^i|_B$, $\pi_T^i|_B$, or $(\pi_T^i|_B)^{-1}$. So this definition is *closed under composition*. This property is crucially used in our proof of the self-reduction lemma (Lemma 5.5).

From \mathcal{P}_m -schemes to m -schemes. Next we show that a \mathcal{P}_m -scheme induces an m -scheme, where \mathcal{P}_m is the system of stabilizers of depth m .

We first prove this fact for the case that G is an m -transitive permutation group on S . In this case, the group G acts transitively on $S^{(k)}$ for $k \leq m$ via the diagonal action. And for $x \in S^{(k)}$, Lemma 2.2 gives a bijection $\lambda_x : S^{(k)} \rightarrow G_x \backslash G$. The idea is using these bijections to construct an m -scheme from a \mathcal{P}_m -scheme. Formally, the construction is given as follows:

Definition 3.16. *Let G be an m -transitive permutation group on a finite set S , where $1 \leq m \leq |S|$, and let \mathcal{P}_m be the corresponding system of stabilizers of depth m . For a \mathcal{P}_m -scheme $\mathcal{C} = \{C_H : H \in \mathcal{P}_m\}$, define $\Pi(\mathcal{C}) = \{P_1, \dots, P_m\}$ where for each $k \in [m]$, P_k is a partition of $S^{(k)}$ defined as follows: pick $x = (x_1, \dots, x_k) \in S^{(k)}$, and let $P_k = \{\lambda_x^{-1}(B) : B \in C_{G_x}\}$.*

Theorem 3.17. *$\Pi(\mathcal{C})$ is a well-defined m -scheme on S . It is strongly antisymmetric, (resp. homogeneous, discrete) iff \mathcal{C} is strongly antisymmetric, (resp. homogeneous on G_x for $x \in S$, discrete on G_x for $x \in S$).*

The proof of Theorem 3.17 is routine, and we defer it to Appendix C. To extend this result to a general permutation group G , we need the following lemma.

Lemma 3.18. *Let G be a permutation group on a finite set S , and let G' be a subgroup of G . Let \mathcal{P}_m (resp. \mathcal{P}'_m) be the system of stabilizers of depth m over G (resp. G') with respect to its action on S . Suppose there exists a strongly antisymmetric \mathcal{P}'_m -scheme \mathcal{C}' . Then there exists a strongly antisymmetric \mathcal{P}_m -scheme \mathcal{C} . Moreover, for $x \in S$,*

- (1) *if \mathcal{C}' is non-discrete on $G'_{x'}$, then \mathcal{C} is non-discrete on G_x , and*
- (2) *if \mathcal{C}' is homogeneous on $G'_{x'}$, and G' is transitive on S , then \mathcal{C} is homogeneous on G_x .*

In particular, we have $d(G') \leq d(G)$ and $d'(G') \leq d'(G)$.

Lemma 3.18 is proved using a technique called *induction of \mathcal{P} -schemes*. We defer its proof to Appendix B.

Corollary 3.19. *Let G be a permutation group on a finite set S , and let \mathcal{P}_m be the corresponding system of stabilizers of depth m , where $1 \leq m \leq |S|$. Then:*

- (1) *If there exists a strongly antisymmetric \mathcal{P}_m -scheme, non-discrete on some $x \in S$, then there exists a strongly antisymmetric non-discrete m -scheme on S .*
- (2) *If there exists a strongly antisymmetric \mathcal{P}_m -scheme, homogeneous on some $x \in S$, and G is transitive, then there exists a strongly antisymmetric homogeneous m -scheme on S .*

Proof. Note $G \leq \text{Sym}(S)$. By Lemma 3.18, we may assume $G = \text{Sym}(S)$. In particular, G is m -transitive. The claims now follow from Theorem 3.17. \square

We use Corollary 3.19 to translate known results on m -schemes to upper bounds for $d(G)$ and $d'(G)$. They are summarized by the following theorem.

Theorem 3.20. *Let G be a permutation group on a finite set S of cardinality $n \in \mathbb{N}^+$. We have*

- (1) $d(G) = O(\log n)$.
- (2) *Suppose $n > 1$. Then $d'(G) \leq \ell$, where ℓ is the least prime factor of n .*
- (3) *Suppose $n > 2$ is a prime number. Then $d'(G) \leq \ell + 1$, where ℓ is the greatest prime factor of $n - 1$.*

Proof. By [IKS09, Lemma 8], for sufficiently large $m = \Theta(\log n)$, all antisymmetric non-discrete m -schemes on S have a matching. So by Lemma A.2, all strongly antisymmetric m -schemes on S are discrete. By Corollary 3.19, all strongly antisymmetric \mathcal{P}_m -schemes are discrete on G_x for all $x \in S$. So $d(G) \leq m = O(\log n)$ by definition.

For the other two claims, we may assume G is transitive on S , since otherwise we have $d'(G) = 1$ by definition and the claims are trivial.

Suppose $n > 1$ and let ℓ be the least prime factor of n . By [IKS09, Lemma 1], all antisymmetric ℓ -schemes on S are inhomogeneous. By Corollary 3.19, all strongly antisymmetric \mathcal{P}_ℓ -schemes are inhomogeneous on G_x for all $x \in S$. So $d'(G) \leq \ell$ by definition.

Now suppose $n > 2$ is a prime number and let ℓ be the greatest prime factor of $n - 1$. By [IKS09, Main Theorem], all antisymmetric homogeneous $(\ell + 1)$ -schemes on S have a matching. So by Lemma A.2, all strongly antisymmetric $(\ell + 1)$ -schemes on S are inhomogeneous. By Corollary 3.19, all strongly antisymmetric $\mathcal{P}_{\ell+1}$ -schemes are inhomogeneous on G_x for all $x \in S$. So $d'(G) \leq \ell + 1$ by definition. \square

From m -schemes to \mathcal{P}_m -schemes. Conversely, an m -scheme also induces a \mathcal{P}_m -scheme when G is the full symmetric group. We omit this since it is not used in this paper. See [Guo17, Section 2.3] for details.

4. The generic factoring algorithm

In this section, we present the generic factoring algorithm, and prove Theorem 1.2 as well as Theorem 1.3.

4.1. Algebraic preliminaries

We give algebraic preliminaries of the algorithm. Standard references include [AM69, Lan02, Mar77].

Commutative algebra. All rings in this paper are commutative rings with unity. An ideal I of a ring R is *prime* (resp. *maximal*) if I is proper and R/I is an integral domain (resp. a field). Two ideals I, I' of R are *coprime* if $I + I' = R$. In particular, distinct maximal ideals are coprime. For pairwise coprime ideals I_1, \dots, I_k , it holds that $\bigcap_{i=1}^k I_i = \prod_{i=1}^k I_i$ [AM69, Proposition 1.10]. We also have

Lemma 4.1 (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.*

A commutative ring R is *semisimple* if it is isomorphic to a finite product of fields. By Lemma 4.1, this is equivalent to that R has finitely many maximal ideals and their intersection is zero.

An element x of a ring is 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 idempotents. We define an *idempotent decomposition* of a ring R to be a set I of nonzero mutually orthogonal idempotents of R satisfying $\sum_{x \in I} x = 1$. We say such an idempotent decomposition I is *proper* if $|I| > 1$ and *complete* if all idempotents in I are primitive.

Let R be a semisimple ring. By the Chinese remainder theorem, 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 ideals $\mathfrak{m}' \neq \mathfrak{m}$. The map $\mathfrak{m} \mapsto \delta_{\mathfrak{m}}$ is a one-to-one correspondence between the set of the maximal ideals of R and the set of the primitive idempotents of R . Every idempotent of R can be expressed uniquely as a sum of distinct primitive idempotents.

Galois theory. Let K/K_0 be a field extension. Denoted by $\text{Aut}(K/K_0)$ the automorphism group of K over K_0 . The field 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 $\text{Gal}(K/K_0)$. When K is the splitting field of a polynomial $g(X) \in K_0[X]$ over K_0 , we also write $\text{Gal}(g/K_0)$ for $\text{Gal}(K/K_0)$.

Theorem 4.2 (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 Galois. 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 $H \mapsto K^H$.*

Number theory. Let K be a number field. Denote by \mathcal{O}_K the ring of integers of K , which is a free \mathbb{Z} -module of rank $[K : \mathbb{Q}]$. Denote by $\bar{\mathcal{O}}_K$ the quotient ring $\mathcal{O}_K/p\mathcal{O}_K$. An ideal of \mathcal{O}_K is a nonzero prime ideal iff it is a maximal ideal, and hence these two notions are interchangeable. By the theory of Dedekind domains [Mar77], the ideal $p\mathcal{O}_K$ 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 extension of $\mathbb{Z}/p\mathbb{Z} \cong \mathbb{F}_p$ of degree $d_i \in \mathbb{N}^+$, 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 $d_i = 1$ for $i \in [k]$, we say p splits completely in K . By the Chinese remainder theorem, this is equivalent to $\bar{\mathcal{O}}_K \cong \mathbb{F}_p^k$ where $k = [K : \mathbb{Q}]$. If p splits completely in K , then it splits completely in every subfield of K . It is also known that if p splits completely in two subfields $K, K' \subseteq L$, then it splits completely in the composite field KK' [Mar77, Theorem 31]. In particular, if p splits completely in K , then it splits completely in the Galois closure of K/\mathbb{Q} , which is the composite field of the conjugates of K .

The following theorem states that if p splits completely in K , the set of the prime ideals of \mathcal{O}_K lying over p can be identified with a right coset space.

Theorem 4.3. *Let L be a Galois extension of \mathbb{Q} in which p splits completely, 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 $H \leq G$ and $K = L^H$, the map $Hg \mapsto {}^g\mathfrak{Q}_0 \cap \mathcal{O}_K$ is a one-to-one correspondence between $H \backslash G$ and the set of the prime ideals of \mathcal{O}_K lying over p .*

See, e.g., [Mar77, Theorem 33]. By passing to the quotient ring $\bar{\mathcal{O}}_K$, we have

Corollary 4.4. *Let L, G, \mathfrak{Q}_0 be as in Theorem 4.3. For $H \leq G$ and $K = L^H$, the map $Hg \mapsto ({}^g\mathfrak{Q}_0 \cap \mathcal{O}_K)/p\mathcal{O}_K$ is a one-to-one correspondence between $H \backslash G$ and the set of the maximal ideals of $\bar{\mathcal{O}}_K$.*

The following lemma, roughly speaking, is a special case of the classical Kummer-Dedekind theorem [Neu99, Proposition I.8.3], which relates factorization of polynomials to splitting of prime ideals. Its proof can be found in Appendix C.

Lemma 4.5. *Suppose $f(X) \in \mathbb{F}_p[X]$ is a monic polynomial that factorizes into $n \in \mathbb{N}^+$ distinct linear factors over \mathbb{F}_p , and $\tilde{f}(X) \in \mathbb{Z}[X]$ is an irreducible lifted polynomial of $f(X)$. Let $F = \mathbb{Q}[X]/(\tilde{f}(X))$ and $\alpha = X + (\tilde{f}(X)) \in \mathcal{O}_F$. Then the ring homomorphism $\tau : \mathbb{F}_p[X]/(f(X)) \rightarrow \bar{\mathcal{O}}_F$ sending $X + (f(X))$ to $\alpha + p\mathcal{O}_F$ is well defined, and is an isomorphism. In addition, p splits completely in F .*

Corollary 4.6. *Suppose $f(X) \in \mathbb{F}_p[X]$ is a monic polynomial that factorizes into $n \in \mathbb{N}^+$ distinct linear factors over \mathbb{F}_p , and $\tilde{f}(X) \in \mathbb{Z}[X]$ is a (not necessarily irreducible) lifted polynomial of $f(X)$. Let L be the splitting field of $\tilde{f}(X)$ over \mathbb{Q} . Then p splits completely in L .*

Proof. Applying Lemma 4.5 to each irreducible factor of $\tilde{f}(X)$, we see that p splits completely in $\mathbb{Q}(\alpha)$ for each root α of $\tilde{f}(X)$ in L . As L is the composite field of these fields $\mathbb{Q}(\alpha)$, p also splits completely in L . \square

4.2. Algorithmic preliminaries

Next we present algorithmic preliminaries.

Encoding. In the algorithm, every number field K is encoded by a polynomial $g(X) \in \mathbb{Q}[X]$ irreducible over \mathbb{Q} such that $K \cong \mathbb{Q}[X]/(g(X))$. We may also assume K is encoded by its *structure constants*⁴ in the standard basis $\{1 + (g(X)), X + (g(X)), \dots, X^{\deg(g)-1} + (g(X))\}$ of $\mathbb{Q}[X]/(g(X)) \cong K$. These two ways of encoding K can be converted into each other in polynomial time [Len92]. An algebraic number in K is represented by its coefficients in the standard basis. The algorithm also uses R -algebras that are free R -modules of finite rank, where $R = \mathbb{Z}$ or $R = \mathbb{F}_p$. Each of these R -algebras is encoded by its structure constants in a certain R -basis B , and its elements are encoded by their coefficients in the basis B . Finally, an R -linear map $\phi : A \rightarrow A'$ between free R -modules of finite rank (where $R \in \{\mathbb{Q}, \mathbb{Z}, \mathbb{F}_p\}$) is encoded by $\phi(x) \in A'$ for $x \in B$, where B is a given R -basis of A .

Computing $\bar{\mathcal{O}}_K$. Given a number field K , we want to compute the \mathbb{F}_p -algebra $\bar{\mathcal{O}}_K$. It is natural to first compute the ring of integers \mathcal{O}_K and then compute its 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 greatest square factor of a given integer [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 . We encode \mathcal{O}'_K in terms of its \mathbb{Z} -basis in K , which can be efficiently computed:

Theorem 4.7. *There exists a polynomial-time algorithm that given K and p , computes a \mathbb{Z} -basis of a p -maximal order \mathcal{O}'_K in K .*

See, e.g., [Coh93, Chapter 6]. The p -maximality of \mathcal{O}'_K guarantees that the map $\mathcal{O}'_K/p\mathcal{O}'_K \rightarrow \bar{\mathcal{O}}_K$ induced by the inclusion $\mathcal{O}'_K \hookrightarrow \mathcal{O}_K$ is an isomorphism. We use this fact to compute $\bar{\mathcal{O}}_K$ and the quotient map $\mathcal{O}'_K \rightarrow \bar{\mathcal{O}}_K$:

Lemma 4.8. *There exists a polynomial-time algorithm that given p and a p -maximal order \mathcal{O}'_K , computes $\bar{\mathcal{O}}_K$ in terms of its structure constants in some \mathbb{F}_p -basis and the quotient map $\mathcal{O}'_K \rightarrow \bar{\mathcal{O}}_K$.*

Proof. Compute the structure constants $c_{ijk} \in \mathbb{Z}$ of \mathcal{O}'_K in its given \mathbb{Z} -basis $\{b_1, \dots, b_d\} \subseteq K$, where $d = [K : \mathbb{Q}]$. As the map $\mathcal{O}'_K/p\mathcal{O}'_K \rightarrow \bar{\mathcal{O}}_K$ sending $b_i + p\mathcal{O}'_K$ to $b_i + p\mathcal{O}_K$ for $i \in [d]$ is an isomorphism, the set $\{b_1 + p\mathcal{O}_K, \dots, b_d + p\mathcal{O}_K\}$ is an \mathbb{F}_p -basis of $\bar{\mathcal{O}}_K$. And the structure constants of $\bar{\mathcal{O}}_K$ in this basis are simply $c_{ijk} \bmod p$. The quotient map simply sends each basis element $b_i \in \mathcal{O}'_K$ to $b_i + p\mathcal{O}_K$. \square

Computing embeddings of number fields. It is well known that embeddings of number fields can be computed efficiently:

Lemma 4.9. *There exists a polynomial-time algorithm that given number fields K and K' , computes all the embeddings of K in K' .*

This is achieved by reducing to polynomial factoring over K' . See, e.g., [Coh93, Proposition 4.5.3].

Computing ring homomorphisms between quotient rings. An embedding $\phi : K \hookrightarrow K'$ between number fields induces a ring homomorphism $\bar{\mathcal{O}}_K \rightarrow \bar{\mathcal{O}}_{K'}$, which can be efficiently computed:

Lemma 4.10. *There exists a polynomial-time algorithm that takes number fields K, K' , p -maximal orders $\mathcal{O}'_K, \mathcal{O}'_{K'}$, quotient maps $\mathcal{O}'_K \rightarrow \bar{\mathcal{O}}_K, \mathcal{O}'_{K'} \rightarrow \bar{\mathcal{O}}_{K'}$, and an embedding $\phi : K \rightarrow K'$, computes the ring homomorphism $\bar{\phi} : \bar{\mathcal{O}}_K \rightarrow \bar{\mathcal{O}}_{K'}$ induced from ϕ .*

Lemma 4.10 is proved by restricting ϕ to \mathcal{O}'_K , clearing the denominators, and then passing to the quotient rings. We include a detailed proof in Appendix C.

⁴For an R -algebra A that is also a free R -module of finite rank d , the *structure constants* of A in an R -basis $B = \{b_1, \dots, b_d\} \subseteq A$ are the constants $c_{ijk} \in R$ defined by $b_i b_j = \sum_{k=1}^d c_{ijk} b_k$ for $i, j, k \in [d]$.

4.3. The reduction to computing an idempotent decomposition

Now we start describing the algorithm. Fix the following notations: let $f(X) \in \mathbb{F}_p[X]$ be the degree- n input polynomial satisfying Assumption 1, with a lifted polynomial $\tilde{f}(X) \in \mathbb{Z}[X]$. Let L be the splitting field of \tilde{f} over \mathbb{Q} , and let $G = \text{Gal}(L/\mathbb{Q}) = \text{Gal}(\tilde{f}/\mathbb{Q})$. Finally, fix a prime ideal \mathfrak{Q}_0 of \mathcal{O}_L lying over p .

The first ingredient of the algorithm is the following lemma, which reduces the problem of factoring $f(X)$ to that of computing an idempotent decomposition.

Lemma 4.11. *There exists a polynomial-time algorithm that given the following data*

- (1) $f(X)$ and $\tilde{f}(X)$, where $\tilde{f}(X)$ is irreducible over \mathbb{Q}
- (2) F , $\bar{\mathcal{O}}_F$, a p -maximal order $\mathcal{O}'_F \subseteq F$, the quotient map $\mathcal{O}'_F \rightarrow \bar{\mathcal{O}}_F$, and an idempotent decomposition I_F of $\bar{\mathcal{O}}_F$, where $F = \mathbb{Q}[X]/(\tilde{f}(X))$

computes a factorization of $f(X)$ in polynomial time. Moreover, the factorization is complete (resp. proper) iff I_F is complete (resp. proper).

Proof. Let $\alpha = X + (\tilde{f}(X)) \in \mathcal{O}_F$. Consider the ring isomorphism $\tau : \mathbb{F}_p[X]/(f(X)) \rightarrow \bar{\mathcal{O}}_F$ in Lemma 4.5 which sends $X + (f(X))$ to $\alpha + p\mathcal{O}_F$. This map as well as its inverse can be evaluated in polynomial time given the data $\mathcal{O}'_F \subseteq F$ and $\mathcal{O}'_F \rightarrow \bar{\mathcal{O}}_F$. For $\delta \in I_F$, compute $1 - \tau^{-1}(\delta) \in \mathbb{F}_p[X]/(f(X))$, lift it to a nonzero polynomial $h_\delta(X) \in \mathbb{F}_p[X]$ of degree at most n , and compute

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

By the choice of $h_\delta(X)$, we have $h_\delta(X) \equiv 1 - \tau^{-1}(\delta) \pmod{g(X)}$ for any factor $g(X)$ of $f(X)$. So $g_\delta(X)$ is the product of the monic irreducible factors $g(X)$ of $f(X)$ satisfying $\tau^{-1}(\delta) \equiv 1 \pmod{g(X)}$. By the Chinese remainder theorem, $\mathbb{F}_p[X]/(f(X))$ is isomorphic to the product of $\mathbb{F}_p[X]/(g(X))$, where $g(X)$ ranges over the set of monic irreducible factors of $f(X)$. As $\{\tau^{-1}(\delta) : \delta \in I_F\}$ is an idempotent decomposition of $\mathbb{F}_p[X]/(f(X))$, we have $f(X) = \prod_{\delta \in I_F} g_\delta(X)$ and each $g_\delta(X) \neq 1$. This factorization is complete (resp. proper) iff $|I_F| = n$ (resp. $|I_F| > 1$), which holds iff I_F is complete (resp. proper). \square

See Algorithm 1 for the pseudocode of the algorithm in Lemma 4.11.

Algorithm 1 The algorithm in Lemma 4.11

Input: $f(X)$, $\tilde{f}(X)$, $F = \mathbb{Q}[X]/(\tilde{f}(X))$, $\bar{\mathcal{O}}_F$, p -maximal order \mathcal{O}'_F and the quotient map $\mathcal{O}'_F \rightarrow \bar{\mathcal{O}}_F$, idempotent decomposition I_F of $\bar{\mathcal{O}}_F$

Output: factorization of f

- 1: compute the ring isomorphism $\tau : \mathbb{F}_p[X]/(f(X)) \rightarrow \bar{\mathcal{O}}_F$ in Lemma 4.5 that sends $X + (f(X))$ to $\alpha + p\mathcal{O}_F$, where $\alpha = X + (\tilde{f}(X)) \in \mathcal{O}_F$
 - 2: **for** $\delta \in I_F$ **do**
 - 3: compute nonzero $h_\delta(X) \in \mathbb{F}_p[X]$ of degree at most n lifting $1 - \tau^{-1}(\delta)$
 - 4: $g_\delta(X) \leftarrow \gcd(f(X), h_\delta(X))$
 - 5: **return** the factorization $f(X) = \prod_{\delta \in I_F} g_\delta(X)$
-

By Theorem 4.7 and Lemma 4.8, all the required data in Lemma 4.11, except I_F , can be efficiently computed from $\tilde{f}(X)$. Thus Lemma 4.11 reduces the problem of factoring $f(X)$ to that of computing I_F .

4.4. Computing idempotent decompositions corresponding to a \mathcal{P} -scheme

Now we describe the main body of the algorithm, which takes a collection \mathcal{F} of subfields of L , and computes an idempotent decomposition of $\bar{\mathcal{O}}_K$ for each $K \in \mathcal{F}$. We call such a collection \mathcal{F} a *subfield system*, and associate with \mathcal{F} a subgroup system \mathcal{P} over G . Formally:

Definition 4.12. A subfield system \mathcal{F} is collection of number fields that are isomorphic to subfields of L . The subgroup system \mathcal{P} over G associated with \mathcal{F} is defined by $\mathcal{P} = \{H \leq G : L^H \cong K \text{ for some } K \in \mathcal{F}\}$.

As conjugate subfields are isomorphic, the set \mathcal{P} is closed under conjugation and hence is indeed a subgroup system. As usual, the fields in \mathcal{F} are encoded by irreducible polynomials over \mathbb{Q} , and we do not specify how these fields embed in L . Choosing different embeddings does not affect \mathcal{P} since the images of the same field under different embeddings are isomorphic.

Idempotent decompositions vs. partitions of a right coset space. To explain the algorithm, first we need to establish a correspondence between idempotent decompositions and partitions of a right coset space. Fix a subfield $K \subseteq L$ and let $H = \text{Gal}(L/K)$. By Corollary 4.6 and Assumption 1, we know p splits completely in L , and hence in K . As $p\mathcal{O}_L \cap \mathcal{O}_K = p\mathcal{O}_K$, the map $\bar{\mathcal{O}}_K \rightarrow \bar{\mathcal{O}}_L$ induced from the natural inclusion $\mathcal{O}_K \hookrightarrow \mathcal{O}_L$ is injective and identifies $\bar{\mathcal{O}}_K$ with a subring of $\bar{\mathcal{O}}_L$. By Corollary 4.4 (with $H = \{e\}$ and $K = L$), the action of G on \mathcal{O}_L induces a regular action on the set of the maximal ideals of $\bar{\mathcal{O}}_L$.

Recall that we fixed a prime ideal \mathfrak{Q}_0 of \mathcal{O}_L lying over p at the beginning of Subsection 4.3. Define $\bar{\mathfrak{Q}}_0 := \mathfrak{Q}_0/p\mathcal{O}_L$, which is a maximal ideal of $\bar{\mathcal{O}}_L$. Let $\delta_{\bar{\mathfrak{Q}}_0}$ be the unique primitive idempotent of $\bar{\mathcal{O}}_L$ satisfying $\delta_{\bar{\mathfrak{Q}}_0} \equiv 1 \pmod{\bar{\mathfrak{Q}}_0}$ and $\delta_{\bar{\mathfrak{Q}}_0} \equiv 0 \pmod{\mathfrak{m}}$ for all maximal ideals $\mathfrak{m} \neq \bar{\mathfrak{Q}}_0$ of $\bar{\mathcal{O}}_L$. The correspondence is described as follows:

Definition 4.13. 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 $g^{-1}\delta \equiv g'^{-1}\delta \pmod{\bar{\mathfrak{Q}}_0}$ holds for all $\delta \in I$.⁵ Conversely, for a partition P of $H \backslash G$, define $I(P) := \{\delta_B : B \in P\}$ where $\delta_B := \sum_{g \in G: Hg \in B} g \delta_{\bar{\mathfrak{Q}}_0}$.

Lemma 4.14. $I(P)$ in Definition 4.13 is an idempotent decomposition of $\bar{\mathcal{O}}_K$.

The proof of Lemma 4.14 can be found in Appendix C. To prove that Definition 4.13 does give a one-to-one correspondence, we need the following definition.

Definition 4.15 (support). For an idempotent δ of $\bar{\mathcal{O}}_K$, the support of δ is the set $B_\delta := \{Hg \in H \backslash G : g^{-1}\delta \equiv 1 \pmod{\bar{\mathfrak{Q}}_0}\}$. It consists of the cosets $Hg \in H \backslash G$ such that δ is not in ${}^g\bar{\mathfrak{Q}}_0$.

The proof of the following lemma is routine and can be found in Appendix C.

Lemma 4.16. 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 $\delta \mapsto B_\delta$ with the inverse map $B \mapsto \delta_B$.

Now we are ready to establish the correspondence:

Lemma 4.17. The map $I \mapsto P(I)$ is a one-to-one correspondence between the set of idempotent decompositions of $\bar{\mathcal{O}}_K$ and the set of partitions of $H \backslash G$, with the inverse map $P \mapsto I(P)$.

Proof. Note $I(P) = \{\delta_B : B \in P\}$ by definition and $P(I) = \{B_\delta : \delta \in I\}$ by Lemma 4.16. So $I = I(P(I))$ by Lemma 4.16. Also note the map $B \mapsto \delta_B$ is injective, and hence $P \mapsto I(P)$ is also injective. So $P \mapsto I(P)$ is the inverse of $I \mapsto P(I)$. \square

The associated \mathcal{P} -collection. For each $K \in \mathcal{F}$, the algorithm maintains an idempotent decomposition I_K of $\bar{\mathcal{O}}_K$. We associate with these I_K a \mathcal{P} -collection $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$, defined as follows: for $H \in \mathcal{P}$, fix a field in \mathcal{F} isomorphic to L^H , denoted by K_H , and fix an isomorphism $\tau_H : K_H \rightarrow L^H$. The isomorphism τ_H induces a ring isomorphism $\bar{\tau}_H : \bar{\mathcal{O}}_{K_H} \rightarrow \bar{\mathcal{O}}_{L^H}$. Then define $I_H := \bar{\tau}_H(I_{K_H})$ and $C_H := P(I_H)$.

⁵When we write ${}^g\delta$ or $g^{-1}\delta$, the idempotent $\delta \in \bar{\mathcal{O}}_K$ is regarded as an element of $\bar{\mathcal{O}}_L$ via the inclusion $\bar{\mathcal{O}}_K \hookrightarrow \bar{\mathcal{O}}_L$, so that the group action of G makes sense. Also note here $g^{-1}\delta$ depends only on Hg , since $\delta \in \bar{\mathcal{O}}_K$ is fixed by H .

Computing idempotent decompositions. The algorithm runs as follows: for $K \in \mathcal{F}$, compute $\mathcal{O}'_K, \bar{\mathcal{O}}_K$, and $\mathcal{O}'_K \rightarrow \bar{\mathcal{O}}_K$ using Theorem 4.7 and Lemma 4.8, and then initialize I_K to $\{1\}$.

We run several subroutines that test various properties of \mathcal{C} , including compatibility, invariance, regularity, and strong antisymmetry. If any of these properties is not satisfied, the subroutines will update the idempotent decompositions I_K so that at least one $C_H \in \mathcal{C}$ is properly refined. Formally, we have the following lemmas.⁶

Lemma 4.18 (compatibility/invariance test). *There exists a subroutine `CompatibilityInvarianceTest` 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 or remain the same, and at least one partition C_H is properly refined if \mathcal{C} is not compatible or invariant.*

Lemma 4.19 (regularity test). *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 or remain the same, and at least one partition C_H is properly refined if \mathcal{C} is compatible but not regular.*

Lemma 4.20 (strong antisymmetry test). *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 or remain the same, and at least one partition C_H is properly refined if \mathcal{C} is a \mathcal{P} -scheme, but not a strongly antisymmetric \mathcal{P} -scheme.*

The proofs of Lemma 4.18–4.20 are deferred to Subsection C.1 in the appendix. We run these subroutines repeatedly until every I_K stabilizes. As each ring $\bar{\mathcal{O}}_K$ has $[K : \mathbb{Q}]$ primitive idempotents, the number of refinements that occur is bounded by $\sum_{K \in \mathcal{F}} [K : \mathbb{Q}]$, which in turn is bounded by the size of \mathcal{F} . We conclude

Theorem 4.21. *Under GRH, there exists an algorithm that given a subfield system \mathcal{F} , computes for each $K \in \mathcal{F}$ a p -maximal order $\mathcal{O}'_K \subseteq K$, the quotient ring $\bar{\mathcal{O}}_K$, the quotient map $\mathcal{O}'_K \rightarrow \bar{\mathcal{O}}_K$, and an idempotent decomposition I_K of $\bar{\mathcal{O}}_K$, such that the associated \mathcal{P} -collection $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ is a strongly antisymmetric \mathcal{P} -scheme. The algorithm runs in time polynomial in $\log p$ and the size of \mathcal{F} .*

See Algorithm 2 for the pseudocode of the algorithm in Theorem 4.21.

Algorithm 2 The algorithm in Theorem 4.21

Input: subfield system \mathcal{F}

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

- 1: **for** $K \in \mathcal{F}$ **do**
 - 2: compute $\mathcal{O}'_K, \bar{\mathcal{O}}_K$, and $\mathcal{O}'_K \rightarrow \bar{\mathcal{O}}_K$ using Theorem 4.7 and Lemma 4.8
 - 3: $I_K \leftarrow \{1\}$
 - 4: **repeat**
 - 5: run `CompatibilityInvarianceTest` to update $\{I_K\}_{K \in \mathcal{F}}$ ▷ See Lemma 4.18 & Algorithm 3
 - 6: run `RegularityTest` to update $\{I_K\}_{K \in \mathcal{F}}$ ▷ See Lemma 4.19 & Algorithm 4
 - 7: run `StrongAntisymmetryTest` to update $\{I_K\}_{K \in \mathcal{F}}$ ▷ See Lemma 4.20 & Algorithm 5
 - 8: **until** I_K remains the same in the last iteration for all $K \in \mathcal{F}$
 - 9: **return** $\mathcal{O}'_K, \bar{\mathcal{O}}_K, \mathcal{O}'_K \rightarrow \bar{\mathcal{O}}_K, I_K$ for $K \in \mathcal{F}$
-

⁶In the following, we use “the size of \mathcal{F} ” to denote the number of bits used to encode \mathcal{F} in the algorithm, not its cardinality $|\mathcal{F}|$.

4.5. Constructing a collection of number fields

The last ingredient is a subroutine that constructs a subfield system \mathcal{F} such that $\mathbb{Q}[X]/(\tilde{f}_i(X)) \in \mathcal{F}$ for all irreducible factors $\tilde{f}_i(X)$ of $\tilde{f}(X)$ over \mathbb{Q} . As mentioned in the introduction, we can choose and construct \mathcal{F} in various ways, leading to specific algorithms with different running time. Here we give two examples.

In the following, denote by S the set of roots of $\tilde{f}(X)$ in L . The Galois group G acts naturally on S . Suppose $\tilde{f}(X)$ factorizes into monic irreducible polynomials $\tilde{f}_1(X), \dots, \tilde{f}_k(X)$ over \mathbb{Q} .

Example 1: constructing the splitting field. Choose $\mathcal{F} = \{\mathbb{Q}[X]/(\tilde{f}_1(X)), \dots, \mathbb{Q}[X]/(\tilde{f}_k(X)), L\}$. For $i \in [k]$, the field $\mathbb{Q}[X]/(\tilde{f}_i(X))$ is isomorphic to $\mathbb{Q}(\alpha)$, where α is a root of $\tilde{f}_i(X)$ in L . So \mathcal{P} consists of the trivial subgroup and the stabilizers G_α for $\alpha \in S$. We can construct the fields $\mathbb{Q}[X]/(\tilde{f}_i(X))$ in polynomial time by factoring $\tilde{f}(X)$ into its irreducible factors $\tilde{f}_i(X)$ using the LLL algorithm [LLL82]. The splitting field L can also be constructed efficiently:

Lemma 4.22. *There exists an algorithm that given $\tilde{f}(X) \in \mathbb{Q}[X]$, computes its splitting field L over \mathbb{Q} in time polynomial in $[L : \mathbb{Q}]$ and the size of $\tilde{f}(X)$.*

This is done by simply adjoining the roots of $\tilde{f}(X)$ to \mathbb{Q} . See Appendix C for the proof.

Example 2: constructing the fields associated with a system of stabilizers. Alternatively, we may choose \mathcal{F} so that it contains the fields obtained by adjoining at most m roots of $\tilde{f}(X)$ to \mathbb{Q} for some $m \in \mathbb{N}^+$. By Galois theory, the associated subgroup system is $\mathcal{P} = \{G_{\alpha_1, \dots, \alpha_k} : \alpha_1, \dots, \alpha_k \in S, 1 \leq k \leq m\}$, which is precisely \mathcal{P}_m , the system of stabilizers of depth m with respect to the action of G on S .

Such \mathcal{F} can be computed in time polynomial in n^m and the size of $\tilde{f}(X)$:

Lemma 4.23. *There exists an algorithm that given $\tilde{f}(X) \in \mathbb{Q}[X]$ of degree n and $m \in [n]$, computes a subfield system \mathcal{F} such that (1) $\mathbb{Q}[X]/(\tilde{f}_i(X)) \in \mathcal{F}$ for all irreducible factors $\tilde{f}_i(X)$ of $\tilde{f}(X)$ over \mathbb{Q} , and (2) the subgroup system associated with \mathcal{F} is \mathcal{P}_m , the system of stabilizers of depth m over G with respect to the action of G on S . Moreover, the algorithm runs in time polynomial in n^m and the size of $\tilde{f}(X)$.*

The proof of Lemma 4.23 can be found in Appendix C.

4.6. Putting it together

Combining the results in previous subsections, we obtain a proof of Theorem 1.2:

Proof of Theorem 1.2. Factorize $\tilde{f}(X)$ into monic irreducible factors $\tilde{f}_1(X), \dots, \tilde{f}_k(X)$ over \mathbb{Q} using the LLL algorithm [LLL82]. We have $\tilde{f}_i(X) \in \mathbb{Z}[X]$ for $i \in [k]$ by Gauss's lemma [Lan02]. Let $f_i(X) = \tilde{f}_i(X) \bmod p$ for $i \in [k]$. Then $f(X) = \prod_{i=1}^k f_i(X)$. It remains to completely factorize each $f_i(X)$.

Construct \mathcal{F} using the hypothetical algorithm. Then use Theorem 4.21 to obtain idempotent decompositions I_K of $\bar{\mathcal{O}}_K$ for $K \in \mathcal{F}$ that correspond to a strongly antisymmetric \mathcal{P} -scheme $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$.

Fix an irreducible factor $\tilde{f}_i(X)$ of $\tilde{f}(X)$. Let $F = \mathbb{Q}[X]/(\tilde{f}_i(X)) \in \mathcal{F}$. Let $H = G_\alpha$, where α is a root of $\tilde{f}_i(X)$ in L , so that $L^H = \mathbb{Q}(\alpha) \cong F$. By the second condition in Theorem 1.2, we have $C_H = \infty_{H \setminus G}$. By the definition of C_H , there exists $K_H \in \mathcal{F}$ isomorphic to $L^H \cong F$ and an isomorphism $\tau_H : K_H \rightarrow L^H$ such that $C_H = P(\bar{\tau}_H(I_{K_H}))$. As $C_H = \infty_{H \setminus G}$, the idempotent decomposition I_{K_H} is complete. If $K_H = F$, then I_F is complete. In general, we find $K \in \mathcal{F}$ isomorphic to F such that I_K is complete, and compute a ring isomorphism $\bar{\tau}' : \bar{\mathcal{O}}_K \cong \bar{\mathcal{O}}_F$ using Lemma 4.9 and Lemma 4.10. By replacing I_F with $\bar{\tau}'(I_K)$, we may assume I_F is complete.⁷ Finally, apply Lemma 4.11 to $(f_i(X), \tilde{f}_i(X))$ to extract the complete factorization of $f_i(X)$ from I_F . \square

⁷In fact, this step is not necessary since the compatibility/invariance test described in Subsection C.1 already guarantees that $I_F = \bar{\tau}'(I_K)$ is satisfied.

We have a similar theorem on the easier problem of computing a proper factorization of $f(X)$:

Theorem 4.24. *Let \mathcal{I} be a family of instances of the input $(f(X), \tilde{f}(X))$ where $f(X)$ satisfies Assumption 1. Suppose there exists a deterministic algorithm that given an instance $s = (f(X), \tilde{f}(X)) \in \mathcal{I}$ where $\tilde{f}(X)$ is irreducible over \mathbb{Q} , constructs in time $T(s)$ a collection \mathcal{F} of number fields that are isomorphic to subfields of the splitting field L of $\tilde{f}(X)$ over \mathbb{Q} , such that*

- $\mathbb{Q}[X]/(\tilde{f}(X)) \in \mathcal{F}$, and
- all strongly antisymmetric \mathcal{P} -schemes are inhomogeneous on $\text{Gal}(L/\mathbb{Q}(\alpha)) \in \mathcal{P}$ for all roots α of $\tilde{f}(X)$ in L , where \mathcal{P} is the subgroup system over $G = \text{Gal}(L/\mathbb{Q})$ associated with \mathcal{F} .

Then under GRH, there exists a deterministic algorithm that given $s = (f(X), \tilde{f}(X)) \in \mathcal{I}$, outputs a proper factorization of $f(X)$ over \mathbb{F}_p in time polynomial in $T(s)$ and the size of s .

Proof. As in the proof of Theorem 1.2, factorize $\tilde{f}(X)$ into monic irreducible factors $\tilde{f}_1(X), \dots, \tilde{f}_k(X)$ over \mathbb{Q} and obtain a factorization $f(X) = \prod_{i=1}^k f_i(X)$ where $f_i(X) = \tilde{f}_i(X) \bmod p$ for $i \in [k]$. This is already a proper factorization if $\tilde{f}(X)$ is reducible over \mathbb{Q} . So assume $\tilde{f}(X)$ is irreducible. Let $F = \mathbb{Q}[X]/(\tilde{f}(X))$. The rest of the proof is the same as that of Theorem 1.2 except that the second condition in Theorem 4.24 only implies $C_H \neq 0_{H \setminus G}$ (instead of $C_H = \infty_{H \setminus G}$) and I_F is proper. Finally, apply Lemma 4.11 to extract a proper factorization of $f(X)$ from I_F . \square

Theorem 1.3 now follows from Theorem 1.2 and Theorem 4.24:

Proof of Theorem 1.3. Let $m = d(G)$ (resp. $d'(G)$). First assume the value of m is known. Construct \mathcal{F} using Lemma 4.23 such that the associated subgroup system is \mathcal{P}_m . Then apply Theorem 1.2 (resp. Theorem 4.24) to obtain the complete factorization (resp. a proper factorization) of $f(X)$.

In general, m may be unknown. In this case, just try $m = 1, 2, \dots$ until the algorithm above succeeds. \square

A unified framework for deterministic polynomial factoring. Our generic factoring algorithm provides a unified framework for deterministic polynomial factoring over finite fields. To illustrate this point, we derive the main results in [Hua91a, Hua91b, Rón88, Rón92, Evd94, IKS09] from our algorithm.

Theorem 4.25 ([Rón92]). *Under GRH, there exists a deterministic algorithm that, given $f(X) \in \mathbb{F}_p[X]$ satisfying Assumption 1 and a lifted polynomial $\tilde{f}(X) \in \mathbb{Z}[X]$, computes the complete factorization of $f(X)$ over \mathbb{F}_p in time polynomial in $|\text{Gal}(\tilde{f}/\mathbb{Q})|$ and the size of the input.*

Proof. Choose \mathcal{F} as in Subsection 4.5, Example 1, so that $\{e\} \in \mathcal{P}$. The theorem then follows from Theorem 1.2, Lemma 4.22 and Lemma 3.9. \square

This subsumes Huang's polynomial-time factoring algorithm for abelian Galois groups $\text{Gal}(\tilde{f}/\mathbb{Q})$ [Hua91a, Hua91b], since in the abelian case, we have $|\text{Gal}(\tilde{f}/\mathbb{Q})| = \deg(f)$.

Theorem 4.26 ([Evd94, IKS09]). *Under GRH, there exists a deterministic algorithm that, given $f(X) \in \mathbb{F}_p[X]$ of degree n satisfying Assumption 1, computes the complete factorization of $f(X)$ over \mathbb{F}_p in time polynomial in $n^{\log n}$ and $\log p$.*

Proof. Choose a lifted polynomial $\tilde{f}(X)$ of $f(X)$. Then apply Theorem 1.3 and Theorem 3.20 (1). \square

For the easier problem of computing a proper factorization of $f(X)$, we have

Theorem 4.27 ([Rón88, IKS09]). *Under GRH, there exists a deterministic algorithm that, given $f(X) \in \mathbb{F}_p[X]$ of degree $n > 1$ satisfying Assumption 1, computes a proper factorization of $f(X)$ over \mathbb{F}_p in time polynomial in n^ℓ and $\log p$, where ℓ is the least prime factor of n . If $n > 2$ is prime, the same holds if we choose ℓ to be the greatest prime factor of $n - 1$ instead.*

Proof. Choose a lifted polynomial $\tilde{f}(X)$ of $f(X)$. The first claim then follows from Theorem 1.3 and Theorem 3.20 (2). The second one follows from Theorem 1.3 and Theorem 3.20 (3). \square

Finally, Evdokimov [Evd92] showed that $f(X)$ can be completely factorized in polynomial time under GRH given $\tilde{f}(X)$ with a solvable Galois group G . The idea is reducing to case that G is primitive solvable, and then using the fact that the order of a primitive solvable permutation group is polynomial in its degree [Pál82]. We do not prove this result here, but remark that it can also be recovered in our framework (see [Guo17, Section 4.3]).

Evdokimov's proof [Evd92] also highlights the importance of primitive permutation groups, as the general problem can be reduced to this case. One natural question is if it is possible to design polynomial-time algorithms for primitive Galois groups of superpolynomial order. In Section 5, we study the case of almost simple primitive Galois groups and answer this question affirmatively.

5. A factoring algorithm for almost simple primitive Galois groups

We prove Theorem 1.1 in this section.

5.1. Restriction of \mathcal{P} -schemes

We need a construction called *restriction of \mathcal{P} -schemes*. Let \mathcal{P} be a subgroup system over a finite group G . For a subgroup G' of G , define $\mathcal{P}|_{G'} := \{H \in \mathcal{P} : H \leq G'\}$, which is a subgroup system over G' .

Definition 5.1 (restriction). *Suppose $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ is a \mathcal{P} -scheme. For $H \in \mathcal{P}|_{G'}$, regard $H \setminus G'$ as a subset of $H \setminus G$ in the obvious way. Then for $H \in \mathcal{P}|_{G'}$, the partition C_H of $H \setminus G$ restricts to a partition of $H \setminus G'$, which we denote by $C_H|_{G'}$. Define $\mathcal{C}|_{G'} := \{C_H|_{G'} : H \in \mathcal{P}|_{G'}\}$, called the restriction of \mathcal{C} to G' .*

Lemma 5.2. *$\mathcal{C}|_{G'}$ in Definition 5.1 is a $\mathcal{P}|_{G'}$ -scheme. Moreover, if \mathcal{C} is strongly antisymmetric, so is $\mathcal{C}|_{G'}$.*

The proof of Lemma 5.2 is routine, and we defer it to Appendix C.

5.2. Self-reduction of discreteness

In this subsection, we prove the *self-reduction lemma*. First, we need the following technical lemma.

Lemma 5.3. *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 $H_0 \leq H_1 \cap H_2$ and the restrictions $\mathcal{C}|_{H_1}, \mathcal{C}|_{H_2}$ are discrete on H_0 . For $i = 0, 1, 2$, let B_i be the block in C_{H_i} containing $H_i e \in H_i \setminus G$, so that we have maps $\pi_{H_0, H_1}|_{B_0} : B_0 \rightarrow B_1$ and $\pi_{H_0, H_2}|_{B_0} : B_0 \rightarrow B_2$. 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. It suffices to show that $\pi_{H_0, H_1}|_{B_0}$ and $\pi_{H_0, H_2}|_{B_0}$ 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 Definition 5.1. It follows from discreteness of $\mathcal{C}|_{H_1}$ on H_0 that $B_0 \cap (H_0 \setminus H_1) = \{H_0 e\}$. 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 π_{H_0, H_1} . So $B_0 \cap (H_0 \setminus H_1) = \{H_0 e\}$ is the preimage of $H_1 e$ under $\pi_{H_0, H_1}|_{B_0}$. As \mathcal{C} is regular, the map $\pi_{H_0, H_1}|_{B_0}$ is injective. Injectivity of $\pi_{H_0, H_2}|_{B_0}$ is proved in the same way. \square

It is also convenient to introduce the following notations.

Definition 5.4. Let G be a finite group acting on a finite set S . Suppose \mathcal{P} is a subgroup system over G and $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ is a \mathcal{P} -scheme. Define the (symmetric) binary relation $\leftrightarrow_{\mathcal{C},S}$ on S such that $x \leftrightarrow_{\mathcal{C},S} y$ iff (1) $G_x, G_y, G_{x,y} \in \mathcal{P}$ and (2) $\mathcal{C}|_{G_x}$ and $\mathcal{C}|_{G_y}$ are both discrete on $G_{x,y}$. Let $\sim_{\mathcal{C},S}$ be the equivalence relation on S generated by $\leftrightarrow_{\mathcal{C},S}$, i.e., $x \sim_{\mathcal{C},S} y$ iff $x = y$ or there exists a finite sequence $x_0, \dots, x_k \in S$ such that $x_0 = x$, $x_k = y$ and $x_{i-1} \leftrightarrow_{\mathcal{C},S} x_i$ for $i \in [k]$.

Now we state the main result of this subsection.

Lemma 5.5 (self-reduction lemma). Let $G, S, \mathcal{P}, \mathcal{C}$ be as in Definition 5.4, and suppose \mathcal{C} is strongly antisymmetric. Let $x \in S$ such that $G_x \in \mathcal{P}$ and $x \sim_{\mathcal{C},S} y$ for all $y \in Gx$. Then \mathcal{C} is discrete on Gx .

Proof. Consider distinct $G_x g, G_x g' \in Gx \setminus G$. Let B (resp. B') be the block of $\mathcal{C}|_{G_x}$ containing $G_x g$ (resp. $G_x g'$). We want to show $B \neq B'$. Let $y = g^{-1}x$ and $z = g'^{-1}x$. By assumption, there exists a finite sequence $x_0, \dots, x_k \in S$ such that $x_0 = y$, $x_k = z$ and $x_{i-1} \leftrightarrow_{\mathcal{C},S} x_i$ for $i \in [k]$. Let B_i be the block of $\mathcal{C}|_{G_{x_i}}$ containing $G_{x_i} e$ for $0 \leq i \leq k$, and let B'_i be the block of $\mathcal{C}|_{G_{x_{i-1}, x_i}}$ containing $G_{x_{i-1}, x_i} e$ for $i \in [k]$. By Lemma 5.3, for each $i \in [k]$, the map $\pi_{G_{x_{i-1}, x_i}, G_{x_i}}|_{B'_i} \circ (\pi_{G_{x_{i-1}, x_i}, G_{x_{i-1}}} |_{B_{i-1}})^{-1}$ is a bijection from B_{i-1} to B_i sending $G_{x_{i-1}} e$ to $G_{x_i} e$. Let $\tau : B_0 \rightarrow B_k$ be the composition of these maps. Then $c_{G_z, g'}|_{B_k} \circ \tau \circ c_{G_x, g^{-1}}|_B$ is a bijection from B to B' sending $G_x g$ to $G_x g'$. So $B \neq B'$ by strongly antisymmetry of \mathcal{C} . \square

By the self-reduction lemma, to prove \mathcal{C} is discrete on Gx , it suffices to prove for all $y \in Gx - \{x\}$ that there exists a finite sequence $x = x_0, x_1, \dots, x_k = y \in S$ such that for $i \in [k]$, we have (1) $G_{x_{i-1}}, G_{x_i}, G_{x_{i-1}, x_i} \in \mathcal{P}$ and (2) $\mathcal{C}|_{G_{x_{i-1}}}$ and $\mathcal{C}|_{G_{x_i}}$ are both discrete on G_{x_{i-1}, x_i} . This reduces discreteness of \mathcal{C} to that of $\mathcal{C}|_{G_z}$ where G_z ranges over a collection of stabilizers.

5.3. Standard actions of symmetric groups

One important special case in the proof of Theorem 1.1 concerns *standard actions* of symmetric groups. Given a symmetric group $\text{Sym}(T)$, its natural action on T induces an action on the set of k -subsets (i.e. subsets of cardinality k) of T , where $1 \leq k \leq |T|$, and also an action on the set of partitions of T . We say an action of $\text{Sym}(T)$ is *standard* if it is equivalent to the action on the set of k -subsets, or to the action on an orbit of some partition of T .

For standard actions of symmetric groups (or their subgroups), we will prove the following theorem.

Theorem 5.6. Under GRH, there exists a deterministic algorithm that, given $f(X) \in \mathbb{F}_p[X]$ satisfying Assumption 1 and a lifted polynomial $\tilde{f}(X) \in \mathbb{Z}[X]$ such that $\text{Gal}(\tilde{f}/\mathbb{Q})$ acting on the set of roots of $\tilde{f}(X)$ is permutation isomorphic to a subgroup of a symmetric group $\text{Sym}(T)$ with a standard action, computes the complete factorization of $f(X)$ over \mathbb{F}_p in time polynomial in $k^{\log k}$ and the size of the input, where $k = |T|$.

To prove Theorem 5.6, we first establish sufficient conditions on \mathcal{P} that imply discreteness of \mathcal{P} -schemes. Then we construct a subfield system whose associated subgroup system satisfies these conditions.

Action on the set of k -subsets. Let $G \leq \text{Sym}(T)$. Denote by S the set of k -subsets of T , where $1 \leq k \leq |T|$. Let $n = |S|$. Regard G as a permutation group on S where the action is induced from that on T .

The basic idea behind our analysis is as follows: we want to choose \mathcal{P} such that all strongly antisymmetric \mathcal{P} -schemes are discrete on G_x for $x \in S$. To achieve this, we apply Lemma 5.5 to reduce to the subproblems for stabilizers G_x acting on $G_x y$, where $x, y \in S$. The idea is to only consider k -subsets x, y that are ‘‘close’’ to each other, so that the underlying set $G_x y$ is much smaller than S . For example, when $|x \cap y| = k - 1$, for all $g \in G_x$ and $z = {}^g y \in G_x y$, we have $|x \cap z| = |{}^g x \cap {}^g y| = |x \cap y| = k - 1$. Therefore $|G_x y|$ is bounded by $|\{z \in S : |x \cap z| = k - 1\}| = k(|T| - k)$. Applying Lemma 5.5 one more time with a careful analysis will reduce the size of the underlying set to $O(|T|)$.

For $x, y \in S$, write $x \sim y$ if $|x \cap y| \geq k - 1$, i.e., either $x = y$ or there exists a transposition in $\text{Sym}(T)$ that sends x to y and vice versa. For $x, y, z \in S$, write $y \sim_x z$ if $x \sim y$, $x \sim z$ and $y \sim z$.

We also need the following two lemmas.

Lemma 5.7. *For distinct $x, y, z \in S$ such that $x \sim y$ and $z \in G_x y$, there exists a finite sequence $x_0, \dots, x_t \in S - \{x\}$ such that $x_0 = y$, $x_t = z$, and $x_{i-1} \sim_x x_i$ for $i \in [t]$.*

Proof. Note $x \sim z$. So $y = (x - \{a\}) \cup \{b\}$ and $z = (x - \{a'\}) \cup \{b'\}$ for some $a, a' \in x$ and $b, b' \in T - x$. Then the sequence $y, (x - \{a'\}) \cup \{b\}, z$ satisfies the requirement. \square

Lemma 5.8. *For distinct $x, y, z \in S$ such that $y \sim_x z$, it holds that $|G_{x,yz}| \leq |T|$.*

Proof. Let $u = x \cap y$. Then there exist $a \in x$ and $b \in T - x$ such that $x = u \cup \{a\}$ and $y = u \cup \{b\}$, and $G_{x,y}$ fixes u setwisely as well as a and b . If $b \notin z$, we have $z = u \cup \{c\}$ for some $c \in T$ since $y \sim z$. In this case, as $G_{x,y}$ fixes $u \subseteq z$ setwisely, we have $|G_{x,yz}| \leq |T|$, as desired. Now assume $b \in z$. As $x \sim z$, we have $z = (x - \{b'\}) \cup \{b\}$ for some $b' \in x$. As $G_{x,y}$ fixes x setwisely and also fixes b , the elements in $G_{x,yz}$ are of the form $(x - \{b''\}) \cup \{b\}$ where $b'' \in x$. In this case, we have $|G_{x,yz}| \leq |x| = k \leq |T|$. \square

We give a criterion for discreteness of strongly antisymmetric \mathcal{P} -schemes:

Lemma 5.9. *Suppose \mathcal{P} is a subgroup system over G such that*

- (1) $G_{x,y} \in \mathcal{P}$ for $x, y \in S$, and
- (2) $(G_{x,y})_U \in \mathcal{P}$ for $x, y, z \in S$ and $U \subseteq G_{x,yz}$ satisfying $|G_{x,yz}| \leq |T|$ and $|U| \leq d(\text{Sym}(|T|))$ (with respect to the natural action of $\text{Sym}(|T|)$).

Then all strongly antisymmetric \mathcal{P} -schemes are discrete on G_x for all $x \in S$.

Proof. Let $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ be a strongly antisymmetric \mathcal{P} -scheme. By Lemma 5.5, we just need to show $x \sim_{\mathcal{C},S} y$ for $x, y \in S$. As $\text{Sym}(T)$ is transitive on S and generated by transpositions, for $x, y \in S$, there exists a finite sequence $x_0, \dots, x_t \in S$ such that $x_0 = x$, $x_t = y$, and $x_{i-1} \sim x_i$ for $i \in [t]$. Therefore, as $\sim_{\mathcal{C},S}$ is generated by $\leftrightarrow_{\mathcal{C},S}$, it suffices to show $x \leftrightarrow_{\mathcal{C},S} y$ for $x, y \in S$ satisfying $x \sim y$.

By definition, we want to prove $\mathcal{C}|_{G_x}$ is discrete on $(G_x)_y$ for $x, y \in S$ satisfying $x \sim y$. Fix such $x, y \in S$. Assume $x \neq y$ as otherwise the claim is trivial. Again by Lemma 5.5, it suffices to prove $y \sim_{\mathcal{C}|_{G_x},S} z$ for $z \in G_x y$. By Lemma 5.7, it suffices to show $z \leftrightarrow_{\mathcal{C}|_{G_x},S} w$ (and hence $z \sim_{\mathcal{C}|_{G_x},S} w$) for all distinct $z, w \in S - \{x\}$ satisfying $z \sim_x w$. Fix such z, w . We have $|G_{x,zw}|, |G_{x,wz}| \leq |T|$ by Lemma 5.8.

By Lemma 3.18, we have $d(G_{x,z}) \leq d(\text{Sym}(G_{x,wz}))$, where $G_{x,z}$ acts on the orbit $G_{x,z}w$ and $\text{Sym}(G_{x,wz})$ acts naturally on $G_{x,wz}$. As $|G_{x,zw}| \leq |T|$, we have $d(\text{Sym}(G_{x,wz})) \leq d(\text{Sym}(|T|))$ by Lemma B.4 in the appendix. By Condition (2) of Lemma 5.9, $\mathcal{P}|_{G_{x,z}}$ contains the system of stabilizers of depth $d(\text{Sym}(|T|)) \geq d(G_{x,z})$, with respect to the action of $G_{x,z}$ on $G_{x,z}w$. So $\mathcal{C}|_{G_{x,z}}$ is discrete on $(G_{x,z})_w$. Similarly, $\mathcal{C}|_{G_{x,w}}$ is discrete on $(G_{x,w})_z$. By definition, we have $z \leftrightarrow_{\mathcal{C}|_{G_x},S} w$, as desired. \square

Next, we analyze the case that $\text{Sym}(T)$ acts on a subset of partitions of T and prove a similar criterion.

Action on a subset of partitions. Again, let G be a subgroup of $\text{Sym}(T)$. Then G permutes the partitions of T . Let S be a G -orbit of some partition of T , and let $n = |S|$. Regard G as a permutation group on S .

For $x, y \in S$, write $x \sim y$ if there exist $a, b \in T$ such that $y = (a \ b)x$. For $x, y, z \in S$, write $y \sim_x z$ if there exist $a, b, c \in T$ such that $y = (a \ b)x$ and $z = (a \ c)x$.

Similar to the previous case, we have the following two lemmas.

Lemma 5.10. *For distinct $x, y, z \in S$ such that $x \sim y$ and $z \in G_x y$, there exists a finite sequence $x_0, \dots, x_t \in S - \{x\}$ such that $x_0 = y$, $x_t = z$, and $x_{i-1} \sim_x x_i$ for $i \in [t]$.*

Proof. Choose $a, b \in T$ such that $y = {}^{(a\ b)}x$. Choose $g \in G_x$ such that $z = {}^g y$. Let $c = {}^g a$ and $d = {}^g b$. Then $z = {}^{g(a\ b)}x = {}^{g(a\ b)g^{-1}}x = {}^{(c\ d)}x$. By swapping c and d if necessary, we may assume $a \neq d$. Let $w = {}^{(a\ d)}x$. We have $y = {}^{(a\ b)}x, w = {}^{(a\ d)}x$, and $z = {}^{(c\ d)}x$. So the sequence y, w, z satisfies the requirement. \square

Lemma 5.11. *For distinct $x, y, z \in S$ such that $y \sim_x z$, it holds that $|G_{x,yz}| \leq 4|T|$.*

Proof. Fix $a, b, c \in T$ such that $y = {}^{(a\ b)}x$ and $z = {}^{(a\ c)}x$. Consider arbitrary $g \in G_{x,y}$ and let $w = {}^g z \in G_{x,yz}$. We want to bound the number of possible values of w when g ranges over $G_{x,y}$. Let $h = g(a\ c)g^{-1} = ({}^g a\ {}^g c) \in \text{Sym}(T)$. Then $w = {}^{g(a\ c)}x = {}^{g(a\ c)g^{-1}}x = {}^h x$. So w can be recovered from ${}^g a$ and ${}^g c$, and we just need to bound the number of possible values of ${}^g a$ and ${}^g c$.

Choose $B, B', B'' \in x$ such that $a \in B, b \in B'$ and $c \in B''$. As $y = {}^{(a\ b)}x \neq x$, we have $B \neq B'$ and $\{B, B'\} = x - y = {}^g x - {}^g y = \{{}^g B, {}^g B'\}$. Let $\tilde{B} = (B - \{a\}) \cup \{b\}$ and $\tilde{B}' = (B' - \{b\}) \cup \{a\}$. Then we also have $\{\tilde{B}, \tilde{B}'\} = y - x = {}^g y - {}^g x = \{{}^g \tilde{B}, {}^g \tilde{B}'\}$. As $\{a\} = B - \tilde{B}$, we have $\{{}^g a\} = {}^g B - {}^g \tilde{B}$, and hence $\{{}^g a\}$ is among $B - \tilde{B}, B - \tilde{B}', B' - \tilde{B}$, and $B' - \tilde{B}'$. Note that B, B', \tilde{B} , and \tilde{B}' are determined by a, b and do not depend on g . So ${}^g a$ can take at most four values in T when g ranges over $G_{x,y}$. And ${}^g c \in T$ can take at most $|T|$ values. So $|G_{x,yz}| \leq 4|T|$. \square

Lemma 5.12. *Suppose \mathcal{P} is a subgroup system over G such that*

- (1) $G_{x,y} \in \mathcal{P}$ for $x, y \in S$, and
- (2) $(G_{x,y})_U \in \mathcal{P}$ for $x, y, z \in S$ and $U \subseteq G_{x,yz}$ satisfying $|G_{x,yz}| \leq 4|T|$ and $|U| \leq d(\text{Sym}(4|T|))$ (with respect to the natural action of $\text{Sym}(4|T|)$).

Then all strongly antisymmetric \mathcal{P} -schemes are discrete on G_x for all $x \in S$.

Proof. The proof is the same as that of Lemma 5.9 except that $|T|$, Lemma 5.7, and Lemma 5.8 are replaced by $4|T|$, Lemma 5.10 and Lemma 5.11 respectively. \square

Now we are ready to prove Theorem 5.6.

Proof of Theorem 5.6. Let L be the splitting field of $\tilde{f}(X)$ over \mathbb{Q} . Let S be the set of roots of $\tilde{f}(X)$ in L . Let $G = \text{Gal}(\tilde{f}/\mathbb{Q})$, which is a permutation group on S . Let $N := \max\{|T|^{d(\text{Sym}(|T|))}, (4|T|)^{d(\text{Sym}(4|T|))}\} = k^{O(\log k)}$. First assume an integer N' satisfying $N \leq N' \leq 2N$ is known. The algorithm runs as follows: compute the subfield system \mathcal{F}_0 whose members are those obtained by adjoining one or two roots of $\tilde{f}(X)$ to \mathbb{Q} . Let $\mathcal{F} = \mathcal{F}_0$. Next, for each $K \in \mathcal{F}_0$ and each nonlinear irreducible factor $g(X)$ of $\tilde{f}(X)$ over K (computed using known factoring algorithms [Len83, Lan85]), add to \mathcal{F} all the number fields that are obtained by adjoining at most d roots of $g(X)$ to K , where d is the largest integer satisfying $\deg(g)^d \leq N'$. The total degree of these fields constructed from a fixed factor $g(X)$ is bounded by $\binom{\deg(g)}{d} \deg(g)^d$, which is polynomial in $\deg(g)^d \leq N' = k^{O(\log k)}$. The number of $g(X)$ is bounded by $\deg(f)$. It follows that the time needed to compute \mathcal{F} is polynomial in $k^{\log k}$ and the size of the input. We then feed \mathcal{F} to the generic factoring algorithm in Theorem 1.2 to factorize $f(X)$.

It remains to verify that the associated subgroup system \mathcal{P} satisfies the conditions in Lemma 5.9 and Lemma 5.12. By Galois theory, the subgroup system associated with \mathcal{F}_0 is $\mathcal{P}_2 = \{G_{x,y} : x, y \in S\}$. Consider $K \in \mathcal{F}_0$. It is isomorphic to L^H for some $H \in \mathcal{P}_2$. We may assume $K = L^H$ by fixing an isomorphism $K \cong L^H$. Let $g(X)$ be an irreducible factor of $\tilde{f}(X)$ over K , and let $z \in S$ be a root of $g(X)$ in L . By Galois theory, $H = \text{Gal}(L/K)$ acts transitively on the set of roots of $g(X)$ in L . So the set of roots of $g(X)$ in L is precisely $H z$. Therefore the number fields obtained by adjoining at most d roots of $g(X)$ to K are those of the form $L^{H'}$ where $H' = H_U$ and U is a subset of $H z$ of cardinality at most d . We conclude

$$\begin{aligned} \mathcal{P} &= \{H_U : H \in \mathcal{P}_2, z \in S, U \subseteq H z, |H z|^{|U|} \leq N'\} \\ &= \{(G_{x,y})_U : x, y, z \in S, U \subseteq G_{x,yz}, |G_{x,yz}|^{|U|} \leq N'\}. \end{aligned}$$

By the choice of N' , the conditions in Lemma 5.9 and Lemma 5.12 are satisfied.

Now suppose N' is not known. Try $N' = 1, 2, 4, 8, \dots$ until $f(X)$ is completely factorized. This increases the time complexity by at most a factor of $\log N = O(\log^2 k)$. \square

5.4. Almost simple primitive permutation groups

Theorem 1.1 is proved by combining Theorem 5.6 with the following result of Liebeck and Shalev.

Theorem 5.13 ([LS99]). *Let G be an almost simple primitive permutation group. Then one of the following holds:*

- (1) G is permutation isomorphic to a symmetric group or an alternating group with a standard action.
- (2) The socle of G is a classical simple group.
- (3) $b(G) \leq c$, where $c \in \mathbb{N}^+$ is a constant.

Proof of Theorem 1.1. Run the algorithm in Theorem 5.6, as well as the algorithm in Theorem 4.25 to factorize $f(X)$. But halt the second algorithm if its running time exceeds $\max\{k^{C \log k}, s^C\}$, where $s = \Omega(n)$ is the size of the input and $C > 0$ is a sufficiently large constant.⁸ Then the running time is polynomial in $k^{\log k}$ and the size of the input.

To prove the correctness of the algorithm, we consider the three cases of Theorem 5.13 separately. Case (1) is already addressed by Theorem 5.6. For the other two cases, it suffices to show $|G| = (\max\{k^{\log k}, n\})^{O(1)}$, and then the algorithm in Theorem 4.25 factorizes $f(X)$ completely.

In Case (2), let H be the socle of G , which is a classical simple group. So $H \leq G \leq \text{Aut}(H)$. Here $|\text{Aut}(H)|/|H|$ is the order of the *outer automorphism group* of H , which is known to be bounded by $O(\log |H|)$ [CCNP85]. So $|G| = O(|H| \log |H|) = |H|^{O(1)}$. By assumption, we can embed H in $\text{Sym}(k)$. Then $|H| = k^{O(\log k)}$ due to the work on the minimal degree for a permutation representation of a classical simple group [Coo78] (see also [KL90, Table 5.2.A]). So $|G| = k^{O(\log k)}$, as desired.

Finally, in Case (3), we have $b(G) \leq c$ for some constant c , and hence $|G| \leq n^{b(G)} = n^{O(1)}$. \square

6. A hierarchy of schemes conjectures

The paper [IKS09] proposed the *schemes conjecture* on m -schemes. In this section, we discuss the analogous (and formally easier) schemes conjectures for permutation groups, as mentioned in the introduction.

Let \mathcal{G} be a family of permutation groups. We restate the schemes conjecture for \mathcal{G} :

Conjecture 6.1 (schemes conjecture for \mathcal{G}). *$d(G)$ is bounded by an absolute constant $c_{\mathcal{G}}$ for $G \in \mathcal{G}$.*

It implies a deterministic polynomial-time factoring algorithm under GRH for the case that the Galois group $\text{Gal}(\tilde{f}/\mathbb{Q})$ is, up to permutation isomorphism, a member of \mathcal{G} :

Theorem 6.2. *Assume the schemes conjecture for \mathcal{G} is true. Then under GRH, there exists a deterministic polynomial-time algorithm that, given $f(X) \in \mathbb{F}_p[X]$ satisfying Assumption 1 and a lifted polynomial $\tilde{f}(X) \in \mathbb{Z}[X]$ whose Galois group $\text{Gal}(\tilde{f}/\mathbb{Q})$ is permutation isomorphic to some $G \in \mathcal{G}$ (as a permutation group on the set of roots of $\tilde{f}(X)$), computes the complete factorization of $f(X)$ over \mathbb{F}_p .*

Proof. This follows immediately from Theorem 1.3. \square

⁸The value of k is not assumed to be known, but we may try $k = 1, 2, \dots$ until it succeeds.

Remark. For every family \mathcal{G} of permutation groups, the schemes conjecture for \mathcal{G} is a relaxation of the original scheme conjecture in [IKS09], which states that there exists $m \in \mathbb{N}^+$ such that for every finite set S with $|S| > 1$, every antisymmetric homogeneous m -scheme on S has a matching. To see this, suppose the original schemes conjecture is true for some m , but the schemes conjecture for \mathcal{G} is false. Then there exists $G \in \mathcal{G}$ acting on a finite set S such that $d(G) > m$. By Corollary 3.19 (1), there exists a strongly antisymmetric non-discrete m -scheme $\Pi = \{P_1, \dots, P_m\}$ on S . So P_1 has a block B with $|B| > 1$. Restricting Π to the subset $B \subseteq S$ gives a strongly antisymmetric homogeneous m -scheme on B , which has no matching by Lemma A.2. But this contradicts the assumption that the schemes conjecture is true for m .

Reductions between the schemes conjectures. For two families \mathcal{G} and \mathcal{G}' , write $\mathcal{G} \preceq \mathcal{G}'$ if every $G \in \mathcal{G}$ is permutation isomorphic to a subgroup of some $G' \in \mathcal{G}'$. Denote by \mathcal{G}_{Sym} the family of the symmetric groups acting naturally on finite sets. We have

Lemma 6.3. *The schemes conjecture for \mathcal{G} is implied by that for \mathcal{G}' if $\mathcal{G} \preceq \mathcal{G}'$. In particular, for every family \mathcal{G} , the schemes conjecture for \mathcal{G} is implied by that for \mathcal{G}_{Sym} .*

Proof. The first claim follows from Lemma 3.18 (1) and the definition of $d(G)$. The second claim follows from the first one and the fact that every finite permutation group is permutation isomorphic to a subgroup of a symmetric group with the natural action. \square

So the schemes conjectures for various families of permutation groups form a hierarchy, partially ordered by the relation \preceq , and the schemes conjecture for the family of symmetric groups is the most difficult one. In addition, all these conjectures relax the original schemes conjecture in [IKS09].

Proving the schemes conjecture for \mathcal{G}_{Sym} would actually solve the general problem of deterministic polynomial factoring (under GRH):

Theorem 6.4. *Assume the scheme conjecture for \mathcal{G}_{Sym} is true. Then under GRH, there exists a deterministic polynomial-time algorithm that given $f(X) \in \mathbb{F}_q[X]$, computes the complete factorization of $f(X)$ over \mathbb{F}_q .*

Proof. We may assume $f(X)$ satisfies Assumption 1 by the standard reduction [Ber70, Yun76]. Let \mathcal{G} be the family of all finite permutation groups. By Lemma 6.3, the schemes conjecture for \mathcal{G} is true. Choose a lifted polynomial $\tilde{f}(X)$ of $f(X)$ in polynomial time. As $\text{Gal}(\tilde{f}/\mathbb{Q}) \in \mathcal{G}$, the claim follows from Theorem 6.2. \square

Lower bounds for $d(G)$. Our knowledge of lower bounds for $d(G)$ is very limited. It was shown in [Guo17] that $d(G) > 3$ for infinitely many $G \in \mathcal{G}_{\text{Sym}}$ (see [Guo17, Lemma 2.9 and Lemma 2.21]). On the other hand, even determining if there exists a permutation group G with $d(G) > 4$ is an open problem.

Acknowledgments. *The author is grateful to Chris Umans, Nitin Saxena, Manuel Arora, and Anand Kumar Narayanan for helpful discussions.*

References

- [AIKS14] M. Arora, G. Ivanyos, M. Karpinski, and N. Saxena. Deterministic polynomial factoring and association schemes. *LMS Journal of Computation and Mathematics*, 17(01):123–140, 2014.
- [AM69] M. F. Atiyah and I. G. Macdonald. *Introduction to Commutative Algebra*. Addison-Wesley, 1969.
- [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*, pages 175–178, 1977.
- [Aro13] M. Arora. *Extensibility of association schemes and GRH-based deterministic polynomial factoring*. PhD thesis, Universitäts- und Landesbibliothek Bonn, 2013.
- [Asc08] M. Aschbacher. The subgroup structure of finite alternating and symmetric groups. *Lecture Notes for the Summer School on Finite Groups and Related Geometrical Structures*, 2008.
- [Bab90] L. Babai. Computational complexity in finite groups. In *Proceedings of the ICM*, pages 1479–1489, 1990.
- [Ber67] E. R. Berlekamp. Factoring polynomials over finite fields. *Bell System Technical Journal*, 46(8):1853–1859, 1967.

- [Ber70] E. R. Berlekamp. Factoring polynomials over large finite fields. *Mathematics of Computation*, 24(111):713–735, 1970.
- [BI84] E. Bannai and T. Ito. *Algebraic Combinatorics*. Benjamin/Cummings, 1984.
- [BKS15] J. Bourgain, S. Konyagin, and I. Shparlinski. Character sums and deterministic polynomial root finding in finite fields. *Mathematics of Computation*, 84(296):2969–2977, 2015.
- [CCNP85] J. H. Conway, R. T. Curtis, S. P. Norton, and R. A. Parker. *ATLAS of Finite Groups*. Oxford University Press, 1985.
- [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*, pages 233–245, 2000.
- [Coh93] H. Cohen. *A Course in Computational Algebraic Number Theory*. Springer, 1993.
- [Coo78] B. N. Cooperstein. Minimal degree for a permutation representation of a classical group. *Israel Journal of Mathematics*, 30(3):213–235, 1978.
- [CZ81] D. G. Cantor and H. Zassenhaus. A new algorithm for factoring polynomials over finite fields. *Mathematics of Computation*, 36(154):587–592, 1981.
- [DM96] J. D. Dixon and B. Mortimer. *Permutation Groups*, volume 163. Springer, 1996.
- [Evd92] S. A. Evdokimov. Factorization of solvable polynomials over finite fields and the generalized Riemann hypothesis. *Journal of Soviet Mathematics*, 59(3):842–849, 1992.
- [Evd94] S. A. Evdokimov. Factorization of polynomials over finite fields in subexponential time under GRH. In *Proceedings of the 1st Algorithmic Number Theory Symposium*, pages 209–219, 1994.
- [Gao01] S. Gao. On the deterministic complexity of factoring polynomials. *Journal of Symbolic Computation*, 31(1):19–36, 2001.
- [GSS98] D. Gluck, Á. Seress, and A. Shalev. Bases for primitive permutation groups and a conjecture of Babai. *Journal of Algebra*, 199(2):367–378, 1998.
- [Gua09] Y. Guan. *Factoring polynomials and Grobner bases*. PhD thesis, Clemson University, 2009.
- [Guo17] Z. Guo. *\mathcal{P} -schemes and deterministic polynomial factoring over finite fields*. PhD thesis, Caltech, 2017.
- [Hua91a] M. A. Huang. Factorization of polynomials over finite fields and decomposition of primes in algebraic number fields. *Journal of Algorithms*, 12(3):482–489, 1991.
- [Hua91b] M. A. Huang. Generalized Riemann hypothesis and factoring polynomials over finite fields. *Journal of Algorithms*, 12(3):464–481, 1991.
- [IKRS12] G. Ivanyos, M. Karpinski, L. Rónyai, and N. Saxena. Trading GRH for algebra: algorithms for factoring polynomials and related structures. *Mathematics of Computation*, 81(277):493–531, 2012.
- [IKS09] G. Ivanyos, M. Karpinski, and N. Saxena. Schemes for deterministic polynomial factoring. In *Proceedings of the International Symposium on Symbolic and Algebraic Computation*, pages 191–198, 2009.
- [KL90] P. B. Kleidman and M. W. Liebeck. *The Subgroup Structure of the Finite Classical Groups*, volume 129. Cambridge University Press, 1990.
- [KS98] E. Kaltofen and V. Shoup. Subquadratic-time factoring of polynomials over finite fields. *Mathematics of Computation*, 67(223):1179–1197, 1998.
- [KU11] K. S. Kedlaya and C. Umans. Fast polynomial factorization and modular composition. *SIAM Journal on Computing*, 40(6):1767–1802, 2011.
- [Lan85] S. Landau. Factoring polynomials over algebraic number fields. *SIAM Journal on Computing*, 14(1):184–195, 1985.
- [Lan02] S. Lang. *Algebra*. Springer, 2002.
- [Len83] A. K. Lenstra. Factoring polynomials over algebraic number fields. *Computer Algebra*, pages 245–254, 1983.
- [Len92] H. W. Lenstra, Jr. Algorithms in algebraic number theory. *Bulletin of the American Mathematical Society*, 26(2):211–244, 1992.
- [LLL82] A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4):515–534, 1982.
- [LPS88] M. W. Liebeck, C. E. Praeger, and J. Saxl. On the O’Nan-Scott theorem for finite primitive permutation groups. *Journal of the Australian Mathematical Society (Series A)*, 44(03):389–396, 1988.
- [LS99] M. W. Liebeck and A. Shalev. Simple groups, permutation groups, and probability. *Journal of the American Mathematical Society*, 12(2):497–520, 1999.
- [Mar77] D. A. Marcus. *Number Fields*. Springer, 1977.
- [Neu99] J. Neukirch. *Algebraic Number Theory*. Springer-Verlag, 1999.
- [Pál82] P. P. Pálffy. A polynomial bound for the orders of primitive solvable groups. *Journal of Algebra*, 77(1):127–137, 1982.
- [Pil90] J. Pila. Frobenius maps of abelian varieties and finding roots of unity in finite fields. *Mathematics of Computation*, 55(192):745–763, 1990.
- [Rón88] L. Rónyai. Factoring polynomials over finite fields. *Journal of Algorithms*, 9(3):391–400, 1988.
- [Rón89] L. Rónyai. Factoring polynomials modulo special primes. *Combinatorica*, 9(2):199–206, 1989.
- [Rón92] L. Rónyai. Galois groups and factoring polynomials over finite fields. *SIAM Journal on Discrete Mathematics*, 5(3):345–365, 1992.
- [Sch85] R. Schoof. Elliptic curves over finite fields and the computation of square roots mod p . *Mathematics of Computation*, 44(170):483–494, 1985.

- [Ser96] Á. Seress. The minimal base size of primitive solvable permutation groups. *Journal of the London Mathematical Society*, 53(2):243–255, 1996.
- [Sho90] V. Shoup. On the deterministic complexity of factoring polynomials over finite fields. *Information Processing Letters*, 33(5):261–267, 1990.
- [Sho91] V. Shoup. Smoothness and factoring polynomials over finite fields. *Information Processing Letters*, 38(1):39–42, 1991.
- [Uma08] C. Umans. Fast polynomial factorization and modular composition in small characteristic. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing*, pages 481–490, 2008.
- [vzG87] J. von zur Gathen. Factoring polynomials and primitive elements for special primes. *Theoretical Computer Science*, 52(1):77–89, 1987.
- [vzGS92] J. von zur Gathen and V. Shoup. Computing Frobenius maps and factoring polynomials. *Computational Complexity*, 2(3):187–224, 1992.
- [Yun76] D. Y.Y. Yun. On square-free decomposition algorithms. In *Proceedings of the 3rd ACM Symposium on Symbolic and Algebraic Computation*, pages 26–35, 1976.

A. Matchings of m -schemes

We discuss *matchings* of m -schemes, introduced in [IKS09]. We use the more general definition in [AIKS14] (where it is called a *generalized matching*).

Definition A.1 (matching). *Let $\Pi = \{P_1, \dots, P_m\}$ be an m -scheme on a finite set 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, AIKS14] designed algorithms leading to m -schemes that have no matching. This property is subsumed by strong antisymmetry of m -schemes by the following lemma.

Lemma A.2. *A strongly antisymmetric m -scheme has no matching.*

Proof. Suppose $\Pi = \{P_1, \dots, P_m\}$ is an m -scheme on a finite set S with a matching $B \in P_k$ for some $k \in [m]$. Let $T, T' \subseteq [k]$ be as in Definition A.1 and let $k' := k - |T|$. Then $B' := \pi_T^k(B) = \pi_{T'}^k(B)$ is a block of $P_{k'}$. We have two bijective maps $\pi_T^k|_B$ and $\pi_{T'}^k|_B$ from B to B' , where bijectivity follows from the condition $|B| = |\pi_T^k(B)| = |\pi_{T'}^k(B)|$. These two maps are different 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' . Therefore Π is not strongly antisymmetric. \square

B. Induction of \mathcal{P} -schemes

In this section, we introduce a technique called *induction of \mathcal{P} -schemes*. It may be seen as an opposite operation of *restriction of \mathcal{P} -schemes* introduced in Subsection 5.1. As an application, we prove Lemma 3.18 using this technique.

Suppose G is a finite group, G' is a subgroup of G , \mathcal{P} is a subgroup system over G , and

$$\mathcal{P}' = \{G' \cap H : H \in \mathcal{P}\},$$

which is a subgroup system over G' . We will show that a \mathcal{P} -scheme can be constructed from a \mathcal{P}' -scheme.

We need the following lemma.

Lemma B.1. *Let H be a subgroup of G . For $g \in G$, define the map $\phi_{H,g} : (G' \cap gHg^{-1}) \setminus G' \rightarrow H \setminus G$ that sends $(G' \cap gHg^{-1})h$ to $Hg^{-1}h$ for $h \in G'$. These maps are well defined injections. Moreover, given $g_1, \dots, g_k \in G$ such that $\{g_1^{-1}, \dots, g_k^{-1}\}$ is a complete set of representatives of $H \setminus G/G'$, the images of $\phi_{H,g_1}, \dots, \phi_{H,g_k}$ form a partition of $H \setminus G$.*

Proof. Consider the action of G' on $H \backslash G$ by inverse right translation. For $g \in G$, let O_g be the G' -orbit of Hg^{-1} . The stabilizer of Hg^{-1} is $G' \cap gHg^{-1}$. So by Lemma 2.2, we have an equivalence of actions of G'

$$\lambda_{Hg^{-1}} : O_g \rightarrow (G' \cap gHg^{-1}) \backslash G'$$

sending ${}^h(Hg^{-1}) = Hg^{-1}h^{-1}$ to $(G' \cap gHg^{-1})h^{-1}$ for $h \in G'$. Its inverse is exactly the map $\phi_{H,g}$. Finally, the partition in the last claim is simply the partition of $H \backslash G$ into its G' -orbits O_{g_1}, \dots, O_{g_k} . \square

Using Lemma B.1, we can construct a partition of a coset space of G by combining partitions of coset spaces of G' . This gives the following construction.

Definition B.2 (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 B.1. 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 constructed above is indeed a \mathcal{P} -scheme:

Theorem B.3. *The \mathcal{P} -collection \mathcal{C} in Definition B.2 is a well defined \mathcal{P} -scheme, independent of the choices of the elements g_i . Moreover, if \mathcal{C}' is strongly antisymmetric, so is \mathcal{C} .*

Proof. Fix $H \in \mathcal{P}$. By Lemma B.1, C_H is indeed a partition of $H \backslash G$. We need to show that C_H is independent of the choices of g_1, \dots, g_k . Consider $g'_1, \dots, g'_k \in G$ such that $\{g'_1{}^{-1}, \dots, g'_k{}^{-1}\}$ is a complete set of representatives of $H \backslash G / G'$ as well. We want to show

$$C_H = \left\{ \phi_{H,g'_i}(B) : i \in [k], B \in C'_{G' \cap g'_i H g'_i{}^{-1}} \right\}. \quad (\text{B.1})$$

As the right hand side of (B.1) 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 $B \in C'_{G' \cap g'_i H g'_i{}^{-1}}$. Fix i and B . Choose $j \in [k]$ such that $Hg_j^{-1}G' = Hg'_i{}^{-1}G'$. Then choose $g \in G'$ such that $Hg_j^{-1} = Hg'_i{}^{-1}g^{-1}$. We have the conjugation

$$c_{G' \cap g'_i H g'_i{}^{-1}, g} : (G' \cap g'_i H g'_i{}^{-1}) \backslash G' \rightarrow (G' \cap g_j H g_j^{-1}) \backslash G'.$$

By invariance of \mathcal{C}' , the set $c_{G' \cap g'_i H g'_i{}^{-1}, g}(B)$ is a block of $C'_{G' \cap g_j H g_j^{-1}}$. So $\phi_{H,g_j} \circ c_{G' \cap g'_i H g'_i{}^{-1}, g}(B)$ is a block of C_H . On the other hand, one can check directly that $\phi_{H,g_j} \circ c_{G' \cap g'_i H g'_i{}^{-1}, g} = \phi_{H,g'_i}$. So we have $\phi_{H,g'_i}(B) \in C_H$, as desired. Therefore 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 $H, H' \in \mathcal{P}$ with $H \leq H'$. For $g \in G$, the following diagram commutes:

$$\begin{array}{ccc} (G' \cap gHg^{-1}) \backslash G' & \xrightarrow{\pi_{G' \cap gHg^{-1}, G' \cap gH'g^{-1}}} & (G' \cap gH'g^{-1}) \backslash G' \\ \phi_{H,g} \downarrow & & \downarrow \phi_{H',g} \\ H \backslash G & \xrightarrow{\pi_{H,H'}} & H' \backslash G. \end{array}$$

Fix $B \in C_H$. Choose $g \in G$ and $\tilde{B} \in C'_{G' \cap gHg^{-1}}$ such that $B = \phi_{H,g}(\tilde{B})$. Note $\pi_{G' \cap gHg^{-1}, G' \cap gH'g^{-1}}(\tilde{B})$ is contained in a block of $C'_{G' \cap gH'g^{-1}}$ by compatibility of \mathcal{C}' . It follows that $\pi_{H,H'}(B) = \phi_{H',g} \circ \pi_{G' \cap gHg^{-1}, G' \cap gH'g^{-1}}(\tilde{B})$ is contained in a block of $C_{H'}$. So \mathcal{C} is compatible.

To prove 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., $|(\pi_{H, H'}|_B)^{-1}(y)|$ 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.

To prove 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

$$\begin{array}{ccc} (G' \cap gHg^{-1}) \backslash G' & \xrightarrow{\text{id}} & (G' \cap gHg^{-1}) \backslash G' \\ \phi_{H, g} \downarrow & & \downarrow \phi_{H', gh^{-1}} \\ H \backslash G & \xrightarrow{c_{H, h}} & H' \backslash G, \end{array}$$

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 3.10). By the two diagrams above, we can choose $g_i \in G$ and $\tilde{B}_i \in C'_{G' \cap g_i H_i g_i^{-1}}$ 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 $i \in [k]$. Define $\tilde{\tau} := \tilde{\sigma}_k \circ \dots \circ \tilde{\sigma}_1$ 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. So $Hg_0^{-1} = Hg_k^{-1}g^{-1}$ for some $g \in G'$. As $\phi_{H, g_0} \circ c_{G' \cap g_k H g_k^{-1}, g} = \phi_{H, g_k}$, by composing $\tilde{\tau}$ with $c_{G' \cap g_k H g_k^{-1}, g}$, 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 C' is not strongly antisymmetric. \square

As an application, we now prove Lemma 3.18.

Proof of Lemma 3.18. Note $\mathcal{P}'_m = \{G' \cap H : H \in \mathcal{P}_m\}$. Suppose $C' = \{C'_H : H \in \mathcal{P}'_m\}$. Let $\mathcal{C} = \{C_H : H \in \mathcal{P}\}$ be the induction of C' to \mathcal{P}_m . Then \mathcal{C} is a strongly antisymmetric \mathcal{P} -scheme by Theorem B.3.

Fix $x \in S$. Suppose C' is non-discrete on G'_x . Then $C'_{G'_x} \neq \infty_{G'_x \backslash G'}$. By construction, the restriction of the partition C_{G_x} to $G'_x \backslash G' \subseteq G_x \backslash G$ is precisely $C'_{G'_x}$. So $C_{G_x} \neq \infty_{G_x \backslash G}$, i.e., \mathcal{C} is non-discrete on G_x .

Now suppose C' is homogeneous on G'_x and G' is transitive on S . Then $C'_{G'_x} = 0_{G'_x \backslash G'}$. As G' is transitive on S , it is also transitive on $G_x \backslash G$ by inverse right translation (the two actions are equivalent by Lemma 2.2). Then the natural inclusion $G'_x \backslash G' \hookrightarrow G_x \backslash G$ is a bijection. By construction, the partition C_{G_x} is precisely $C'_{G'_x}$ if we identify $G_x \backslash G$ with $G'_x \backslash G'$. So $C_{G_x} = 0_{G_x \backslash G}$, i.e., \mathcal{C} is homogeneous on G_x .

The claims $d(G') \leq d(G)$ and $d'(G') \leq d'(G)$ now follows from the definition of $d(\cdot)$ and $d'(\cdot)$. \square

Lemma 3.18 also implies monotonicity of the function $d(G)$ for symmetric groups, which is used in the proof of Lemma 5.9:

Lemma B.4. *Let $\text{Sym}(n)$ and $\text{Sym}(n')$ act naturally on $[n]$ and $[n']$ respectively, where $n \leq n'$. Then $d(\text{Sym}(n)) \leq d(\text{Sym}(n'))$.*

Proof. Let $G = \text{Sym}(n)$, but regard it as a permutation group on $[n']$ that acts naturally on $[n] \subseteq [n']$ and fixes $[n'] - [n]$ pointwisely. Then $d(G) \leq d(\text{Sym}(n'))$ by Lemma 3.18. So it suffices to show $d(\text{Sym}(n)) \leq d(G)$.

For $m \in \mathbb{N}^+$, let \mathcal{P}_m (resp. \mathcal{P}'_m) be the system of stabilizers of depth m with respect to the action of $\text{Sym}(n)$ on $[n]$ (resp. G on $[n']$). Clearly $\mathcal{P}_m \subseteq \mathcal{P}'_m$. Any group in \mathcal{P}'_m is of the form G_T where $T \subseteq [n']$ and $1 \leq |T| \leq m$. As G fixes $[n'] - [n]$ pointwisely, we have $G_T = \text{Sym}(n)_{T \cap [n]}$. So $G_T \in \mathcal{P}_m$ unless $T \cap [n] = \emptyset$, in which case $G_T = G$. So $\mathcal{P}'_m - \mathcal{P}_m = \{G\}$ if it is nonempty. Then a strongly antisymmetric \mathcal{P}_m -scheme can be (uniquely) extended to a strongly antisymmetric \mathcal{P}'_m -scheme by choosing the partition of the singleton $G \setminus G$ to be the unique one. It follows that $d(\text{Sym}(n)) \leq d(G)$. \square

C. Omitted proofs

Proof of Theorem 3.17. We want to show that $P_k = \{\lambda_x^{-1}(B) : B \in C_{G_x}\}$ does not depend on the choice of $x \in S^{(k)}$. Consider arbitrary $x, x' \in S^{(k)}$. Choose $h \in G$ such that $x' = hx$. Such h exists since G acts transitively on $S^{(k)}$. 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 $\lambda_{x'} = c_{G_x, h} \circ \lambda_x$. As \mathcal{C} is invariant, $c_{G_x, h}^{-1}$ 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 . So $\Pi(\mathcal{C})$ is well defined.

Next we check that $\Pi(\mathcal{C})$ is an m -scheme. Fix $1 < k \leq m$, $i \in [k]$ and $x \in S^{(k)}$. Let $x' = \pi_i^k(x) \in S^{(k-1)}$. Then the following diagram commutes:

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

Also note that λ_x and $\lambda_{x'}$ are bijections, sending blocks to blocks. Compatibility (resp. regularity) of $\Pi(\mathcal{C})$ then follows from compatibility (resp. regularity) of \mathcal{C} .

Now fix $k \in [m]$, $g \in \text{Sym}(k)$ and $x = (x_1, \dots, x_k) \in S^{(k)}$. Let $x' = c_g^k(x) \in S^{(k)}$. Choose $h \in G$ such that $x' = hx$, which exists by m -transitivity of G . Then $G_{x'} = hG_x h^{-1}$. We have the following commutative diagram:

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

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

Using the two diagrams above and the fact $\lambda_{hx} = c_{G_x, h} \circ \lambda_x$ proved at the beginning, we see that a nontrivial permutation of some $B_0 \in P_k$ for some $k \in [m]$ can be obtained by composing maps of the form $c_g^i|_B$, $\pi_T^i|_B$, or $(\pi_T^i|_B)^{-1}$ iff a nontrivial permutation of some $B_0 \in C_{G_x}$ for some $x \in S^{(k)}$ can 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 $(\pi_{H_i, H_{i-1}}|_{B_i})^{-1}$. Also note every $H \in \mathcal{P}_m$ is of the form $H = G_x$ for some $x = (x_1, \dots, x_k) \in S^{(k)}$. So $\Pi(\mathcal{C})$ is strongly antisymmetric iff \mathcal{C} is strongly antisymmetric.

Finally, note $P_1 = \{\lambda_x^{-1}(B) : B \in C_{G_x}\}$ for $x \in S$. So $\Pi(\mathcal{C})$ is homogeneous (resp. discrete) iff \mathcal{C} is homogeneous on G_x (resp. discrete on G_x) for $x \in S$ by definition. \square

Proof of Lemma 4.5. The last claim follows from the first one since $\mathbb{F}_p[X]/(f(X))$ is a semisimple ring with n distinct maximal ideals by the Chinese remainder theorem. For the first claim, note $\tilde{f}(\alpha) = 0$ and $\tilde{f}(X) \bmod p = f(X)$. Then τ is well defined as $X + (f(X))$ is sent to zero. Identify $\mathbb{F}_p[X]/(f(X))$ with $\mathbb{Z}[\alpha]/p\mathbb{Z}[\alpha]$, where $X + (f(X))$ is identified with $\alpha + p\mathbb{Z}[\alpha]$, so that τ becomes the map $\mathbb{Z}[\alpha]/p\mathbb{Z}[\alpha] \rightarrow \bar{\mathcal{O}}_F$ induced from the natural inclusion $\mathbb{Z}[\alpha] \hookrightarrow \mathcal{O}_F$. As \mathcal{O}_F is integral over $\mathbb{Z}[\alpha]$, for every prime ideal \mathfrak{p} of $\mathbb{Z}[\alpha]$, there exists a prime ideal \mathfrak{q} of \mathcal{O}_F such that $\mathfrak{q} \cap \mathbb{Z}[\alpha] = \mathfrak{p}$ by the *lying-over theorem* [AM69, Theorem 5.10]. So for every maximal ideal \mathfrak{p} of $\mathbb{Z}[\alpha]/p\mathbb{Z}[\alpha]$, there exists a maximal ideal \mathfrak{q} of $\bar{\mathcal{O}}_F$ such that $\tau^{-1}(\mathfrak{q}) = \mathfrak{p}$. As 0 is in every maximal ideal of $\bar{\mathcal{O}}_F$, the kernel $\tau^{-1}(0)$ is contained in the intersection of all the maximal ideals of $\mathbb{Z}[\alpha]/p\mathbb{Z}[\alpha]$, which is zero (as $\mathbb{Z}[\alpha]/p\mathbb{Z}[\alpha] \cong \mathbb{F}_p[X]/(f(X))$ is semisimple). So τ is injective. It is 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

Proof of Lemma 4.10. Let $d = [K : \mathbb{Q}]$. Suppose $\{x_1, \dots, x_d\} \subseteq K$ is the given \mathbb{Z} -basis of \mathcal{O}'_K , so that $\{x_1 + p\mathcal{O}_K, \dots, x_d + p\mathcal{O}_K\}$ is an \mathbb{F}_p -basis of $\bar{\mathcal{O}}_K$. For $i \in [d]$, we want to compute $\bar{\phi}(x_i + p\mathcal{O}_K) = \phi(x_i) + p\mathcal{O}_{K'} \in \mathcal{O}_{K'}$. Note $\phi(x_i) \in \phi(\mathcal{O}_K) \subseteq \mathcal{O}_{K'}$, and we can compute it as an element in K' since ϕ is given. So it suffices to prove the following statement: given an element $\alpha \in K'$ that lies in $\mathcal{O}_{K'}$ and the data in the lemma, the residue $\alpha + p\mathcal{O}_{K'} \in \bar{\mathcal{O}}_{K'}$ can be computed in polynomial time.

Let $d' = [K' : \mathbb{Q}]$ and let $B = \{y_1, \dots, y_{d'}\} \subseteq K'$ be the given \mathbb{Z} -basis of $\mathcal{O}'_{K'}$. Then $\{y_1 + p\mathcal{O}_{K'}, \dots, y_{d'} + p\mathcal{O}_{K'}\}$ is an \mathbb{F}_p -basis of $\bar{\mathcal{O}}_{K'}$ since $\mathcal{O}'_{K'}$ is a p -maximal order. As B is also a \mathbb{Q} -basis of K' , we can compute in polynomial time the coefficients $r_i \in \mathbb{Q}$ for $i \in [d']$ such that $\alpha = \sum_{i=1}^{d'} r_i y_i$. Write each r_i in the form a_i/b_i where a_i, b_i are coprime integers and $b_i \neq 0$. Let m be the least common multiple of all the denominators b_i . Then we have $m\alpha = \sum_{i=1}^{d'} m r_i y_i$ with the coefficients $m r_i \in \mathbb{Z}$. So $m\alpha \in \mathcal{O}'_{K'} \subseteq \mathcal{O}_{K'}$. Therefore

$$m\alpha + p\mathcal{O}_{K'} = \sum_{i=1}^{d'} c_i (y_i + p\mathcal{O}_{K'}), \quad \text{where } c_i = m r_i \bmod p \in \mathbb{F}_p.$$

Suppose $m = p^e m'$ where $e \in \mathbb{N}$, $m' \in \mathbb{Z}$ and $p \nmid m'$. Then $p^e | b_j$ and $p^{e+1} \nmid b_j$ hold for some $j \in [d']$. We claim $e = 0$. Assume to the contrary that $e > 0$. We have $p \nmid a_j$ since a_j, b_j are coprime. So $p \nmid m r_j$, i.e., $c_j \neq 0$. Therefore $m\alpha + p\mathcal{O}_{K'} \neq 0$. But as $\alpha + p\mathcal{O}_{K'} \in \bar{\mathcal{O}}_{K'}$, we have $m\alpha + p\mathcal{O}_{K'} \in p^e \bar{\mathcal{O}}_{K'} = 0$, 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 $c'_i = s c_i$ for $i \in [d']$. Then $\alpha + p\mathcal{O}_{K'}$ is represented in the \mathbb{F}_p -basis $\{y_1 + p\mathcal{O}_{K'}, \dots, y_{d'} + p\mathcal{O}_{K'}\}$ by

$$\alpha + p\mathcal{O}_{K'} = \sum_{i=1}^{d'} c'_i (y_i + p\mathcal{O}_{K'}).$$

\square

Proof of Lemma 4.14. Note $\sum_{\delta \in I(P)} \delta = \sum_{B \in P} \delta_B = \sum_{g \in G} {}^g \delta_{\bar{\Omega}_0} = 1$, and $\delta_B \neq 0$ for $B \in P$. So we only need to show $\delta_B \in \bar{\mathcal{O}}_K$ for $B \subseteq H \setminus G$. For $x = Hg \in H \setminus G$, consider the corresponding maximal ideal of $\bar{\mathcal{O}}_K$

$$\mathfrak{P}_x := ({}^g \bar{\Omega}_0 \cap \mathcal{O}_K) / p\mathcal{O}_K.$$

Choose δ to be an idempotent 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 $x \notin B$. Such an element δ exists and is unique by Corollary 4.4 and the Chinese remainder theorem. Then $\delta \equiv \delta_B \equiv 1 \pmod{{}^g \bar{\Omega}_0}$ for $g \in G$ satisfying $Hg \in B$, and $\delta \equiv \delta_B \equiv 0 \pmod{{}^g \bar{\Omega}_0}$ for $g \in G$ satisfying $Hg \notin B$. It follows that $\delta_B = \delta \in \bar{\mathcal{O}}_K$, again by the Chinese remainder theorem. \square

Proof of Lemma 4.16. Consider $B \in P(I)$ and $\delta := \delta_B = \sum_{g \in G: Hg \in B} g \delta_{\bar{\mathcal{Q}}_0}$. We verify that $B_\delta = B$: for $Hh \in H \setminus G$, we have

$$h^{-1} \delta = \sum_{g \in G: Hg \in B} h^{-1} g \delta_{\bar{\mathcal{Q}}_0}.$$

Note that $h^{-1} g \delta_{\bar{\mathcal{Q}}_0} \bmod \bar{\mathcal{Q}}_0$ equals one if $h = g$, and zero otherwise. So $h^{-1} \delta \bmod \bar{\mathcal{Q}}_0$ equals one if $Hh \in B$ and zero otherwise. It follows by definition that $B_\delta = B$.

It remains to show that the map $\delta \mapsto B_\delta$ is injective. Assume to the contrary that $B_\delta = B_{\delta'}$ for distinct $\delta, \delta' \in I$. Choose $Hg \in B_\delta$. We have $g^{-1} \delta \equiv g^{-1} \delta' \equiv 1 \pmod{\bar{\mathcal{Q}}_0}$ by definition. But then $g^{-1}(\delta \delta') \equiv 1 \pmod{\bar{\mathcal{Q}}_0}$, contradicting the fact $\delta \delta' = 0$. \square

Proof of Lemma 4.22. This is done by adjoining roots of $\tilde{f}(X)$ to \mathbb{Q} repeatedly until $\tilde{f}(X)$ factorizes into linear factors. Formally, we maintain a number field K which is initially \mathbb{Q} . Each time we factorize $\tilde{f}(X)$ over K using known factoring algorithms for polynomials over number fields [Len83, Lan85]. Output K if $\tilde{f}(X)$ factorizes into linear polynomials over K . Otherwise, pick a nonlinear irreducible factor $h(X)$, replace K by $K' = K[X]/(h(X))$, and repeat. To encode K' by an irreducible polynomial over \mathbb{Q} , we need to find a primitive element of K' . It was shown in [Rón92] that this can be solved in the desired running time by an effective version of the primitive element theorem. We refer the readers to [Rón92] for details. \square

Proof of Lemma 4.23. For $i = 1, 2, \dots, m$, we inductively compute \mathcal{F}_i such that its associated subgroup system is \mathcal{P}_i . For $i = 1$, just compute $\mathcal{F}_1 = \{\mathbb{Q}[X]/(f_1(X)), \dots, \mathbb{Q}[X]/(f_k(X))\}$, whose associated subgroup system is $\{G_\alpha : \alpha \in S\} = \mathcal{P}_1$.

For $i = 2, \dots, m$, compute $\mathcal{F}_i = \mathcal{F}_{i-1} \cup \mathcal{K}_i$, where \mathcal{K}_i consists of the fields $K[X]/(h(X))$, K ranges over \mathcal{F}_{i-1} , and $h(X)$ ranges over the set of nonlinear irreducible factors of $\tilde{f}(X)$ over K (these factors are computed using known factoring algorithms [Len83, Lan85]). Again, we need to encode each field $K[X]/(h(X))$ by an irreducible polynomial over \mathbb{Q} , which is done by finding a primitive element [Rón92]. By the induction hypothesis, the subgroup system associated with \mathcal{F}_{i-1} is \mathcal{P}_{i-1} . So the subgroup system associated with \mathcal{K}_i is

$$\mathcal{P}'_i := \{H_\alpha : H \in \mathcal{P}_{i-1}, \alpha \in S, H_\alpha \neq H\}$$

and the subgroup system associated with \mathcal{F}_i is $\mathcal{P}_{i-1} \cup \mathcal{P}'_i = \mathcal{P}_i$, as desired. The total degree of the fields in \mathcal{F}_m over \mathbb{Q} is bounded by $\sum_{H \in \mathcal{P}_m} [G : H] \leq \sum_{i=1}^m \binom{n}{i} n(n-1) \cdots (n-i+1)$, which is polynomial in n^m . So the running time is polynomial in n^m and the size of $\tilde{f}(X)$. \square

Proof of Lemma 5.2. We have projections $\pi_{H,H'}$ and conjugations $c_{H,g}$ between coset spaces of the group G , as well as those between coset spaces of G' . We use $\pi'_{H,H'}$ and $c'_{H,g}$ for the latter maps to distinguish them from the former.

For $H \in \mathcal{P}|_{G'}$, partition $H \setminus G$ into the “fibers” of the projection $\pi_{H,G'} : H \setminus G \rightarrow G' \setminus G$:

$$H \setminus G = \coprod_{y \in G' \setminus G} \pi_{H,G'}^{-1}(y).$$

For $y \in G' \setminus G$, we call $\pi_{H,G'}^{-1}(y)$ the y -fiber of $H \setminus G$. Note that $H \setminus G' \subseteq H \setminus G$ is precisely the y -fiber for $y = G'e \in G' \setminus G$. For $x \in H \setminus G$, call $\pi_{H,G'}(x) \in G' \setminus G$ the index of x .

Consider $H, H' \in \mathcal{P}|_{G'}$ and a map $\tau : H \setminus G \rightarrow H' \setminus G$ that is either a projection $\pi_{H,H'}$ with $H \leq H'$, or a conjugation $c_{H,g}$ with $g \in G'$ and $H' = gHg^{-1}$. We claim $\pi_{H,G'} = \pi_{H',G'} \circ \tau$, i.e., the map τ preserves indices. This can be checked directly: if $\tau = \pi_{H,H'}$, we have $\pi_{H',G'} \circ \tau(Hh) = \pi_{H',G'}(H'h) = G'h = \pi_{H,G'}(Hh)$. And if $\tau = c_{H,g}$ with $g \in G'$, we have $\pi_{H',G'} \circ \tau(Hh) = \pi_{H',G'}(H'gh) = G'gh = G'h = \pi_{H,G'}(Hh)$. So the claim holds.

So τ is also fibered over $G' \setminus G$ such that its “ y -fiber” $\tau_y := \tau|_{\pi_{H,G'}^{-1}(y)}$ maps the y -fiber of $H \setminus G$ to the y -fiber of $H' \setminus G$ for $y \in G' \setminus G$. Setting $y = G'e$ yields the map $\tau_y : H \setminus G' \rightarrow H' \setminus G'$ 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, and regularity of $\mathcal{C}|_{G'}$ follows from the corresponding properties of \mathcal{C} .

Assume $\mathcal{C}|_{G'}$ is not strongly antisymmetric. Then there exists a nontrivial permutation τ of a block $B_0 \in C_{H_0}|_{G'}$ for some subgroup $H_0 \in \mathcal{P}|_{G'}$ such that τ is a composition of bijective 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 3.10). 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 $\pi'_{H_{i-1},H_i}|_{B_{i-1}}$ (resp. $(\pi'_{H_i,H_{i-1}}|_{B_i})^{-1}$), the map $\pi_{H_{i-1},H_i}|_{\tilde{B}_{i-1}} : \tilde{B}_{i-1} \rightarrow \tilde{B}_i$ (resp. $(\pi_{H_i,H_{i-1}}|_{\tilde{B}_i})^{-1} : \tilde{B}_{i-1} \rightarrow \tilde{B}_i$) is also a bijection. Then $\tau = \tilde{\tau}|_{B_0}$ holds for the nontrivial permutation $\tilde{\tau} = \tilde{\sigma}_k \cdots \circ \tilde{\sigma}_1$ of the block $\tilde{B}_0 \in C_H$, 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

C.1. Proofs of Lemma 4.18–4.20

In this subsection, we prove Lemma 4.18–4.20 and describe the corresponding subroutines.

Spectrum of a ring. Before presenting the proofs, we describe the intuitions using the *spectrum of a ring* [AM69, Exercise 1.15], which is a fundamental object in commutative algebra and modern algebraic geometry. The discussion here is not necessary for following the formal proofs, which are given in a self-contained way, but may help the reader understand these proofs conceptually.

For every commutative ring A , one can associate a set $\text{Spec}(A)$, called the *spectrum* of A , which is defined to be the set of all prime ideals of A . A ring homomorphism $\phi : A \rightarrow B$ induces a map $\phi^* : \text{Spec}(B) \rightarrow \text{Spec}(A)$ that sends $\mathfrak{p} \in \text{Spec}(B)$ to $\phi^{-1}(\mathfrak{p}) \in \text{Spec}(A)$, and the composition of ring homomorphisms $\phi \circ \psi$ induces the composition of maps $\psi^* \circ \phi^*$ (this means $\text{Spec}(\cdot)$ is a *contravariant functor*).

When we consider only rings A that are finite products of copies of \mathbb{F}_p (like the rings $\bar{\mathcal{O}}_K$), A can be recovered from $\text{Spec}(A)$ as follows: for $a \in A$ and $\mathfrak{m} \in \text{Spec}(A)$, $a \bmod \mathfrak{m} \in A/\mathfrak{m} \cong \mathbb{F}_p$ may be viewed as the value of the “function” a at the “point” \mathfrak{m} . The Chinese remainder theorem (Lemma 4.1) implies that a is determined by its values at all the points $\mathfrak{m} \in \text{Spec}(A)$. So we may identify A with the ring of functions $f : \text{Spec}(A) \rightarrow \mathbb{F}_p$. In this way, we can recover A from $\text{Spec}(A)$. A ring homomorphism $\phi : A \rightarrow B$ can also be recovered from $\phi^* : \text{Spec}(B) \rightarrow \text{Spec}(A)$ via $\phi : f \mapsto f \circ \phi^*$. This means the ring-theoretic language and the set-theoretic language using spectra of rings are equivalent.

Now we apply the theory to our problem: let $K = L^H$ for a subgroup $H \leq G$. Corollary 4.4 states that there is a one-to-one correspondence between $H \setminus G$ and the set of maximal (i.e. prime) ideal of $\bar{\mathcal{O}}_K$. So we may identify $\text{Spec}(\bar{\mathcal{O}}_K)$ with $H \setminus G$. Each subset $B \in C_H$ is then the spectrum of the quotient ring $\bar{\mathcal{O}}_K/(1 - \delta_B)$ (this is an analogue of the ideal-variety correspondence in classical algebraic geometry). Define R_H to be the subring of $\bar{\mathcal{O}}_K$ generated by the idempotents δ_B , $B \in C_H$. The quotient set $(H \setminus G)/\sim_H$ (where \sim_H is the equivalence relation defined by the partition C_H) is then exactly the spectrum of R_H .

Motivated by the equivalence between the ring-theoretic language and the set-theoretic one, one may expect that the properties of \mathcal{P} -schemes can be reformulated and tested in a ring-theoretic way. This is indeed how our subroutines work. For example, compatibility (resp. invariance) of \mathcal{P} -schemes means a map $\tau : H \setminus G \rightarrow H' \setminus G'$ of the form $\pi_{H,H'}$ (resp. $c_{H,g}$) induces a map $(H \setminus G)/\sim_H \rightarrow (H' \setminus G')/\sim_{H'}$ between the quotient sets. This is equivalent to the statement that a certain ring homomorphism $\phi : \bar{\mathcal{O}}_{LH'} \rightarrow \bar{\mathcal{O}}_{LH}$ (for which $\tau = \phi^*$ holds) induces a map $R_{H'} \rightarrow R_H$ between the subrings (i.e., $\phi(R_{H'}) \subseteq R_H$). This is further equivalent to $\phi(\delta_{B'})\delta_B \in \{0, \delta_B\}$ for $B \in C_H$ and $B' \in C_{H'}$, which is precisely what the subroutine in Lemma 4.18 checks. Similarly, to test regularity of a \mathcal{P} -scheme, we check if each map $\bar{\mathcal{O}}_{LH'}/(1 - \delta_{B'}) \rightarrow \bar{\mathcal{O}}_{LH}/(1 - \delta_B)$ makes $\bar{\mathcal{O}}_{LH}/(1 - \delta_B)$ a free module over $\bar{\mathcal{O}}_{LH'}/(1 - \delta_{B'})$. Finally,

to check strong antisymmetry, which states that no nontrivial permutation can be obtained by composing certain maps, we check if a nontrivial ring automorphism can be obtained by composing certain ring homomorphisms.

We remark that the analysis in [IKS09] uses similar ideas. One major difference is that [IKS09] considers certain quotient rings $\mathcal{A}^{(k)}$ of tensor powers $\mathcal{A}^{\otimes k}$, where $\mathcal{A} = \mathbb{F}_p[X]/(f(X))$. In our algorithm, they are replaced by the rings $\bar{\mathcal{O}}_K$ that come from number rings.

C.1.1. Compatibility/invariance test

The subroutine `CompatibilityInvarianceTest` in Lemma 4.18 is given as follows (see Algorithm 3 for the pseudocode): enumerate $K, K' \in \mathcal{F}$, embeddings $\phi : K' \rightarrow K$, $\delta \in I_K$, $\delta' \in I_{K'}$, and check if $\bar{\phi}(\delta')\delta \in \{0, \delta\}$ holds, where $\bar{\phi} : \bar{\mathcal{O}}_{K'} \rightarrow \bar{\mathcal{O}}_K$ is induced from ϕ . If it fails to hold, replace $\delta \in I_K$ by the two nonzero idempotents $\bar{\phi}(\delta')\delta$ and $(1 - \bar{\phi}(\delta'))\delta$, and then return.

Algorithm 3 `CompatibilityInvarianceTest`

Input: idempotent decompositions I_K for $K \in \mathcal{F}$

Output: updated I_K

```

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  $\triangleright \bar{\phi} : \bar{\mathcal{O}}_{K'} \rightarrow \bar{\mathcal{O}}_K$  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

```

We prove that this subroutine satisfies the claim in Lemma 4.18:

Proof of Lemma 4.18. Assume that no proper refinement is made. Recall that we fixed a field $K_H \in \mathcal{F}$ isomorphic to L^H for each $H \in \mathcal{P}$. By identifying L^H with K_H for $H \in \mathcal{P}$, we see that for all $H, H' \in \mathcal{P}$, $\delta \in I_H = I(C_H)$, $\delta' \in I_{H'} = I(C_{H'})$ and embeddings $\phi : L^{H'} \hookrightarrow L^H$, it holds that $\bar{\phi}(\delta')\delta \in \{0, \delta\}$, where $\bar{\phi} : \bar{\mathcal{O}}_{L^{H'}} \rightarrow \bar{\mathcal{O}}_{L^H}$ is induced from ϕ .

Consider H, H' and ϕ be as above and arbitrary $B \in C_H$. Let $\delta = \delta_B$, so that $B_\delta = B$ (see Lemma 4.16). As $\sum_{\delta' \in I_{H'}} \bar{\phi}(\delta') = 1$, we may choose $\delta' \in I_{H'}$ such that $B_{\bar{\phi}(\delta')} \cap B_\delta \neq \emptyset$. As B_δ is defined to be the support of δ (see Definition 4.15), the fact $\bar{\phi}(\delta')\delta \in \{0, \delta\}$ implies $B_{\bar{\phi}(\delta')} \cap B_\delta \in \{\emptyset, B_\delta\}$. But we know $B_{\bar{\phi}(\delta')} \cap B_\delta \neq \emptyset$. So $B = B_\delta \subseteq B_{\bar{\phi}(\delta')}$. Let $B' = B_{\delta'} \in C_{H'}$.

Consider the case that $H \leq H'$ and $\phi : L^{H'} \rightarrow L^H$ is the natural inclusion. Then we have

$$\begin{aligned}
B \subseteq B_{\bar{\phi}(\delta')} &= \{Hg \in H \setminus G : g^{-1} \bar{\phi}(\delta') \equiv 1 \pmod{\bar{\mathfrak{Q}}_0}\} \\
&= \{Hg \in H \setminus G : g^{-1} \delta' \equiv 1 \pmod{\bar{\mathfrak{Q}}_0}\} \\
&= \{Hg \in H \setminus G : H'g \in B_{\delta'}\} \\
&= \pi_{H, H'}^{-1}(B').
\end{aligned} \tag{C.1}$$

So $\pi_{H, H'}(B) \subseteq B'$. Therefore \mathcal{C} is compatible.

Similarly, consider the case that $H' = hHh^{-1}$ for some $h \in G$ and ϕ sends $x \in L^{H'}$ to $h^{-1}x \in L^H$. Then we have

$$\begin{aligned}
B \subseteq B_{\bar{\phi}(\delta')} &= \{Hg \in H \setminus G : g^{-1} \bar{\phi}(\delta') \equiv 1 \pmod{\bar{\mathfrak{Q}}_0}\} \\
&= \{Hg \in H \setminus G : g^{-1} h^{-1} \delta' \equiv 1 \pmod{\bar{\mathfrak{Q}}_0}\} \\
&= \{Hg \in H \setminus G : H'hg \in B_{\delta'}\} \\
&= c_{H, h}^{-1}(B').
\end{aligned} \tag{C.2}$$

So $c_{H,h}(B) \subseteq B'$. Applying the same argument to $c_{H',h^{-1}} = c_{H,h}^{-1}$, we see $c_{H,h}^{-1}(B') \subseteq B$, and hence $c_{H,h}(B) = B'$. Therefore \mathcal{C} is invariant. \square

C.1.2. Regularity test

Now we prove Lemma 4.19. First, we need the following result from [IKS09, IKRS12].

Lemma C.1 ([IKS09, IKRS12]). *There exists a polynomial-time algorithm `FreeModuleTest` that given a semisimple \mathbb{F}_p -algebra A and a finitely generated A -module M , tests whether M is a free A -module. If it is not, the algorithm returns a nonzero zero-divisor of A .*

We also need the following two technical lemmas.

Lemma C.2. *There exists a polynomial-time algorithm `SplitIdempotent` that given a semisimple \mathbb{F}_p -algebra R , a nonzero idempotent δ of R , and an element $a \in R$ such that $a + (1 - \delta) \in R/(1 - \delta)$ is a nonzero zero-divisor, outputs two nonzero idempotents $\delta_1, \delta_2 \in R$ satisfying $\delta_1 + \delta_2 = \delta$.*

Proof. Compute the ideal $(a) \subseteq R$. As R is semisimple, we have $(a) = (\delta')$ for some idempotent δ' of R . Note that δ' is the unique element $b \in (a)$ satisfying $bx = x$ for all $x \in (a)$. So we can compute δ' by solving a system of linear equations over \mathbb{F}_p . Then output $\delta'\delta$ and $(1 - \delta')\delta$.

It remains to prove $\delta'\delta \notin \{0, \delta\}$. Let $I = (1 - \delta) \subseteq R$. As $a + I$ is a nonzero zero divisor of R/I , it generates a nonzero proper ideal J of R/I . As $(a) = (\delta')$, $\delta'\delta + I = \delta' + I$ also generates J . So $\delta'\delta \not\equiv 0 \pmod{I}$ and $\delta'\delta \not\equiv 1 \pmod{I}$. But $1 \equiv \delta \pmod{I}$. So $\delta'\delta \notin \{0, \delta\}$. \square

Lemma C.3. *Suppose $\phi : A' \rightarrow A$ is a ring homomorphism. Let δ, δ' be idempotents of A and A' respectively satisfying $\phi(\delta')\delta = \delta$. Then ϕ induces a ring homomorphism $\phi_{\delta, \delta'} : A'/(1 - \delta') \rightarrow A/(1 - \delta)$ sending $x + (1 - \delta')$ to $\phi(x) + (1 - \delta)$ for $x \in A'$.*

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

The subroutine `RegularityTest` in Lemma 4.19 runs as follows (see Algorithm 4 for the pseudocode): enumerate $K, K' \in \mathcal{F}$, embeddings $\phi : K' \rightarrow K$, $\delta \in I_K$, and $\delta' \in I_{K'}$ satisfying $\bar{\phi}(\delta')\delta = \delta$, where $\bar{\phi} : \bar{\mathcal{O}}_{K'} \rightarrow \bar{\mathcal{O}}_K$ is induced from ϕ . Compute $A = \bar{\mathcal{O}}_{K'}/(1 - \delta')$ and $M = \bar{\mathcal{O}}_K/(1 - \delta)$. Then compute the map $\phi_{\delta, \delta'} : A \rightarrow M$ induced from $\bar{\phi}$ as in Lemma C.3. It makes M an A -algebra and hence an A -module. Run the algorithm in Lemma C.1 to test whether M is a free A -module. If it is not, we obtain a nonzero zero-divisor $\bar{a} \in A$. In this case, we update $I_{K'}$ as follows: lift \bar{a} to $a \in \bar{\mathcal{O}}_{K'}$. Run the algorithm in Lemma C.2 on $\bar{\mathcal{O}}_{K'}$, δ' and a to obtain nonzero idempotents $\delta_1, \delta_2 \in \bar{\mathcal{O}}_{K'}$ satisfying $\delta_1 + \delta_2 = \delta'$. Replace $\delta' \in I_{K'}$ by δ_1 and δ_2 , and then return.

Proof of Lemma 4.19. Assume that no proper refinement is made. By identifying L^H with $K_H \in \mathcal{F}$ for each $H \in \mathcal{P}$, we see that for all $H, H' \in \mathcal{P}$ with $H \leq H'$, $\delta \in I_H = I(C_H)$, $\delta' \in I_{H'} = I(C_{H'})$ satisfying $\bar{\phi}(\delta')\delta = \delta$ where $\bar{\phi} : \bar{\mathcal{O}}_{L^{H'}} \rightarrow \bar{\mathcal{O}}_{L^H}$ is the natural inclusion, the $\bar{\mathcal{O}}_{L^{H'}}/(1 - \delta')$ -module $\bar{\mathcal{O}}_{L^H}/(1 - \delta)$ is a free module.

Consider arbitrary $H, H' \in \mathcal{P}$ with $H \leq H'$ and $B \in C_H$. Then there exists a unique idempotent $\delta \in I_H$ satisfying $B = B_\delta$. Choose $\delta' \in I_{H'}$ such that $\pi_{H, H'}(B_\delta) \cap B_{\delta'} \neq \emptyset$. By compatibility of \mathcal{C} , we have $B_\delta \subseteq \pi_{H, H'}^{-1}(B_{\delta'})$. It suffices to prove that, when $H'g$ ranges over $B_{\delta'}$, $|\pi_{H, H'}^{-1}(H'g) \cap B|$ is a constant.

By Equation (C.1) in the proof of Lemma 4.18, we have $\pi_{H, H'}^{-1}(B_{\delta'}) = B_{\bar{\phi}(\delta')}$. So $B_{\bar{\phi}(\delta')} \cap B_\delta = B_\delta$, or equivalently, $\bar{\phi}(\delta')\delta = \delta$. Let $M = \bar{\mathcal{O}}_{L^H}/(1 - \delta)$ and $A = \bar{\mathcal{O}}_{L^{H'}}/(1 - \delta')$. By assumption, we know M is a free A -module. Denote its rank by k . Then for every maximal ideal \mathfrak{m} of $\bar{\mathcal{O}}_{L^{H'}}$ containing $(1 - \delta')$ (i.e. $\delta' \equiv 1 \pmod{\mathfrak{m}}$) and $\bar{\mathfrak{m}} := \mathfrak{m}/(1 - \delta') \subseteq A$, the ring $M \otimes_A A/\bar{\mathfrak{m}} = \bar{\mathcal{O}}_{L^H}/((1 - \delta) + \bar{\phi}(\mathfrak{m})\bar{\mathcal{O}}_{L^H})$ is a vector space of dimension k over $A/\bar{\mathfrak{m}} \cong \bar{\mathcal{O}}_{L^{H'}}/\mathfrak{m} \cong \mathbb{F}_p$.

Algorithm 4 RegularityTest

Input: idempotent decompositions I_K for $K \in \mathcal{F}$ **Output:** updated I_K

- 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** $\triangleright \bar{\phi} : \bar{\mathcal{O}}_{K'} \rightarrow \bar{\mathcal{O}}_K$ is induced from ϕ
 - 3: compute $A = \bar{\mathcal{O}}_{K'}/(1 - \delta')$, $M = \bar{\mathcal{O}}_K/(1 - \delta)$, and $\phi_{\delta, \delta'} : A \rightarrow M$ induced from $\bar{\phi}$
 - 4: run `FreeModuleTest` in Lemma C.1 to test if M is a free A -module, and obtain a nonzero divisor $a \in A$ if M is not
 - 5: **if** M is not a free A -module **then**
 - 6: run `SplitIdempotent` in Lemma C.2 on $\bar{\mathcal{O}}_{K'}$, δ' and a to obtain $\delta_1, \delta_2 \in \bar{\mathcal{O}}_{K'}$
 - 7: $I_{K'} \leftarrow I_{K'} - \{\delta'\}$
 - 8: $I_{K'} \leftarrow I_{K'} \cup \{\delta_1, \delta_2\}$
 - 9: **return**
-

Consider arbitrary $H'g \in B_{\delta'}$. Let \mathfrak{m} be the maximal ideal $({}^g\mathcal{Q}_0 \cap \mathcal{O}_{LH'})/p\mathcal{O}_{LH'}$ of $\bar{\mathcal{O}}_{LH'}$, and let $R = \bar{\mathcal{O}}_{LH}/((1 - \delta) + \bar{\phi}(\mathfrak{m})\bar{\mathcal{O}}_{LH})$. We have $\delta' \equiv 1 \pmod{\mathfrak{m}}$ by the definition of $B_{\delta'}$. By the previous paragraph, the ring R is a vector space of dimension k over \mathbb{F}_p . As $\bar{\mathcal{O}}_{LH}$ is isomorphic to a finite product of copies of \mathbb{F}_p as a ring, so is its quotient ring R . Thus R has k maximal ideals. Its maximal ideals correspond one-to-one to the maximal ideals of $\bar{\mathcal{O}}_{LH}$ containing both $1 - \delta$ and $\bar{\phi}(\mathfrak{m})$. Consider such a maximal ideal $\mathfrak{m}' = ({}^h\mathcal{Q}_0 \cap \mathcal{O}_{LH})/p\mathcal{O}_{LH}$, which corresponds to some coset $Hh \in H \setminus G$. The condition $1 - \delta \in \mathfrak{m}'$ (i.e., $\delta \equiv 1 \pmod{\mathfrak{m}'}$) is equivalent to $Hh \in B_\delta = B$ by the definition of B_δ . The condition $\bar{\phi}(\mathfrak{m}) \subseteq \mathfrak{m}'$ (i.e., $\bar{\phi}^{-1}(\mathfrak{m}') = \mathfrak{m}$) is equivalent to $\pi_{H, H'}(Hh) = H'g$.⁹ So the number of the maximal ideals of R is

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

which is independent of the choice of $H'g \in B_{\delta'}$. □

C.1.3. Strong antisymmetry test

We need the following result in [Rón92]:

Lemma C.4 ([Rón92]). *Under GRH, there exists a polynomial-time algorithm `FindZeroDivisor` that given a ring A isomorphic to a finite product of copies of \mathbb{F}_p , and a nontrivial ring automorphism σ of A , returns a nonzero zero-divisor of A .*

The subroutine `StrongAntisymmetryTest` in Lemma 4.20 is implemented as follows (see Algorithm 5 for the pseudocode): construct an edge-labeled directed graph (V, E) , where the vertex set is

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

and the edge set E is initially the empty set. For each $(K, \delta) \in V$, compute the ring $\bar{\mathcal{O}}_K/(1 - \delta)$. For each $((K, \delta), (K', \delta')) \in V^2$ and embedding $\phi : K' \hookrightarrow K$ satisfying $\bar{\phi}(\delta')\delta = \delta$, where $\bar{\phi} : \bar{\mathcal{O}}_{K'} \rightarrow \bar{\mathcal{O}}_K$ is induced from ϕ , compute the map $\phi_{\delta, \delta'} : \bar{\mathcal{O}}_{K'}/(1 - \delta') \rightarrow \bar{\mathcal{O}}_K/(1 - \delta)$ induced from $\bar{\phi}$ as in Lemma C.3. If $\phi_{\delta, \delta'}$ is invertible, add to E an edge e from (K', δ') to (K, δ) with label $\phi_{\delta, \delta'}$, and an edge e' from (K, δ) to (K', δ') with label $\phi_{\delta, \delta'}^{-1}$. Then we check if the graph contains a cycle C such that the composition of the labels along C gives a nontrivial ring automorphism σ of $\bar{\mathcal{O}}_K/(1 - \delta)$ for some $(K, \delta) \in V$. This step can be implemented using standard algorithms in graph theory. If such C , σ , and (K, δ) are found, we

⁹This is because $\bar{\phi}^{-1}(\mathfrak{m}') = ({}^h\mathcal{Q}_0 \cap \mathcal{O}_{LH'})/p\mathcal{O}_{LH'}$ and $\mathfrak{m} = ({}^g\mathcal{Q}_0 \cap \mathcal{O}_{LH'})/p\mathcal{O}_{LH'}$ are equal iff $H'h = H'g$ by Corollary 4.4.

update I_K as follows: run the algorithm in Lemma C.4 on $\bar{\mathcal{O}}_K/(1 - \delta)$ and σ to find a nonzero zero-divisor $\bar{a} \in \bar{\mathcal{O}}_K/(1 - \delta)$, and lift it to $a \in \bar{\mathcal{O}}_K$. Then we run the algorithm in Lemma C.2 on $\bar{\mathcal{O}}_K$, δ and a to obtain nonzero idempotents $\delta_1, \delta_2 \in \bar{\mathcal{O}}_K$ satisfying $\delta_1 + \delta_2 = \delta$. Finally, replace $\delta \in I_K$ by δ_1 and δ_2 , and then return.

Algorithm 5 StrongAntisymmetryTest

Input: idempotent decompositions I_K for $K \in \mathcal{F}$

Output: updated I_K

- 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 $\bar{\mathcal{O}}_K/(1 - \delta)$
 - 4: **for** $((K, \delta), (K', \delta')) \in V^2$ and $\phi : K' \hookrightarrow K$ satisfying $\bar{\phi}(\delta')\delta = \delta$ **do**
 - 5: compute $\phi_{\delta, \delta'} : \bar{\mathcal{O}}_{K'}/(1 - \delta') \rightarrow \bar{\mathcal{O}}_K/(1 - \delta)$ induced from $\bar{\phi}$
 - 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: check if G contains a cycle C such that the composition of the labels along C gives a nontrivial ring automorphism σ of $\bar{\mathcal{O}}_K/(1 - \delta)$ for some $(K, \delta) \in V$
 - 9: **if** C , σ , and (K, δ) are found at Line 8 **then**
 - 10: run FindZeroDivisor in Lemma C.4 on $\bar{\mathcal{O}}_K/(1 - \delta)$ and σ to obtain $\bar{a} \in \bar{\mathcal{O}}_K/(1 - \delta)$
 - 11: lift \bar{a} to $a \in \bar{\mathcal{O}}_K$
 - 12: run SplitIdempotent in Lemma C.2 on $\bar{\mathcal{O}}_K$, δ and a to obtain $\delta_1, \delta_2 \in \bar{\mathcal{O}}_K$
 - 13: $I_K \leftarrow I_K - \{\delta\}$
 - 14: $I_K \leftarrow I_K \cup \{\delta_1, \delta_2\}$
 - 15: **return**
-

For $H \in \mathcal{P}$ and $Hg \in H \setminus G$, denote by \mathfrak{m}_{Hg} the maximal ideal $({}^g\Omega_0 \cap \mathcal{O}_{LH})/p\mathcal{O}_{LH}$ of $\bar{\mathcal{O}}_K$ corresponding to the coset Hg . To prove Lemma 4.20, we need the following technical lemma.

Lemma C.5. *Suppose $H, H' \in \mathcal{P}$, $B \in C_H$, $B' \in C_{H'}$, $\sigma : B \rightarrow B'$, and $\phi : L^{H'} \hookrightarrow L^H$ are in one of the following two cases:*

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

Let $\delta = \delta_B \in I_H$ and $\delta' = \delta_{B'} \in I_{H'}$ (see Definition 4.13). 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, and the map $\phi_{\delta, \delta'} : \bar{\mathcal{O}}_{K'}/(1 - \delta') \rightarrow \bar{\mathcal{O}}_K/(1 - \delta)$ satisfies

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

for $Hg \in B$.¹⁰ Moreover, $\phi_{\delta, \delta'}$ is a ring isomorphism if σ is a bijection.

Proof. As $B = B_\delta$, the statement $\bar{\phi}(\delta')\delta = \delta$ is equivalent to $B \subseteq B_{\bar{\phi}(\delta')}$. By Equation (C.1) (resp. Equation (C.2)) in the proof of Lemma 4.18, we have $B_{\bar{\phi}(\delta')} = \pi_{H, H'}^{-1}(B')$ (resp. $B_{\bar{\phi}(\delta')} = c_{H, h}^{-1}(B')$) in Case (1) (resp. Case (2)). As $\sigma = \pi_{H, H'}|_B$ in Case (1) (resp. $\sigma = c_{H, h}|_B$ in Case (2)) and $\sigma(B) \subseteq B'$, we have $B \subseteq B_{\bar{\phi}(\delta')}$, as desired.

¹⁰Note that for $Hg \in B$, we have $\delta \equiv 1 \pmod{\mathfrak{m}_{Hg}}$ and hence $(1 - \delta) \subseteq \mathfrak{m}_{Hg}$.

Next we prove $\phi_{\delta, \delta'}^{-1}(\mathfrak{m}_{Hg}/(1-\delta)) = \mathfrak{m}_{\sigma(Hg)}/(1-\delta')$, or equivalently, $\bar{\phi}(\mathfrak{m}_{\sigma(Hg)}) \subseteq \mathfrak{m}_{Hg}$. In Case (1), we have $\sigma(Hg) = \pi_{H, H'}(Hg) = H'g$, and

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

As ϕ is the natural inclusion, we have $\bar{\phi}(\mathfrak{m}_{\sigma(Hg)}) \subseteq \mathfrak{m}_{Hg}$, as desired.

In Case (2), we have $\sigma(Hg) = c_{H, h}(Hg) = H'hg$, and

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

As ϕ sends $x \in L^{H'}$ to $h^{-1}x \in L^H$, again we have $\bar{\phi}(\mathfrak{m}_{\sigma(Hg)}) \subseteq \mathfrak{m}_{Hg}$. This proves $\phi_{\delta, \delta'}^{-1}(\mathfrak{m}_{Hg}/(1-\delta)) = \mathfrak{m}_{\sigma(Hg)}/(1-\delta')$.

The maximal ideals of $\bar{\mathcal{O}}_K/(1-\delta)$ are those of the form $\mathfrak{m}_{Hg}/(1-\delta)$ where $(1-\delta) \subseteq \mathfrak{m}_{Hg}$. The condition $(1-\delta) \subseteq \mathfrak{m}_{Hg}$ holds iff $\delta \equiv 1 \pmod{\mathfrak{m}_{Hg}}$, which holds iff $Hg \in B_\delta = B$. So the maximal ideals of $\bar{\mathcal{O}}_K/(1-\delta)$ are of the form $\mathfrak{m}_{Hg}/(1-\delta)$ where $Hg \in B$. Similarly, the maximal ideals of $\bar{\mathcal{O}}_{K'}/(1-\delta')$ are of the form $\mathfrak{m}_{H'g}/(1-\delta')$ where $H'g \in B'$. As $\bar{\mathcal{O}}_K$ and $\bar{\mathcal{O}}_{K'}$ are finite products of copies of \mathbb{F}_p , so are $\bar{\mathcal{O}}_K/(1-\delta)$ and $\bar{\mathcal{O}}_{K'}/(1-\delta')$. So their dimensions over \mathbb{F}_p equals their numbers of maximal ideals. It follows that $\dim_{\mathbb{F}_p} \bar{\mathcal{O}}_K/(1-\delta) = |B|$ and $\dim_{\mathbb{F}_p} \bar{\mathcal{O}}_{K'}/(1-\delta') = |B'|$.

Now assume $\sigma : B \rightarrow B'$ is a bijection. So $|B| = |B'|$. We know

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

As σ is surjective, this shows $\phi_{\delta, \delta'}^{-1}(0)$ is contained in the intersection of all the maximal ideals of $\bar{\mathcal{O}}_{K'}/(1-\delta')$, which is zero since $\bar{\mathcal{O}}_{K'}/(1-\delta')$ is semisimple. So $\phi_{\delta, \delta'}$ is injective. As $\dim_{\mathbb{F}_p} \bar{\mathcal{O}}_K/(1-\delta) = |B|$ equals $\dim_{\mathbb{F}_p} \bar{\mathcal{O}}_{K'}/(1-\delta') = |B'|$, $\phi_{\delta, \delta'}$ is a ring isomorphism. \square

Lemma 4.20 now follows naturally:

Proof of Lemma 4.20. Assume \mathcal{C} is a \mathcal{P} -scheme but not a strongly antisymmetric \mathcal{P} -scheme. Then there exist a sequence of subgroups $H_0, \dots, H_k \in \mathcal{P}$, blocks $B_0 \in C_{H_0}, \dots, B_k \in C_{H_k}$, and bijective maps $\sigma_i : B_{i-1} \rightarrow B_i$ for $i = 1, \dots, k$, such that (1) each σ_i is of the form $c_{H_{i-1}, g}|_{B_{i-1}}$, $\pi_{H_{i-1}, H_i}|_{B_{i-1}}$, or $(\pi_{H_i, H_{i-1}}|_{B_i})^{-1}$, (2) $H_0 = H_k$ and $B_0 = B_k$, and (3) $\tau := \sigma_k \circ \dots \circ \sigma_1$ is a nontrivial permutation of $B_0 = B_k$. For $i = 0, \dots, k$, let $\delta_i = \delta_{B_i} \in I_{H_i}$, and define $\phi_i : \bar{\mathcal{O}}_{L^{H_i}}/(1-\delta_i) \rightarrow \bar{\mathcal{O}}_{L^{H_{i-1}}}/(1-\delta_{i-1})$ as follows:

$$\phi_i := \begin{cases} \phi_{\delta_{i-1}, \delta_i} & \sigma_i \text{ is of the form } c_{H_{i-1}, g}|_{B_{i-1}} \text{ or } \pi_{H_{i-1}, H_i}|_{B_{i-1}}, \\ \phi_{\delta_i, \delta_{i-1}}^{-1} & \sigma_i \text{ is of the form } (\pi_{H_i, H_{i-1}}|_{B_i})^{-1}. \end{cases}$$

Each ϕ_i is a well defined ring isomorphism by the last claim in Lemma C.5. Let $\phi := \phi_1 \circ \dots \circ \phi_k$. Applying Lemma C.5 to each σ_i , we see $\phi_i^{-1}(\mathfrak{m}_{H_{i-1}g}/(1-\delta_{i-1})) = \mathfrak{m}_{\sigma_i(H_{i-1}g)}/(1-\delta_i)$ for $H_{i-1}g \in B_{i-1}$. Composing the maps σ_i and the maps ϕ_i , we see $\phi^{-1}(\mathfrak{m}_{H_0g}/(1-\delta_0)) = \mathfrak{m}_{\tau(H_0g)}/(1-\delta_0)$ for $H_0g \in B_0$. As τ is a nontrivial permutation of B_0 , ϕ is a nontrivial ring automorphism of $\bar{\mathcal{O}}_{L^{H_0}}/(1-\delta_0)$.

By identifying L^H with $K_H \in \mathcal{F}$ for each $H \in \mathcal{P}$, we see that the graph (V, E) contains a cycle C such that the composition of the labels along C is a nontrivial ring automorphism of $\bar{\mathcal{O}}_K/(1-\delta)$ for some $(K, \delta) \in V$. The subroutine then finds such a ring automorphism and uses it to modify I_K . \square