

TRADING GRH FOR ALGEBRA: ALGORITHMS FOR FACTORING POLYNOMIALS AND RELATED STRUCTURES

GÁBOR IVANYOS, MAREK KARPINSKI, LAJOS RÓNYAI, AND NITIN SAXENA

ABSTRACT. In this paper we develop a general technique to eliminate the assumption of the Generalized Riemann Hypothesis (GRH) from various deterministic polynomial factoring algorithms over finite fields. It is the first bona fide progress on that issue for more than 25 years of study of the problem. Our main results are basically of the following form: either we construct a nontrivial factor of a given polynomial or compute a nontrivial automorphism of the factor algebra of the given polynomial. Probably the most notable application of such automorphisms is efficiently finding zero divisors in noncommutative algebras. The proof methods used in this paper exploit virtual roots of unity and lead to efficient actual polynomial factoring algorithms in special cases.

1. INTRODUCTION

Factoring polynomials over finite fields (FPPF, for short) belongs to the fundamental computational problems. There are many computational tasks for which known algorithms first require factoring polynomials. Thus, polynomial factoring was an intensely studied question and various randomized polynomial time algorithms are known [Be67], [Rab80], [CZ81], [GS92], [KS98], [KU08]. As the polynomial is assumed to be given as an array of its coefficients, the input size is approximately $n \log |k|$ where k stands for the ground field and n is the degree of the polynomial. Thus polynomial time means time polynomial in both n and $\log |k|$. In addition to its practical significance, FPPF occupies a very special place in the landscape of complexity classes. Together with polynomial identity testing (see for e.g. [KI03]), it is one of the two major specific problems related to the celebrated $BPP = P$ question. In fact, FPPF is known to be $RP \cap coRP$ -easy, and indeed admits nice and practical randomized algorithms (whose roots can be traced as far back as Legendre), but has resisted decades of efforts to devise deterministic polynomial time algorithms. Note that in [Be67], a deterministic algorithm is given which runs in time polynomial in n and $|k|$ (more precisely, polynomial in n , $\log |k|$ and p , where p is the characteristic of k).

On the basis of the Generalized Riemann Hypothesis (GRH) several important subproblems and special cases can be solved in deterministic polynomial time. These results fit well the line of inquiry put forward by [KI03] (and many others): try to provide derandomization without (complexity theoretic) hardness assumptions. Interestingly enough, here a central open problem of Computer Science (circuit lower bounds) gives way to a central open problem of pure mathematics

Received by the editor September 12, 2009 and, in revised form, November 10, 2010.
2010 *Mathematics Subject Classification.* Primary 68W30, 12Y05, 16Z05.

(the Riemann Hypothesis). The surprising connection of GRH with polynomial factoring is based on the fact that if GRH is true and r is a prime dividing $(|k| - 1)$, then one can find r -th nonresidues in the finite field k , which can then be used to factor “special” polynomials, $x^r - a$ over k , in deterministic polynomial time (see [Hua85]).

Based on GRH, many deterministic factoring algorithms are known, but all of them are super-polynomial time except on special instances.

Degree has a small factor. The special instance when the degree n of the input polynomial $f(x)$ has a “small” prime factor r has been particularly interesting. Rónyai [Ró87] showed that under GRH one can find a nontrivial factor of $f(x)$ in deterministic polynomial time. Later it was shown by Evdokimov [Ev94] that Rónyai’s algorithm can be modified to get under GRH a deterministic algorithm that factors *any* input polynomial $f(x) \in k[x]$ of degree n in *sub-exponential* time $\text{poly}(n^{\log n}, \log |k|)$. This line of approach has since been investigated, in an attempt to remove GRH or improve the time complexity, leading to several algebraic-combinatorial conjectures and quite special case solutions [CH00, Gao01, IKS08].

Galois group. Some other instances studied have been related to the *Galois group* of the given polynomial over rationals. Rónyai [Ró89b] showed under GRH that any polynomial $f(x) \in \mathbb{Z}[x]$ can be factored modulo p deterministically in time polynomial in the size of the Galois group over \mathbb{Q} of f and $\log p$, except for finitely many primes p . Other results of a similar flavor are: Huang [Hua85] showed under GRH that $f(x)$ can be factored in deterministic polynomial time if it has an *Abelian* Galois group while Evdokimov [Ev89] showed under GRH that $f(x)$ can be factored in deterministic polynomial time if it has a *solvable* Galois group.

Special fields. Another instance studied is that of “special” finite fields. Bach, von zur Gathen and Lenstra [BGL01] showed under GRH that polynomials over finite fields of characteristic p can be factored in deterministic polynomial time if $\Phi_k(p)$ is “smooth” for some integer k , where $\Phi_k(x)$ is the k -th cyclotomic polynomial. This result generalizes the previous works of Rónyai [Ró89a], Mignotte and Schnorr [MS88], von zur Gathen [G87], Camion [Cam83] and Moenck [Moe77].

Application to finite algebra questions. Polynomial factoring has several applications both in the real world — coding theory and cryptography — and in fundamental computational algebra problems. The latter kind of application is relevant to this work. Friedl and Rónyai [FR85] studied the computational problem of finding the simple components and a zero divisor of a given finite algebra over a finite field. They showed that all these problems depend on factoring polynomials over finite fields and hence have randomized polynomial time algorithms. Furthermore, they have under GRH deterministic quasipolynomial time algorithms.

As we saw above there are several results on polynomial factoring that assume the truth of the GRH. Of course one would like to eliminate the need of GRH but that goal is still elusive. Most notably, at present we cannot give an unconditional polynomial time algorithm even for computing square roots in finite fields. However, we are able to make progress in the desired direction: While during the course of most of the algorithms mentioned above, GRH is used to take r -th roots of field elements at several places (and for various numbers r), the typical GRH-free versions of this paper come up either with a proper factor or with an automorphism

of the algebra closely related to the polynomial. As such automorphisms can be used to factor polynomials under GRH, our results can be interpreted as pushing GRH to the end of the factoring algorithms. Also, our techniques turn out to be powerful enough to achieve efficient GRH-free *factoring* algorithms in special cases. But probably the most interesting application is finding zero divisors in noncommutative algebras over finite fields in deterministic quasipolynomial time *without* needing GRH.

1.1. Our main results and techniques. Results related to given groups of automorphisms of algebras analogous to Galois theory of finite fields play a crucial role in the algorithms of the present paper.

Commutative algebras. The most notable among results of this type is the following.

Theorem 1.1. *Let \mathcal{A} be a finite dimensional commutative and associative algebra over the finite field k . Assume that we are given t automorphisms (as matrices in terms of a basis of \mathcal{A}) which generate a noncyclic group. Then in deterministic polynomial time (in $\log |k|$, t and $\dim_k \mathcal{A}$) we can find a zero divisor in \mathcal{A} .*

For every integer m it is straightforward to construct a group of automorphisms of $\mathbb{F}_p[X]/\Phi_m(X)\mathbb{F}_p[X]$ which is isomorphic to the multiplicative group $(\mathbb{Z}/m\mathbb{Z})^*$ of the reduced residue classes modulo m . This group is cyclic iff $m \leq 4$ or m is an odd prime power or m is two times an odd prime power whence we obtain the following.

Corollary 1.2. *Assume that $m > 0$ is an integer which is neither a power of an odd prime nor two times a power of an odd prime. Then one can find a proper divisor of the cyclotomic polynomial $\Phi_m(X)$ in $\mathbb{F}_p[X]$ in deterministic $\text{poly}(m, \log p)$ time.*

(Note: for $m \leq 4$ we can completely factor $\Phi_m(X)$ as square roots of “small” numbers can be found by [Sch85].) To our knowledge the above result gives the first deterministic polynomial time algorithm to nontrivially factor “most” of the cyclotomic polynomials without assuming GRH. (There are some results known for very restricted cyclotomic polynomials; see [S96, S01].)

Our proof for Theorem 1.1, following the seminal work of Lenstra [L91] on constructing isomorphisms between finite fields, is based on further generalizations of classical Galois theory constructs like cyclotomic extensions, Kummer extensions, Teichmüller subgroups, to the case of commutative semisimple algebras with automorphisms. In turn, Theorem 1.1 can be considered as a generalization of Lenstra’s result; see Subsection 4.4 for a more formal discussion showing this.

It turns out that in many cases we are able to develop unconditional counterparts of known deterministic factoring algorithms which rely on GRH. The time complexity of the new algorithm is polynomially equivalent to the original one, the tradeoff for dispensing with GRH is that we either find a nontrivial factor or a nontrivial automorphism of a related algebra. Our most notable result of this flavor is the following.

Theorem 1.3. *Let $f(X)$ be a polynomial of degree n over the finite field k . Then there is a deterministic algorithm which in quasipolynomial time $\text{poly}(n^{\log n}, \log |k|)$ computes either a proper divisor of $f(X)$ in $k[X]$ or a k -automorphism of order n of the algebra $k[X]/f(X)k[X]$.*

This theorem can be considered as a GRH-free version of Evdokimov’s factoring result [Ev94]. Besides its application to noncommutative algebras the result is of interest in its own right. It is the first unconditional deterministic quasipolynomial time algorithm to find a nontrivial automorphism of a given commutative semisimple algebra over a finite field. Finding a nontrivial automorphism of a given arbitrary ring is in general as hard as integer factoring [KS05], but our result shows that it might be a lot easier for a commutative semisimple algebra over a finite field. Note that in the case when $f(X)$ splits over k as $\prod_{j=1}^n (X - \alpha_j)$, with $\alpha_1, \dots, \alpha_n$ all distinct, the above algorithm either finds a nontrivial factor of $f(X)$ — or it gives an automorphism σ of $\mathcal{A} = k[X]/f(X)k[X]$ of order n , thus yielding n distinct “roots” of $f(X) - x$, $\sigma(x), \dots, \sigma^{n-1}(x)$ — all living in $\mathcal{A} \setminus k$. This latter case can be interpreted as finding roots over finite fields in terms of “radicals”, in analogy to classical Galois theory where one studies rational polynomials whose roots can be expressed by radicals; see Section 4 for details. We also remark that — using arguments similar to those we used to derive Lenstra’s result from Theorem 1.1 — it is easy to derive Evdokimov’s result from Theorems 1.1 and 1.3. In view of this, Theorem 1.3 can also be interpreted essentially as pushing the use of GRH to the final step in Evdokimov’s algorithm.

Proof idea of Theorem 1.3. Evdokimov’s algorithm uses GRH for taking r -th roots of field elements in recursion for various primes r . To obtain a GRH-free version, the first idea would be adjoining virtual roots to the ground field and working with the algebra obtained this way in place of the base field. As we do not have satisfactory control over the set of primes r for which we need r -th roots, the dimension of the algebra replacing the base field can become exponentially large (or at least we are unable to prove that such blowup does not happen.) Therefore, we do not add roots permanently, instead we work with automorphisms of algebras and use virtual roots locally: only when they are needed for operations with automorphisms, e.g., computing zero divisors when we encounter noncyclic groups of automorphism, “bringing down” automorphisms to subalgebras or gluing an automorphism to another one, given on the subalgebra of the elements fixed by the former.

Our method uses a recursive process. During an iteration we work with a pair of semisimple algebras $\mathcal{B} \leq \mathcal{A}$ over the base field k . Initially, $\mathcal{A} = k[X]/(f(X))$ and $\mathcal{B} = k$, in each subsequent recursive call the algebras themselves might get bigger, but $\text{rk}_{\mathcal{B}} \mathcal{A}$ will be at least halved. This corresponds to Evdokimov’s main idea of attempting to factor polynomials over algebras obtained by adjoining some roots of the original polynomial to be factored to the base field. We attempt to find a nontrivial automorphism of \mathcal{A} which acts on \mathcal{B} trivially. The key idea in finding such an automorphism is to consider a special ideal \mathcal{A}' (what we call the *essential part* in Section 5.2) of the tensor product $\mathcal{A} \otimes_{\mathcal{B}} \mathcal{A}$. The ideal \mathcal{A}' is just the kernel of the standard homomorphism of $\mathcal{A} \otimes_{\mathcal{B}} \mathcal{A}$ onto \mathcal{A} given by the multiplication in \mathcal{A} and has rank (“dimension”) $\text{rk}_{\mathcal{B}} \mathcal{A}(\text{rk}_{\mathcal{B}} \mathcal{A} - 1)$ over \mathcal{B} if \mathcal{A} is a free \mathcal{B} -module. The algebra \mathcal{A} is naturally embedded in \mathcal{A}' by a map ϕ , hence \mathcal{A}' is an extension algebra of $\phi(\mathcal{A}) \cong \mathcal{A}$ which in turn is an extension algebra of $\phi(\mathcal{B}) \cong \mathcal{B}$. The advantage of working with \mathcal{A}' is that we know a natural automorphism of \mathcal{A}' fixing \mathcal{B} — the map $\tau : x \otimes y \mapsto y \otimes x$. A lot of technical effort goes into “bringing down” this automorphism (or a certain other automorphism σ of order 2 obtained by recursion) from \mathcal{A}' to \mathcal{A} , i.e., getting a \mathcal{B} -automorphism σ' of \mathcal{A} . The technical

arguments fall into two cases, depending on whether $\text{rk}_{\mathcal{A}} \mathcal{A}' = \text{rk}_{\mathcal{B}} \mathcal{A}' / \text{rk}_{\mathcal{B}} \mathcal{A}$ is odd or even.

(1) If the rank $\text{rk}_{\mathcal{B}} \mathcal{A}$ is even, then $\text{rk}_{\mathcal{A}} \mathcal{A}'$ is odd. We find an element $u \in \mathcal{A}'$ with $u^\tau = -u$. If $u \in \mathcal{A}$, then the restriction of τ is a nontrivial \mathcal{B} -automorphism of the subalgebra $\mathcal{B}[u]$ of \mathcal{A} generated by \mathcal{B} and u . If $u \notin \mathcal{A}$, then either the subalgebra $\mathcal{A}[u]$ of \mathcal{A}' is not a free \mathcal{A} -module or \mathcal{A}' is not a free $\mathcal{A}[u]$ -module. Both cases give us a zero divisor in \mathcal{A}' , and allow us to go to a smaller ideal \mathcal{I} of \mathcal{A}' such that we know an automorphism of \mathcal{I} , it contains a “copy” of \mathcal{A} and $\text{rk}_{\mathcal{A}} \mathcal{I}$ is odd. Thus we can continue this “descent” (from \mathcal{A}' to \mathcal{I}) till we have a \mathcal{B} -automorphism of \mathcal{A} or of a subalgebra of \mathcal{A} (this process appears in Section 5.1). In the former case we are done, while in the latter case we use two recursive calls and certain techniques to “glue” the three available automorphisms. (The gluing process is described in Section 4.6.)

(2) If the rank $\text{rk}_{\mathcal{B}} \mathcal{A}$ is odd, then $\text{rk}_{\mathcal{A}} \mathcal{A}'$ is even and we can use the technique above to find an \mathcal{A} -automorphism σ of \mathcal{A}' . It turns out that σ and τ generate a group of automorphisms of \mathcal{A}' which is big enough to find a proper ideal \mathcal{I} of \mathcal{A}' efficiently. We may further assume that the rank of \mathcal{I} over \mathcal{A} is at most $\text{rk}_{\mathcal{A}} \mathcal{A}' / 2 = (\text{rk}_{\mathcal{B}} \mathcal{A} - 1) / 2$. This allows us a recursive call with $(\mathcal{I}, \mathcal{A})$ in place of $(\mathcal{A}, \mathcal{B})$ to get an \mathcal{A} -automorphism of \mathcal{I} , which we eventually show is enough to extract an automorphism of \mathcal{A} using tensor properties and a recursive call (this case gets handled in Section 5.3).

This algebraic-extensions jugglery *either* goes through and yields a nontrivial automorphism σ' of \mathcal{A} fixing \mathcal{B} or it “fails” and yields a zero divisor in \mathcal{A} which we use to “break” \mathcal{A} into smaller subalgebras and continue working from there. As in each recursive call, in the above two cases, the rank of the bigger algebra over the subalgebra is at most half of the original one (the *invariant* condition), the depth of the recursion is at most $\log \text{rk}_{\mathcal{B}} \mathcal{A}$. The *termination* condition is: the rank of the bigger algebra over the subalgebra is one. This gives the dominating $n^{\log n}$ term in the time complexity analysis.

Galois group. The techniques used to prove Theorems 1.1 and 1.3 can be applied to the instance of polynomial factoring over prime fields when we know the Galois group of the input polynomial. The following theorem can be seen as the GRH-free version of the main theorem of Rónyai [Ró89b].

Theorem 1.4. *Let $F(X) \in \mathbb{Z}[X]$ be a polynomial irreducible over \mathbb{Q} with Galois group of size m and let L be the maximum length of the coefficients of $F(X)$. Let p be a prime not dividing the discriminant of $F(X)$ and let $f(x) = F(X) \pmod{p}$. Then by a deterministic algorithm of running time $\text{poly}(m, L, \log p)$ we can find either a nontrivial factor of $f(x)$ or a nontrivial automorphism of $\mathbb{F}_p[x]/(f(x))$ of order $\deg f$.*

Rational polynomials known to have small but noncommutative Galois groups also emerge in various branches of mathematics and its applications. For example, the six roots of the polynomial $F_j(X) = (X^2 - X + 1)^3 - \frac{j}{28} X^2 (X - 1)^2$ are the possible parameters λ of the elliptic curves from the *Legendre family* E_λ having prescribed j -invariant j ; see [Hu86]. (Recall that the curve E_λ is defined by the equation $Y^2 = X(X - 1)(X - \lambda)$.) The Galois group of $F_j(X)$ is S_3 , whence Theorem 1.1 gives a nontrivial factorization of the polynomial $F_j(X)$ modulo p where p is odd and j is coprime to p .

Special fields. The next application of the techniques used to prove Theorems 1.1 and 1.3 is in the instance of polynomial factoring over \mathbb{F}_p when p is a prime with smooth $(p-1)$. The following theorem can be seen as the GRH-free version of the main theorem of Rónyai [Ró89a].

Theorem 1.5. *Let $f(x)$ be a polynomial of degree n , that splits into linear factors over \mathbb{F}_p . Let $r_1 < \dots < r_t$ be the prime factors of $(p-1)$. Then by a deterministic algorithm of running time $\text{poly}(r_t, n, \log p)$, we can find either a nontrivial factor of $f(x)$ or a nontrivial automorphism of $\mathbb{F}_p[x]/(f(x))$ of order n . In fact, we always find a nontrivial factor of $f(x)$ in the case $n \nmid \text{lcm}\{r_i - 1 \mid 1 \leq i \leq t\}$.*

Thus over “special” fields (i.e., when $p-1$ has only small prime factors) the above result actually gives a deterministic polynomial time algorithm, a significant improvement over Theorem 1.3.

We succeeded in obtaining GRH-free versions of most of the known GRH-dependent results we considered so far. The most notable exception which withstood our efforts is the result of Bach, von zur Gathen and Lenstra [BGL01] for the case when $\Phi_k(p)$ is smooth. An even more important limitation of our results is that they do not provide (direct) tools for computing square, cubic, etc., roots in general finite fields. Rather, they provide methods for circumventing explicit computations of those.

Noncommutative algebras. Theorem 1.1 and Corollary 1.2 demonstrate that finding automorphisms of algebras can be a useful tool in certain factoring algorithms. The following result gives a direct evidence for the power of this tool in computing the structure of noncommutative algebras.

Theorem 1.6. *Let \mathcal{A} be a finite dimensional associative algebra over the finite field k . Assume that we are given a commutative subalgebra \mathcal{B} of \mathcal{A} as well as a nontrivial automorphism σ of \mathcal{B} whose restriction to the intersection of \mathcal{B} with the center of \mathcal{A} is the identity map. Then in deterministic polynomial time (in $\log |k| + \dim_k \mathcal{A}$) we can find a zero divisor in \mathcal{A} .*

Theorem 1.3 and its proof techniques have important applications. The first one is that — together with Theorem 1.6 — it gives a quasipolynomial time deterministic algorithm for finding zero divisors in noncommutative algebras.

Theorem 1.7. *Let \mathcal{A} , an associative algebra of dimension n over the finite field k , be given. Assume that \mathcal{A} is noncommutative. Then there is a deterministic algorithm which finds a zero divisor in \mathcal{A} in time $\text{poly}(n^{\log n}, \log |k|)$.*

The previous best result in this direction was due to Rónyai [Ró90] who gave an algorithm invoking polynomial factorization over finite fields and hence taking quasipolynomial time assuming GRH. Our result removes the GRH assumption. It is interesting to note that if we prove such a result for *commutative* algebras as well, then we would basically be able to factor polynomials in quasipolynomial time without needing GRH.

If \mathcal{A} is a finite simple algebra over the finite field k , then by a theorem of Wedderburn, it is isomorphic to the algebra $M_m(K)$ of the $m \times m$ matrices with entries from an extension field K of k . By Theorem 1.7 we find a proper left ideal of \mathcal{A} . A recursive call to a certain subalgebra of the left ideal will ultimately give a minimal left ideal of \mathcal{A} and using this minimal one-sided ideal an isomorphism

with $M_m(K)$ can be efficiently computed. Actually, if m has small prime factors, instead of the method of Theorem 1.7, we can also use a variant which is based on our unconditional version of [R687]. We obtain the following.

Theorem 1.8. *Let K be a finite field. Given an algebra \mathcal{A} which is isomorphic to $M_m(K)$ one can construct an isomorphism of \mathcal{A} with $M_m(K)$ in time $\text{poly}(m^{\min(r, \log m)}, \log |K|)$, where r is the largest prime factor of m .*

In particular, one can solve the explicit isomorphism problem in polynomial time (that is, in time polynomial in m and $\log |K|$) for algebras isomorphic to $M_m(K)$ where m is a power of two.

For constant m , or more generally, for m having prime factors of constant size only, Theorem 1.8 extends Lenstra's result (on computing isomorphisms between input fields) to noncommutative simple algebras, i.e, the *explicit isomorphism problem* is solved in this case. We note that, in general, the problem of finding an isomorphism between finite algebras over a finite field is not "believed" to be NP-hard but it is at least as hard as the graph isomorphism problem [KS05]. We also remark that the analogous problem over the rationals has a surprising application to rational parametrization of curves; see [GHPS06].

1.2. Organization. In Section 2 we fix the notation and terminology used throughout the paper and recall various standard concepts and structural facts associated to algebras. We also discuss the three basic methods that lead to discovering a zero divisor in an algebra: finding discrete logs for elements of prime-power order, finding a free basis of a module and refining an ideal by a given automorphism.

In this work we use methods for finding zero divisors in algebras in the case when certain groups of automorphisms are given. One of these methods is computing fixed subalgebras and testing freeness over them. In Section 3 we give a characterization of algebras and groups which survive these kinds of attacks. These algebras, called *semiregular* with respect to the group, behave like fields in the sense that the whole algebra is a free module over the subalgebra of fixed points of the group and the rank equals the size of the group.

In Section 4 we build a small theory for the main algebraic construction, *Kummer-type extensions* of algebras, that we are going to use. We investigate there the action of the automorphisms of an algebra \mathcal{A} on a certain subgroup, the *Teichmüller subgroup*, of the multiplicative group of a Kummer-type extension of \mathcal{A} . This theory leads to the proof of Theorem 1.1. The proof of Theorem 1.4 is also completed in this section using Theorem 4.7, which is a technical tool for bringing down large automorphism groups of finite algebras to ideals of subalgebras.

In Section 5 we apply the machinery of Section 4 to the tensor power algebras to obtain automorphism of algebras as stated in Theorem 5.6, which is actually a GRH-free version of the result of [R687]. The other main technical results proved in Section 5 are Theorem 5.8, a slightly stronger version of Theorem 1.3 and Theorem 5.9, a result of iterated application of the former theorem.

In Section 6 we use the techniques developed for Theorems 1.1 and 1.3 in the case of special finite fields and prove Theorem 6.3 which is a slight generalization of Theorem 1.5.

In Section 7 we prove Theorem 1.6, find suitable subalgebras of given noncommutative algebras to use our tools for finding automorphisms, and invoke Theorem 1.6 to finish the proof of Theorems 1.7 and 1.8.

2. PRELIMINARIES

We assume that the reader is familiar with basic algebraic notions such as fields, commutative and noncommutative rings, modules, homomorphisms, automorphisms. In this section we fix terminology and notation and recall the most important standard notions and facts that we use in this work. These can be found in standard algebra texts, for example, [La80].

We denote the set of numbers $\{1, \dots, n\}$ by $[n]$. Throughout this paper, unless stated otherwise, by a ring we mean a commutative ring with identity. If R is a ring, then by R^* we denote its group of units, i.e., the (multiplicative) group of elements of R that have a multiplicative inverse. Modules over R are assumed to be unital and finitely generated. (An R -module M is called unital if the identity element of R acts on M as the identity map.) An *associative R -algebra* or just R -algebra for short is a not necessarily commutative ring \mathcal{A} which is an R -module at the same time where the ring and module addition coincide and multiplication by elements of R commutes with multiplication by elements of \mathcal{A} (from both sides). Throughout this paper we assume that algebras have identity elements and — unless explicitly stated otherwise — by a subalgebra we mean a subalgebra containing the identity element of the whole algebra. Note that if \mathcal{B} is a commutative subalgebra of \mathcal{A} , then \mathcal{A} is a \mathcal{B} -module in a natural way. If, furthermore, \mathcal{B} is contained in the center of \mathcal{A} (that is, $ab = ba$ for every $a \in \mathcal{A}$ and for every $b \in \mathcal{B}$), then \mathcal{A} is a \mathcal{B} -algebra. An element $x \in \mathcal{A}$ is called a *zero divisor* if $x \neq 0$ and there exist nonzero $y, y' \in \mathcal{A}$ such that $yx = xy' = 0$.

For a finitely generated R -module M , a finite set $B \subset M$ is called a *free basis* of M if every element of M can be written in a unique way such as a sum $\sum_{b \in B} r_b b$ with $r_b \in R$. A *free module* is a module with a free basis. $|B|$ is called the *rank* of the free module M over R . Clearly, a vector space is a free module. A module is called a *cyclic* module if it is generated by one element.

In this work we will consider finite dimensional algebras \mathcal{A} over a finite field k . We assume that an algebra \mathcal{A} is always presented in the input-output in terms of a k -linear basis of \mathcal{A} i.e., there are *basis elements* $b_1, \dots, b_n \in \mathcal{A}$ such that $\mathcal{A} = kb_1 + \dots + kb_n$ and furthermore an array $(\alpha_{ij\ell}) \in k^{n \times n \times n}$ of scalars is given such that $b_i \cdot b_j = \sum_{\ell=1}^n \alpha_{ij\ell} b_\ell$ ($i, j \in [n]$). The scalars $\alpha_{ij\ell}$ are referred to as the *structure constants* of \mathcal{A} with respect to the basis b_1, \dots, b_n .

If \mathcal{B} is a subalgebra of the commutative k -algebra \mathcal{A} such that \mathcal{A} is also a free module over \mathcal{B} , then we call \mathcal{A} an *algebra extension* or an *extension algebra* over \mathcal{B} . We denote the rank (“dimension”) of \mathcal{A} as a \mathcal{B} -module by $\text{rk}_{\mathcal{B}} \mathcal{A}$ or $[\mathcal{A} : \mathcal{B}]$. We sometimes use this notation also when there is an implicit embedding of \mathcal{B} in \mathcal{A} .

We will make use of *tensor products*. If \mathcal{B} is a commutative algebra and $\mathcal{A}_1, \mathcal{A}_2$ are free \mathcal{B} -modules of ranks n_1, n_2 , respectively, then their tensor product $\mathcal{A}_1 \otimes_{\mathcal{B}} \mathcal{A}_2$ is a free \mathcal{B} -module of rank $n_1 n_2$. It is generated as a \mathcal{B} -module by the elements of the form $a_1 \otimes a_2$ ($a_i \in \mathcal{A}_i$). Furthermore, if \mathcal{A}_1 and \mathcal{A}_2 are \mathcal{B} -algebras, then the map $(a_1 \otimes a_2) \cdot (a'_1 \otimes a'_2) := (a_1 a'_1 \otimes a_2 a'_2)$ has a \mathcal{B} -homomorphic extension to $\mathcal{A}_1 \otimes \mathcal{A}_2$ making $\mathcal{A}_1 \otimes \mathcal{A}_2$ a \mathcal{B} -algebra.

In an algebra \mathcal{A} we call an element $x \in \mathcal{A}$ *nilpotent* if $x^m = 0$ for some $0 < m \in \mathbb{Z}$, while we call x *idempotent* if $x^2 = x \neq 0$. It is called a *primitive* idempotent if it cannot be expressed as the sum of two idempotents whose product is zero. It is called *nontrivial* if it is not 1.

An *ideal* of an R -algebra \mathcal{A} is an R -submodule which is at the same time a ring theoretic (two-sided) ideal. Note that if \mathcal{A} has an identity element, a ring theoretic ideal is automatically an algebra ideal. Note that $\{0\}$ and \mathcal{A} are ideals of \mathcal{A} , we call them *trivial* ideals. Also note that proper ideals are not subalgebras in the strict sense used in this paper.

An algebra \mathcal{A} is called *simple* if it has no nontrivial ideal. A finite dimensional algebra over a field is called *semisimple* if it is a direct sum of finitely many simple algebras. Finite dimensional commutative simple algebras are finite extensions of the base field and hence commutative semisimple algebras are isomorphic to direct sums of such extensions. A finite dimensional algebra \mathcal{A} over a field has a smallest ideal J such that the factor algebra \mathcal{A}/J is semisimple. It is called the *radical* of \mathcal{A} . The radical consists of nilpotent elements and if the ground field is finite it can be computed in deterministic polynomial time; see [R690, CIW96].

We will make use of some standard facts about idempotents and ideals in semisimple algebras.

Fact 2.1 (Ideals of commutative semisimple algebras). *Let \mathcal{A} be a commutative semisimple algebra over a field and let I be an ideal of \mathcal{A} . Then $I^\perp := \{a \in \mathcal{A} \mid aI = 0\}$ is also an ideal of \mathcal{A} (called the complement of I) and $\mathcal{A} = I \oplus I^\perp$. Furthermore, there exists an idempotent e of the center of \mathcal{A} such that $I = e\mathcal{A}$ and $I^\perp = (1 - e)\mathcal{A}$ thus giving an explicit projection from \mathcal{A} to I and I^\perp , respectively.*

Following is the celebrated *Artin-Wedderburn Theorem* that classifies semisimple algebras over finite fields.

Fact 2.2 (Artin-Wedderburn). *Any semisimple algebra \mathcal{A} over the finite field k is isomorphic to a direct sum of $n_i \times n_i$ matrix algebras over finite extensions K_i of k . Both the n_i 's and K_i 's are uniquely determined up to permutation of the indices i .*

2.1. Discrete Log for r -elements. Given two r -elements (i.e., having order a power of the prime r) in a commutative semisimple algebra, there is an algorithm that computes the discrete logarithm or finds a zero divisor (of a special form) in \mathcal{A} . We describe this algorithm below, it is a variant of the Pohlig-Hellman [PH78] algorithm with the equality testing of elements replaced by testing whether their difference is a zero divisor.

Lemma 2.3. *Given a prime r distinct from the characteristic of a finite field k , a commutative semisimple algebra \mathcal{A} over k and two r -elements $a, b \in \mathcal{A}^*$, such that the order of a is greater than or equal to the order of b . Then there is a deterministic algorithm which computes in time $\text{poly}(r, \log |\mathcal{A}|)$:*

- (1) *either two nonnegative integers s, s' such that $a^s - b^{s'}$ is a zero divisor in \mathcal{A} ,*
- or*
- (2) *an integer $s \geq 0$ with $a^s = b$.*

Proof. Let t_a be the smallest nonnegative integer such that $a^{r^{t_a}} - 1$ is zero or a zero divisor in \mathcal{A} . Since $t_a \leq \log_r |\mathcal{A}|$ we can compute $a^{r^0} - 1, a^{r^1} - 1, \dots, a^{r^{t_a}} - 1$ in $\text{poly}(\log |\mathcal{A}|)$ time via fast exponentiation. We are done if $0 \neq a^{r^{t_a}} - 1 = a^{r^{t_a}} - b^0$ is a zero divisor. Therefore we may assume that $a^{r^{t_a}} = 1$, i.e., the order of a is r^{t_a} . Let t_b be the smallest nonnegative integer such that $b^{r^{t_b}} - 1$ is a zero divisor. Like t_a, t_b can be computed in polynomial time and we may again assume that r^{t_b} is the

order of b . Replacing a with $a^{r^{t_a-t_b}}$ we may assume that $t_a = t_b = t$. In this case for every primitive idempotent e of \mathcal{A} : ea, eb have order r^t in the finite field $e\mathcal{A}$. As the multiplicative group of a finite field is cyclic, this means that there exists a nonnegative integer $s < r^t$ such that $(ea)^s = eb$. So we now attempt to find this discrete log, s , and the corresponding idempotent e as well.

We iteratively compute the consecutive sections of the base r expansion of s . To be more specific, we compute integers $s_0 = 0, s_1, s_2, \dots, s_t$ together with idempotents e_1, \dots, e_t of \mathcal{A} such that, for all $1 \leq j \leq t$: $0 \leq s_j < r^j$, $s_j \equiv s_{j-1} \pmod{r^{j-1}}$ and $a^{s_j r^{t-j}} e_j = b^{r^{t-j}} e_j$.

In the initial case $j = 1$ we find by an exhaustive search, in at most r rounds, $s_1 \in \{1, \dots, r-1\}$ such that $z_1 = (a^{r^{t-1} s_1} - b^{r^{t-1}})$ is zero or a zero divisor. If it is zero, then we set $e_1 = 1$; otherwise we compute, and set e_1 equal to, the identity element of the annihilator ideal $\{x \in \mathcal{A} | z_1 x = 0\}$.

Assume that for some $j < t$ we have found already s_j and e_j with the desired property. Then we find by an exhaustive search, in at most r rounds, an integer $d_{j+1} \in \{0, \dots, r-1\}$ such that $z_{j+1} = (a^{(s_j+r^j d_{j+1})r^{t-j-1}} - b^{r^{t-j-1}})$ is zero or a zero divisor. We set $s_{j+1} = (s_j + d_{j+1}r^j)$ and take as e_{j+1} the identity element of the annihilator ideal $\{x \in e_j \mathcal{A} | x z_{j+1} = 0\}$.

The above procedure clearly terminates in t rounds and using fast exponentiation can be implemented in $\text{poly}(r, \log |\mathcal{A}|)$ time. \square

2.2. Free bases of modules. One of the possible methods for finding zero divisors in algebras is attempting to compute a free basis of a module over it. The following lemma describes a basic tool to do that.

Lemma 2.4. *Let V be a finitely generated module over a finite dimensional algebra \mathcal{A} over a finite field k . If V is not a free \mathcal{A} -module, then one can find a zero divisor in \mathcal{A} deterministically in time $\text{poly}(\dim_k V, \log |\mathcal{A}|)$.*

Proof. We give an algorithm that attempts to find a free basis of V over \mathcal{A} , but as there is no free basis it ends up finding a zero divisor.

Pick a nonzero $v_1 \in V$. We can efficiently check whether a nonzero $x \in \mathcal{A}$ exists such that $xv_1 = 0$, and also find it by linear algebra over k . If we get such an x , then it is a zero divisor, for otherwise x^{-1} would exist implying $v_1 = 0$. So suppose such an x does not exist, hence $V_1 := \mathcal{A}v_1$ is a free \mathcal{A} -module. Now $V_1 \neq V$ so find a $v_2 \in V \setminus V_1$ by linear algebra over k . Again we can efficiently check whether a nonzero $x \in \mathcal{A}$ exists such that $xv_2 \in V_1$, and also find it by linear algebra over k . If we get such an x , then it is a zero divisor, for otherwise x^{-1} would exist implying $v_2 \in V_1$. So suppose such an x does not exist, hence $V_2 := \mathcal{A}v_1 + \mathcal{A}v_2$ is a free \mathcal{A} -module. Now $V_2 \neq V$, so we can find a $v_3 \in V \setminus V_2$ by linear algebra over k and continue this process. This process will, in at most $\dim_{\mathcal{A}} V$ iterations, yield a zero divisor as V is not a free \mathcal{A} -module. \square

2.3. Automorphisms and invariant ideal decompositions. Automorphisms of a semisimple k -algebra \mathcal{A} are assumed to be given as linear transformations of the k -vector space \mathcal{A} in terms of a k -linear basis of \mathcal{A} . For images we use the superscript notation while for the fixed points the subscript notation: if σ is an automorphism of \mathcal{A} , then the image of $x \in \mathcal{A}$ under σ is denoted by x^σ . If Γ is a set of automorphisms of \mathcal{A} , then \mathcal{A}_Γ denotes the set of the elements of \mathcal{A} fixed by

every $\sigma \in \Gamma$. It is obvious that \mathcal{A}_Γ is a subalgebra of \mathcal{A} . For a single automorphism σ we use \mathcal{A}_σ in place of $\mathcal{A}_{\{\sigma\}}$.

Given an ideal I of \mathcal{A} and an automorphism σ of \mathcal{A} we usually try to find zero divisors from the action of σ on I . Note that, by Fact 2.1, $\mathcal{A} = I \oplus I^\perp$. Now I^σ is an ideal of \mathcal{A} , and if it is neither I nor I^\perp , then we try computing $I \cap I^\sigma$. This can be easily computed by first finding the identity element e of I , and then $I \cap I^\sigma$ is simply $\mathcal{A}ee^\sigma$. By the hypothesis this will be a proper ideal of I , thus leading to a *refinement* of the decomposition: $\mathcal{A} = I \oplus I^\perp$. This basic idea can be carried all the way to give the following tool that finds a refined, invariant, ideal decomposition.

Lemma 2.5. *Given \mathcal{A} , a commutative semisimple algebra over a finite field k together with a set of k -automorphisms Γ of \mathcal{A} and a decomposition of \mathcal{A} into a sum of pairwise orthogonal ideals J_1, \dots, J_s , there is a deterministic algorithm of time complexity $\text{poly}(|\Gamma|, \log |\mathcal{A}|)$ that computes a decomposition of \mathcal{A} into a sum of pairwise orthogonal ideals I_1, \dots, I_t such that:*

- (1) *the new decomposition is a refinement of the original one — for every $j \in \{1, \dots, t\}$, there exists $i \in \{1, \dots, s\}$ such that $I_j \subseteq J_i$, and*
- (2) *the new decomposition is invariant under Γ — the group generated by Γ permutes the ideals I_1, \dots, I_t , i.e., for every $\sigma \in \Gamma$ and for every index $j \in \{1, \dots, t\}$, we have $I_j^\sigma = I_{j^\sigma}$ for some index $j^\sigma \in \{1, \dots, t\}$.*

For a subalgebra (or ideal) \mathcal{B} of an algebra \mathcal{A} and $G \leq \text{Aut}(\mathcal{A})$, we denote the restriction of G to \mathcal{B} by $G|_{\mathcal{B}} := \{g \in G \mid \mathcal{B} \text{ is } g\text{-invariant i.e. } \mathcal{B}^g = \mathcal{B}\}$. Clearly, it is a subgroup of $\text{Aut}(\mathcal{B})$.

3. SEMIREGULARITY

In this section we assume that \mathcal{A} is a commutative semisimple algebra over a finite field k . Given $\Gamma \subseteq \text{Aut}_k(\mathcal{A})$, a basis of \mathcal{A}_Γ can be computed by solving a system of linear equations in \mathcal{A} . Thus, we can apply the method of Lemma 2.4 considering \mathcal{A} as a \mathcal{A}_Γ -module with respect to the multiplication in \mathcal{A} . In this section we describe a class of algebras, together with automorphisms, that are free modules over the subalgebra of the fixed points of the corresponding set of automorphisms, i.e., on which the tool of Lemma 2.4 is ineffective.

Let σ be a k -automorphism of \mathcal{A} . We say that σ is *fix-free* if there is no non-trivial ideal I of \mathcal{A} such that σ fixes I elementwise. We call a group $G \leq \text{Aut}(\mathcal{A})$ *semiregular* if every nonidentity element of G is fix-free. A single automorphism σ of \mathcal{A} is *semiregular* if σ generates a semiregular group of automorphisms of \mathcal{A} .

Example 3.1. *Consider the semisimple algebra $\mathcal{A} = \mathbb{F}_p \oplus \mathbb{F}_p \oplus \mathbb{F}_{p^2} \oplus \mathbb{F}_{p^2}$. It has an automorphism σ that swaps the two \mathbb{F}_p components, and also the two \mathbb{F}_{p^2} components. Then $G = \{1, \sigma\}$ is a semiregular group of automorphisms of \mathcal{A} .*

Note that $\mathcal{A}_G \cong \mathbb{F}_p \oplus \mathbb{F}_{p^2}$, and \mathcal{A} is a free \mathcal{A}_G -module.

We have the following characterization of semiregularity. It can be seen as a generalization of classical *Galois extension*.

Lemma 3.2. *Let \mathcal{A} be a commutative semisimple algebra over a finite field k and let G be a group of k -automorphisms of \mathcal{A} . Then $\dim_k \mathcal{A} \leq |G| \cdot \dim_k \mathcal{A}_G$, where equality holds if and only if G is semiregular. This condition is also equivalent to saying that \mathcal{A} is a free \mathcal{A}_G -module of rank $|G|$.*

Proof. The proof is based on the observation that \mathcal{A} is a direct sum of fields and a k -automorphism of \mathcal{A} just *permutes* these component fields. Note that an automorphism of \mathcal{A} will map a component field to one of the same size.

Let e be a primitive idempotent of \mathcal{A} . We denote the stabilizer of e in G by G_e , i.e., $G_e = \{\sigma \in G \mid e^\sigma = e\}$. Let C be a complete set of right coset representatives modulo G_e in G . The orbit of e under G is $\{e^\gamma \mid \gamma \in C\}$ and they are $|G : G_e|$ many pairwise orthogonal primitive idempotents in \mathcal{A} . (Note: there may be more primitive idempotents in \mathcal{A} in total.) This means that the component field $e\mathcal{A}$ is sent to the other component fields $\{e^\gamma\mathcal{A} \mid \gamma \in C\}$ by G . Thus, the element $f := \sum_{\gamma \in C} e^\gamma$ is a primitive idempotent of \mathcal{A}_G and equivalently $f\mathcal{A}_G$ is a field.

The subgroup G_e acts as a group of field automorphisms of $e\mathcal{A}$. This gives a restriction map $\lambda : G_e \rightarrow \text{Aut}_k(e\mathcal{A})$. The kernel $N_e = \{\sigma \in G_e \mid \sigma \text{ fixes } e\mathcal{A}\}$ of λ is a normal subgroup of G_e and the elements of the factor group G_e/N_e are distinct k -automorphisms of the field $e\mathcal{A}$. We claim that $(e\mathcal{A})_{G_e} = e\mathcal{A}_G$. The inclusion $e\mathcal{A}_G \subseteq (e\mathcal{A})_{G_e}$ is trivial. To see the reverse inclusion, let $x \in (e\mathcal{A})_{G_e}$ and consider $y := \sum_{\gamma \in C} x^\gamma$. Since $x \in e\mathcal{A}$ we get $ex = x$ and $y = \sum_{\gamma \in C} e^\gamma x^\gamma$, whence using the orthogonality of the idempotents e^γ , we infer $ey = x$. The fact that $y \in \mathcal{A}_G$ completes the proof of the claim. As G_e is a group of automorphisms of the field $e\mathcal{A}$, this claim implies $e\mathcal{A}_G$ is a field too and also by Galois theory $[e\mathcal{A} : e\mathcal{A}_G] = |G_e/N_e|$.

Observe that $ef = e$ and this makes multiplication by e , a surjective homomorphism from $f\mathcal{A}_G$ to $e\mathcal{A}_G$. This homomorphism is also injective as $e\mathcal{A}_G, f\mathcal{A}_G$ are fields, thus making $f\mathcal{A}_G \cong e\mathcal{A}_G$. Together with the fact that $f\mathcal{A}$ is a free $e\mathcal{A}$ -module of rank $|G : G_e|$ this implies that $\dim_{f\mathcal{A}_G} f\mathcal{A} = |G : G_e| \dim_{e\mathcal{A}_G} e\mathcal{A}$. Furthermore, from the last paragraph $\dim_{e\mathcal{A}_G} e\mathcal{A} = |G_e : N_e|$, thus $\dim_{f\mathcal{A}_G} f\mathcal{A} = |G : N_e| \leq |G|$. Finally, this gives $\dim_k f\mathcal{A} \leq \dim_k f\mathcal{A}_G \cdot |G|$. Applying this for all the primitive idempotents e of \mathcal{A} (and thus to all the corresponding primitive idempotents f of \mathcal{A}_G), we obtain the asserted inequality.

Observe that equality holds iff $|N_e| = 1$ for every primitive idempotent e of \mathcal{A} . In that case for every primitive idempotent e of \mathcal{A} , there is no nonidentity automorphism in G that fixes $e\mathcal{A}$, thus equivalently for every nontrivial ideal I of \mathcal{A} there is no nonidentity automorphism in G that fixes I elementwise. This means that equality holds iff G is semiregular.

Also, equality holds iff $\dim_{f\mathcal{A}_G} f\mathcal{A} = |G|$ for every primitive idempotent e of \mathcal{A} . The latter condition is equivalent to saying that every component field of \mathcal{A}_G has multiplicity $|G|$ in the \mathcal{A}_G -module \mathcal{A} , this in turn is equivalent to saying that \mathcal{A} is a free \mathcal{A}_G -module of rank $|G|$. \square

Using the above lemma we can decide semiregularity in an efficient way.

Proposition 3.3 (Checking semiregularity). *Given a commutative semisimple algebra \mathcal{A} over a finite field k , together with a set Γ of k -automorphisms of \mathcal{A} . Let G be the group generated by Γ . In deterministic $\text{poly}(|\Gamma|, \log |\mathcal{A}|)$ time one can list all the elements of G if G is semiregular, or one can find a zero divisor of \mathcal{A} if G is not semiregular.*

Proof. We first compute \mathcal{A}_Γ by linear algebra over k . We can assume that \mathcal{A} is a free \mathcal{A}_Γ -module; otherwise the algorithm in Lemma 2.4 finds a zero divisor. By Lemma 3.2, $|G| \geq \dim_{\mathcal{A}_\Gamma} \mathcal{A} =: m$ so try to enumerate $(m+1)$ different elements in the group G . If we fail, then, by Lemma 3.2, G is semiregular and we end up with a list of m elements that exactly comprise G .

If we do get a set S of $(m + 1)$ elements, then G is clearly not semiregular. Let e be a primitive idempotent of \mathcal{A} such that the subgroup $N_e \leq G$, consisting of automorphisms that fix $e\mathcal{A}$, is of maximal size. Then clearly $|G : N_e| \leq m$, which means, by the pigeon-hole principle, that in the set S there are two different elements σ_1, σ_2 such that $\sigma := \sigma_1\sigma_2^{-1} \in N_e$, thus σ fixes $e\mathcal{A}$. We now compute \mathcal{A}_σ and we know from this discussion that $e\mathcal{A} \subseteq \mathcal{A}_\sigma$. Thus we get two orthogonal component algebras $e\mathcal{A}_\sigma$ and $(1-e)\mathcal{A}_\sigma$ of \mathcal{A}_σ . We have from the proof of Lemma 3.2 that $e\mathcal{A}_\sigma = (e\mathcal{A})_\sigma = e\mathcal{A}$ while $(1-e)\mathcal{A}_\sigma = ((1-e)\mathcal{A})_\sigma \neq (1-e)\mathcal{A}$ (if $((1-e)\mathcal{A})_\sigma = (1-e)\mathcal{A}$, then σ would fix every element in \mathcal{A} and would be a trivial automorphism). As a result, \mathcal{A} is not a free module over \mathcal{A}_σ and hence we can find a zero divisor of \mathcal{A} using the method of Lemma 2.4. \square

As a warmup application of semiregularity we now show how to efficiently compute the size of the group of units of a given commutative semisimple algebra.

Lemma 3.4 (Computing $|\mathcal{A}^*|$). *Given a commutative semisimple finite algebra \mathcal{A} over a field k , we can compute $|\mathcal{A}^*|$ in deterministic $\text{poly}(\log |\mathcal{A}|)$ time.*

Proof. For concreteness we assume $k = \mathbb{F}_q$ and $n = \dim_k \mathcal{A}$. By the hypothesis there are integers e_i 's such that

$$\mathcal{A} \cong \bigoplus_{i=1}^n \mathbb{F}_{q^i}^{e_i},$$

where the notation $\mathbb{F}_{q^i}^{e_i}$ refers to a direct sum of e_i copies of the field. Let ϕ_q be the Frobenius automorphism of \mathcal{A} , i.e., $\phi_q(a) = a^q$ for all $a \in \mathcal{A}$, and define the group $G := \langle \phi_q \rangle$. Note that $\mathcal{A}_G \cong \mathbb{F}_q^e$, where $e := (e_1 + \dots + e_n)$.

If G is semiregular, then \mathcal{A} is a free \mathcal{A}_G -module, hence a free \mathbb{F}_q^e -module. In other words, all the component fields of \mathcal{A} are of the same size. Say, $\mathcal{A} \cong \mathbb{F}_{q^i}^e$. We can easily compute i , as $i = [\mathcal{A} : \mathcal{A}_G] = |G|$, and then e , as $e = (\dim_{\mathbb{F}_q} \mathcal{A})/i$. Thus, we can compute $|\mathcal{A}^*| = (q^i - 1)^e$.

If G is not semiregular, then by Proposition 3.3, we can find a zero divisor z in \mathcal{A} , and hence a nontrivial ideal $I := \mathcal{A}z$. By Fact 2.1, we get a nontrivial decomposition $\mathcal{A} = I \oplus J$. Now we can recursively compute $|I^*|$ and $|J^*|$. Finally, we output $|\mathcal{A}^*| = |I^*| \cdot |J^*|$. \square

Subgroup $G_{\mathcal{B}}$. Let G be a semiregular group of k -automorphisms of \mathcal{A} and let \mathcal{B} be a subalgebra of \mathcal{A} . We define $G_{\mathcal{B}}$ to be the subgroup of automorphisms of G that fix \mathcal{B} elementwise. We give below a Galois theory-like characterization of $G_{\mathcal{B}}$.

Proposition 3.5 (Subgroup-subalgebra correspondence). *Given a semiregular group G of automorphisms of a commutative semisimple algebra \mathcal{A} over a finite field k and a subalgebra \mathcal{B} of \mathcal{A} containing \mathcal{A}_G , one can find a zero divisor in \mathcal{A} in deterministic polynomial time unless $\mathcal{B} = \mathcal{A}_{G_{\mathcal{B}}}$.*

Proof. If \mathcal{A} is a field extension of k , then by Galois theory $\mathcal{B} = \mathcal{A}_{G_{\mathcal{B}}}$. If $|k| < (\dim_k \mathcal{A})^2$ and \mathcal{A} is not a field, then we can find a zero divisor in \mathcal{A} using Berlekamp's deterministic polynomial time algorithm. So for the rest of the proof we may assume that $|k| \geq (\dim_k \mathcal{A})^2$ and then the usual proof of existence of primitive elements in field extensions gives a deterministic polynomial time algorithm for finding a k -algebra generator x for \mathcal{A} ; see [GI00], i.e. $\mathcal{A} = k[x]$.

Let $|G| = d$. Compute a minimal relation between $\{1, x, \dots, x^d\}$ over \mathcal{A}_G . Say it is a polynomial (in x) of degree i . If it is a polynomial with the leading coefficient not a unit, then we have a zero divisor in \mathcal{A} , otherwise \mathcal{A} is a free \mathcal{A}_G -module of rank i . As G is semiregular we deduce $i = d$. Thus, the elements $1, x, x^2, \dots, x^{d-1}$ form a free basis of \mathcal{A} over \mathcal{A}_G . Let $x^d = \sum_{i=0}^{d-1} a_i x^i$ with $a_i \in \mathcal{A}_G$ and let $f(X) := X^d - \sum_{i=0}^{d-1} a_i X^i \in \mathcal{A}_G[X]$. Obviously, x is a root of $f(X)$ and as any $\sigma \in G$ fixes the coefficients of $f(X)$ we get that x^σ is also a root of $f(X)$. By a similar argument as before, we may assume that \mathcal{A} is a \mathcal{B} -module with $\{1, x, \dots, x^{m-1}\}$ as a free basis, where $m := \dim_{\mathcal{B}} \mathcal{A}$. Let $x^m = \sum_{i=0}^{m-1} b_i x^i$ with $b_i \in \mathcal{B}$, thus x is a root of the polynomial $g(X) := X^m - \sum_{i=0}^{m-1} b_i X^i \in \mathcal{B}[X]$.

Let us consider $f(X)$ as a polynomial in $\mathcal{B}[X]$. As $g(X)$ is monic we can apply the usual polynomial division algorithm to obtain polynomials $h(X)$ and $r(X)$ from $\mathcal{B}[X]$ such that the degree of $h(X)$ is $(d - m)$, the degree of $r(X)$ is less than m , and $f(X) = g(X)h(X) + r(X)$. We have $r(x) = 0$ which together with the freeness of the basis $\{1, \dots, x^{m-1}\}$ implies that $r(X) = 0$ and $f(X) = g(X)h(X)$. We know from the last paragraph that for all $\sigma \in G$, x^σ is a root of $g(X)h(X)$. If neither $g(x^\sigma)$ nor $h(x^\sigma)$ is zero, then we have a pair of zero divisors. If $g(x^\sigma) = 0$, then we can perform the division of $g(X)$ by $(X - x^\sigma)$ obtaining a polynomial $g_1(X) \in \mathcal{B}[X]$ with $g(X) = (X - x^\sigma)g_1(X)$ and can then proceed with a new automorphism $\sigma' \in G$ and with $g_1(X)$ in place of $g(X)$. In d rounds we either find a zero divisor in \mathcal{A} or two disjoint subsets K, K' of G with $g(X) = \prod_{\sigma \in K} (X - x^\sigma)$ and $h(X) = \prod_{\sigma' \in K'} (X - x^{\sigma'})$.

For $\sigma \in K$, let $\phi_\sigma : \mathcal{B}[X] \rightarrow \mathcal{A}$ be the homomorphism which fixes \mathcal{B} but sends X to x^σ . As $g(x^\sigma) = 0$, ϕ_σ induces a homomorphism from $\mathcal{B}[X]/(g(X))$ to \mathcal{A} , which we denote again by ϕ_σ . We know that ϕ_1 is actually an isomorphism $\mathcal{B}[X]/(g(X)) \cong \mathcal{A}$, therefore the maps $\mu_\sigma = \phi_\sigma \circ \phi_1^{-1}$ ($\sigma \in K$) are \mathcal{B} -endomorphisms of \mathcal{A} . Note that we can find a zero divisor in \mathcal{A} if any μ_σ is not an automorphism; also by Proposition 3.3 we can find a zero divisor in \mathcal{A} if the maps μ_σ ($\sigma \in K$) generate a nonsemiregular group of \mathcal{B} -automorphisms of \mathcal{A} . Thus, we can assume that μ_σ , for all $\sigma \in K$, generate a semiregular group of \mathcal{B} -automorphisms of \mathcal{A} . As $|K| = \dim_{\mathcal{B}} \mathcal{A}$ this means, by Lemma 3.2, that the set $\{\mu_\sigma | \sigma \in K\}$ is a group, say, H . We will now show that H is, essentially, $G_{\mathcal{B}}$ and that $\mathcal{A}_H = \mathcal{B}$.

We can as well assume that the group of k -automorphisms of \mathcal{A} generated by G and H is semiregular, for otherwise we find a zero divisor in \mathcal{A} . Again, as $|G| = \dim_k \mathcal{A}$ this means, by Lemma 3.2, that H is a subgroup of G . Thus, by Lemma 3.2, $[\mathcal{A} : \mathcal{A}_H] = |H| = |K| = [\mathcal{A} : \mathcal{B}]$ which together with the fact $\mathcal{B} \leq \mathcal{A}_H$ gives $\mathcal{A}_H = \mathcal{B}$. As $H \leq G_{\mathcal{B}}$ we also get $H = G_{\mathcal{B}}$ (if $H < G_{\mathcal{B}}$, then, by their semiregularity, $[\mathcal{A} : \mathcal{A}_H] < [\mathcal{A} : \mathcal{A}_{G_{\mathcal{B}}}] \leq [\mathcal{A} : \mathcal{B}]$, which is a contradiction). Thus, if none of the above steps yield a zero divisor, then $\mathcal{B} = \mathcal{A}_{G_{\mathcal{B}}}$. \square

Corollary 3.6 (Normal subgroup). *If $G_{\mathcal{B}}$ is a normal subgroup, then one can find a zero divisor in \mathcal{A} in deterministic polynomial time, unless \mathcal{B} is G -invariant and $G|_{\mathcal{B}} \cong G/G_{\mathcal{B}}$.*

Proof. Assume $G_{\mathcal{B}}$ to be a normal subgroup of G . We can also assume $\mathcal{B} = \mathcal{A}_{G_{\mathcal{B}}}$ as otherwise Proposition 3.5 gives a zero divisor in \mathcal{A} .

Let $g \in G$, $h \in G_{\mathcal{B}}$ and $b \in \mathcal{B}$. By the first assumption, $g^{-1}hg(b) = b$, thus $h(g(b)) = g(b)$. This means $g(b)$ is fixed by $G_{\mathcal{B}}$, or $g(b) \in \mathcal{A}_{G_{\mathcal{B}}}$, thus $g(b) \in \mathcal{B}$. As g, b are arbitrary, we deduce that \mathcal{B} is G -invariant.

Now consider the restriction map $\tau : G \rightarrow \text{Aut}_k(\mathcal{B})$ that maps g to $g|_{\mathcal{B}}$. Clearly, the kernel of τ is $G_{\mathcal{B}}$ and the image is $G|_{\mathcal{B}}$. Thus, $G/G_{\mathcal{B}} \cong G|_{\mathcal{B}}$. \square

4. KUMMER EXTENSIONS AND AUTOMORPHISMS OF AN ALGEBRA OVER A FINITE FIELD

In classical field theory a field extension L over k is called a *Kummer extension* if k has, say, a primitive r -th root of unity and $L = k(\sqrt[r]{a})$. Kummer extensions are the building blocks in field theory because they have a cyclic Galois group. In the previous section we developed a notion of semiregular groups to mimic the classical notion of Galois groups, now in this section we extend the classical notion of Kummer extensions to commutative semisimple algebras \mathcal{A} over a finite field k . The properties of Kummer extensions of \mathcal{A} , that we prove in the next three subsections, are the reason why we can get polynomial factoring-like results without invoking GRH.

4.1. Kummer-type extensions. We generalize below several tools and results in field theory, from the seminal paper of Lenstra [L91], to commutative semisimple algebras.

$k[\zeta_r]$ and Δ_r . Let k be a finite field and let r be a prime different from $\text{char } k$. By $k[\zeta_r]$ we denote the factor algebra $k[X]/(\sum_{i=0}^{r-1} X^i)$ and $\zeta_r := X \pmod{\sum_{i=0}^{r-1} X^i}$. Then $k[\zeta_r]$ is an $(r - 1)$ -dimensional k -algebra with basis $\{1, \zeta_r, \dots, \zeta_r^{r-2}\}$ and for every integer a coprime to r , there exists a unique k -automorphism ρ_a of $k[\zeta_r]$ which sends ζ_r to ζ_r^a . Let Δ_r denote the set of all ρ_a 's. Clearly, Δ_r is a group isomorphic to the multiplicative group of integers modulo r , therefore it is a cyclic group of order $(r - 1)$. Note that for $r = 2$, we have $\zeta_2 = -1$, $k[\zeta_2] = k$ and $\Delta_2 = \{id\}$.

$\mathcal{A}[\zeta_r]$ and Δ_r . Let \mathcal{A} be a commutative semisimple algebra over k ; then by $\mathcal{A}[\zeta_r]$ we denote $\mathcal{A} \otimes_k k[\zeta_r]$. We consider \mathcal{A} as embedded into $\mathcal{A}[\zeta_r]$ via the map $x \mapsto x \otimes 1$ and $k[\zeta_r]$ embedded into $\mathcal{A}[\zeta_r]$ via the map $x \mapsto 1 \otimes x$. Every element ρ_a of the group Δ_r can be extended in a unique way to an automorphism of $\mathcal{A}[\zeta_r]$ which acts as an identity on \mathcal{A} . These extended automorphisms of $\mathcal{A}[\zeta_r]$ are also denoted by ρ_a and their group by Δ_r . Note that if $\mathcal{A} = \mathcal{A}_1 \oplus \dots \oplus \mathcal{A}_t$, then $\mathcal{A}[\zeta_r] = \mathcal{A}_1[\zeta_r] \oplus \dots \oplus \mathcal{A}_t[\zeta_r]$, thus \mathcal{A} 's semisimplicity implies that $\mathcal{A}[\zeta_r]$ is semisimple as well. We can also easily see the fixed points in $\mathcal{A}[\zeta_r]$ of Δ_r just like Proposition 4.1 of [L91]:

Lemma 4.1. $\mathcal{A}[\zeta_r]_{\Delta_r} = \mathcal{A}$.

Proof. Observe that $\mathcal{A}[\zeta_r]$ is a free \mathcal{A} -module with basis $\{\zeta_r, \dots, \zeta_r^{r-1}\}$. As r is prime this basis is transitively permuted by Δ_r , thus an $x = \sum_{i=1}^{r-1} a_i \zeta_r^i \in \mathcal{A}[\zeta_r]$ is fixed by Δ_r iff all the a_i 's are equal iff $x \in \mathcal{A}$. \square

Next we consider the multiplicative group $\mathcal{A}[\zeta_r]^*$ of units in $\mathcal{A}[\zeta_r]$.

Sylow subgroup $\mathcal{A}[\zeta_r]_r^*$. Let $\mathcal{A}[\zeta_r]_r^*$ be the subgroup of the elements of $\mathcal{A}[\zeta_r]^*$ whose orders are powers of r . Note that $\mathcal{A}[\zeta_r]_r^*$ is of an r -power size and is the r -Sylow subgroup of the group $\mathcal{A}[\zeta_r]^*$. Let $|\mathcal{A}[\zeta_r]_r^*| =: r^t$.

Automorphism $\omega(a)$. Let a be coprime to r . Observe that the residue class of $a^{r^{t-1}}$ modulo r^t depends only on the residue class of a modulo r , because the map $a \mapsto a^{r^{t-1}}$ corresponds just to the projection of the multiplicative group $\mathbb{Z}_{r^t}^* \cong (\mathbb{Z}_{r-1}, +) \oplus (\mathbb{Z}_{r^{t-1}}, +)$ on the first component. Together with the fact that $x^{r^t} = 1$, for any $x \in \mathcal{A}[\zeta_r]^*$, we get that the element $x^{a^{r^{t-1}}}$ depends only on the residue class of a modulo r . This motivates the definition of the map, following [L91], $\omega(a) : x \mapsto x^{\omega(a)} := x^{a^{r^{u-1}}}$ (where $\text{ord}(x) =: r^u$) from $\mathcal{A}[\zeta_r]^*$ to itself. Note that the map $\omega(a)$ is an automorphism of the group $\mathcal{A}[\zeta_r]^*$ and it commutes with all the endomorphisms of the group $\mathcal{A}[\zeta_r]^*$. Also, the map $a \mapsto \omega(a)$ is a group embedding $\mathbb{Z}_r^* \rightarrow \text{Aut}(\mathcal{A}[\zeta_r]^*)$.

Teichmüller subgroup. Notice that if $x \in \mathcal{A}[\zeta_r]$ has order r^u , then $x^{\omega(a)} = x^{a^{r^{u-1}}}$. Thus, $\omega(a)$ can be considered as an extension of the map ρ_a that raised elements of order r to the a -th power. The elements on which the actions of $\omega(a)$ and ρ_a are the same, for all a , form the *Teichmüller subgroup*, $T_{\mathcal{A},r}$, of $\mathcal{A}[\zeta_r]^*$:

$$T_{\mathcal{A},r} := \{x \in \mathcal{A}[\zeta_r]^* \mid x^{\rho_a} = x^{\omega(a)} \text{ for every } \rho_a \in \Delta_r\}.$$

Note that $\zeta_r \in T_{\mathcal{A},r}$. For $r = 2$, $T_{\mathcal{A},2}$ is just the Sylow 2-subgroup of \mathcal{A}^* .

By [L91], Proposition 4.2, if \mathcal{A} is a field, then $T_{\mathcal{A},r}$ is cyclic. We show in the following lemma that, in our general case, given a witness of noncyclicity of $T_{\mathcal{A},r}$, we can compute a zero divisor in \mathcal{A} .

Lemma 4.2. *Given $u, v \in T_{\mathcal{A},r}$ such that the subgroup generated by u and v is not cyclic, we can find a zero divisor in \mathcal{A} in deterministic $\text{poly}(r, \log |\mathcal{A}|)$ time.*

Proof. Suppose the subgroup generated by u and v is not cyclic. Then, by Lemma 2.3 we can efficiently find a zero divisor z , in the semisimple algebra $\mathcal{A}[\zeta_r]$, of the form $z = (u^s - v^{s'})$. Next we compute the annihilator ideal I of z in $\mathcal{A}[\zeta_r]$ and its identity element e , thus $I = e\mathcal{A}[\zeta_r]$. If we can show that I is invariant under Δ_r , then Δ_r is a group of algebra automorphisms of I which of course would fix the identity element e of I . Thus, e is in $\mathcal{A}[\zeta_r]_{\Delta_r}$ and hence e is in \mathcal{A} by Lemma 4.1, so we have a zero divisor in \mathcal{A} .

Now we show that the annihilator ideal $I = e\mathcal{A}[\zeta_r]$ of z in $\mathcal{A}[\zeta_r]$ is invariant under Δ_r . By definition e is an idempotent such that $e(u^s - v^{s'}) = 0$. Observe that for any $a \in \{1, \dots, r-1\}$, we have that $(eu^s)^{\omega(a^{-1})} = (ev^{s'})^{\omega(a^{-1})}$. Using this together with the fact that $u^s, v^{s'} \in T_{\mathcal{A},r}$ we obtain $e^{\rho_a}(u^s - v^{s'}) = (e((u^s)^{\rho_a^{-1}} - (v^{s'})^{\rho_a^{-1}}))^{\rho_a} = (e((u^s)^{\omega(a^{-1})} - (v^{s'})^{\omega(a^{-1})}))^{\rho_a} = ((eu^s)^{\omega(a^{-1})} - (ev^{s'})^{\omega(a^{-1})})^{\rho_a} = 0^{\rho_a} = 0$. Thus, for all $a \in \{1, \dots, r-1\}$, $e^{\rho_a} \in I$ which means that I is invariant under Δ_r . \square

Now we are in a position to define what we call a Kummer extension of an algebra \mathcal{A} .

Kummer extension $\mathcal{A}[\zeta_r][\sqrt[s]{c}]$. For $c \in \mathcal{A}[\zeta_r]^*$ and a power s of r , by $\mathcal{A}[\zeta_r][\sqrt[s]{c}]$ we denote the factor algebra $\mathcal{A}[\zeta_r][Y]/(Y^s - c)$ and $\sqrt[s]{c} := Y \pmod{Y^s - c}$.

Remark. Given $c, c_1 \in T_{\mathcal{A},r}$ such that the order of c is greater than or equal to the order of c_1 , and c_1 is not a power of c , by Lemma 4.2, we can find a zero divisor in \mathcal{A} in $\text{poly}(r, \log |\mathcal{A}|)$ time. Therefore, the really interesting Kummer extensions are of the form $\mathcal{A}[\zeta_r][\sqrt[s]{c}]$, where $c \in T_{\mathcal{A},r}$ and ζ_r is a power of $\sqrt[s]{c}$ (as otherwise we compute a zero divisor in \mathcal{A}).

Clearly, $\mathcal{A}[\zeta_r][\sqrt[r]{c}]$ is a free $\mathcal{A}[\zeta_r]$ -module of rank s with basis $\{1, \sqrt[r]{c}, \dots, \sqrt[r]{c}^{s-1}\}$. If $c \in T_{\mathcal{A},r}$, then $\sqrt[r]{c}$ is an r -element of $\mathcal{A}[\zeta_r][\sqrt[r]{c}]^*$ and for any integer a coprime to r , we now identify an automorphism of the Kummer extension. Extending [L91], Proposition 4.3, we obtain:

Lemma 4.3. *Let $c \in T_{\mathcal{A},r}$. Then we can extend every $\rho_a \in \Delta_r$ to a unique automorphism of $\mathcal{A}[\zeta_r][\sqrt[r]{c}]$ that sends $\sqrt[r]{c}$ to $(\sqrt[r]{c})^{\omega(a)}$.*

In the rest of the paper we will use ρ_a also to refer to this extension.

Proof. For a $\rho_a \in \Delta_r$ let $\tilde{\rho}_a$ denote the map from $\mathcal{A}[\zeta_r][Y]$ to $\mathcal{A}[\zeta_r][\sqrt[r]{c}]$ that fixes \mathcal{A} , sends ζ_r to ζ_r^a and Y to $(\sqrt[r]{c})^{\omega(a)}$. (Recall, $(\sqrt[r]{c})^{\omega(a)} = (\sqrt[r]{c})^{a^{r^u-1}}$ where $\text{ord}(\sqrt[r]{c}) = r^u$.) As $c \in T_{\mathcal{A},r}$, $\tilde{\rho}_a$ maps c to $c^{\omega(a)}$ and thus maps $(Y^s - c)$ to zero. This means that $\tilde{\rho}_a$ can be seen as an endomorphism of $\mathcal{A}[\zeta_r][\sqrt[r]{c}]$ that sends $\sqrt[r]{c}$ to $(\sqrt[r]{c})^{\omega(a)}$. Clearly, $\tilde{\rho}_b \cdot \tilde{\rho}_{b'}$ is the same endomorphism as $\tilde{\rho}_{bb'}$ if b, b' are both coprime to r . Now as $\tilde{\rho}_a \cdot \tilde{\rho}_{a^{-1}} = \tilde{\rho}_1$ is the identity automorphism of $\mathcal{A}[\zeta_r][\sqrt[r]{c}]$ we get that $\tilde{\rho}_a$ is also an automorphism of $\mathcal{A}[\zeta_r][\sqrt[r]{c}]$, completing the proof. \square

We saw above automorphisms of the Kummer extension $\mathcal{A}[\zeta_r][\sqrt[r]{c}]$ that fixed \mathcal{A} . When $s = r$ we can also identify automorphisms that fix $\mathcal{A}[\zeta_r]$:

Proposition 4.4. *Let $c \in T_{\mathcal{A},r}$ and Δ_r be the group of automorphisms of $\mathcal{A}[\zeta_r][\sqrt[r]{c}]$ identified in Lemma 4.3. Then there is a unique automorphism σ of $\mathcal{A}[\zeta_r][\sqrt[r]{c}]$ such that:*

- (1) σ fixes $\mathcal{A}[\zeta_r]$ and maps $\sqrt[r]{c}$ to $\zeta_r \sqrt[r]{c}$.
- (2) σ commutes with the action of Δ_r .
- (3) σ is a semiregular automorphism of $\mathcal{A}[\zeta_r][\sqrt[r]{c}]_{\Delta_r}$ of order r and $(\mathcal{A}[\zeta_r][\sqrt[r]{c}]_{\Delta_r})_{\sigma} = \mathcal{A}$.

Proof. The map fixing $\mathcal{A}[\zeta_r]$ and mapping Y to $\zeta_r Y$ is clearly an automorphism of $\mathcal{A}[\zeta_r][Y]/(Y^r - c)$. This implies the existence and uniqueness of σ .

Let $\rho_a \in \Delta_r$ be an automorphism of $\mathcal{A}[\zeta_r][\sqrt[r]{c}]$. Clearly, the action of σ and ρ_a is commutative on any element $x \in \mathcal{A}[\zeta_r]$. Also, $(\sqrt[r]{c})^{\sigma \rho_a} = (\zeta_r \sqrt[r]{c})^{\rho_a} = (\zeta_r \sqrt[r]{c})^{\omega(a)} = \zeta_r^{\omega(a)} (\sqrt[r]{c})^{\omega(a)} = ((\sqrt[r]{c})^{\omega(a)})^{\sigma} = (\sqrt[r]{c})^{\rho_a \sigma}$. This implies the commutativity of the actions of σ and Δ_r on $\mathcal{A}[\zeta_r][\sqrt[r]{c}]$.

From commutativity it follows that $(\mathcal{A}[\zeta_r][\sqrt[r]{c}]_{\Delta_r})^{\sigma} = \mathcal{A}[\zeta_r][\sqrt[r]{c}]_{\Delta_r}$, thus σ is an automorphism of $\mathcal{A}[\zeta_r][\sqrt[r]{c}]_{\Delta_r}$. Let G be the group generated by Δ_r and σ . Then G is a commutative group of order $r(r - 1)$. As $\mathcal{A}[\zeta_r][\sqrt[r]{c}]_G = (\mathcal{A}[\zeta_r][\sqrt[r]{c}]_{\sigma})_{\Delta_r} = \mathcal{A}[\zeta_r]_{\Delta_r} = \mathcal{A}$, Lemma 3.2 implies that G is semiregular on $\mathcal{A}[\zeta_r][\sqrt[r]{c}]$. But then the subgroup Δ_r is semiregular as well and by Lemma 3.2, $\dim_k \mathcal{A}[\zeta_r][\sqrt[r]{c}]_{\Delta_r} = \dim_k \mathcal{A}[\zeta_r][\sqrt[r]{c}]/|\Delta_r| = r \dim_k \mathcal{A} = |(\sigma)| \dim_k \mathcal{A}$. This again implies that σ is a semiregular automorphism of $\mathcal{A}[\zeta_r][\sqrt[r]{c}]_{\Delta_r}$. \square

4.2. \mathcal{A} and the Kummer extension of \mathcal{A}_{τ} , where $\tau \in \text{Aut}_k(\mathcal{A})$. In this subsection we show how to express $\mathcal{A}[\zeta_r]$ as a Kummer extension of \mathcal{A}_{τ} given a semiregular $\tau \in \text{Aut}_k(\mathcal{A})$ of order r . The Lagrange resolvent technique of [Ró87] remains applicable in our context as well and leads to the following:

Lemma 4.5 (Lagrange resolvent). *Given a commutative semisimple algebra \mathcal{D} over a finite field k , a k -automorphism τ of \mathcal{D} of prime order $r \neq \text{char } k$ and a root $\xi \in \mathcal{D}_{\tau}$ of the cyclotomic polynomial $\frac{X^r - 1}{X - 1}$. We can find in deterministic $\text{poly}(r, \log |\mathcal{D}|)$ time a nonzero $x \in \mathcal{D}$ such that $x^{\tau} = \xi x$.*

Proof. Observe that if $\xi \in \mathcal{D}$ is a root of $1 + X + \dots + X^{r-1}$, then so is every power ξ^i ($i = 1, \dots, r - 1$). Take an element $y \in \mathcal{D} \setminus \mathcal{D}_\tau$ and compute the *Lagrange-resolvents* for $0 \leq j \leq r - 1$:

$$(y, \xi^j) := \sum_{i=0}^{r-1} \xi^{ij} y^{\tau^i}.$$

It is easy to see that $(y, \xi^0) = y + y^\tau + \dots + y^{\tau^{r-1}} \in \mathcal{D}_\tau$ as $\tau^r = id$, while $\sum_{j=0}^{r-1} (y, \xi^j) = ry + \sum_{i=1}^{r-1} \sum_{j=0}^{r-1} \xi^{ij} y^{\tau^i} = ry + \sum_{i=1}^{r-1} y^{\tau^i} \sum_{j=0}^{r-1} (\xi^i)^j = ry \notin \mathcal{D}_\tau$. It follows that for some $1 \leq j \leq (r-1)$, $(y, \xi^j) \notin \mathcal{D}_\tau$, let us keep this j fixed. In particular, $(y, \xi^j) \neq 0$ and taking $l := (-j)^{-1} \pmod r$ we find $x := (y, \xi^j)^l$ is also nonzero as commutative semisimple algebras do not contain nilpotent elements. This x is then the element promised in the claim as: $x^\tau = ((y, \xi^j)^\tau)^l = (\xi^{-j}(y, \xi^j))^l = \xi x$. \square

We now proceed to describe an algorithm that given a k -automorphism τ of \mathcal{A} of prime order r , expresses $\mathcal{A}[\zeta_r]$ as a Kummer extension of \mathcal{A}_τ .

Embedding $\text{Aut}_k(\mathcal{A})$ in $\text{Aut}_k(\mathcal{A}[\zeta_r])$. Given a semiregular automorphism τ of \mathcal{A} we extend τ to an automorphism of $\mathcal{A}[\zeta_r]$ by letting $\zeta_r^\tau := \zeta_r$. It is easy to see that the extension (denoted again by τ) is a semiregular automorphism of $\mathcal{A}[\zeta_r]$ as well and it commutes with Δ_r .

Application of Lemma 4.5, techniques from [L91] and a careful treatment of cases when we find zero divisors, give the following.

Proposition 4.6 (Canonical embedding of \mathcal{A}). *Given a commutative semisimple algebra \mathcal{A} over a finite field k together with a semiregular k -automorphism τ of \mathcal{A} of prime order $r \neq \text{char } k$, we can find in deterministic $\text{poly}(\log |\mathcal{A}|)$ time an element $x \in T_{\mathcal{A}, r}$ such that $x^\tau = \zeta_r x$.*

Any such x satisfies $c := x^r \in T_{\mathcal{A}, r}$ and defines an isomorphism $\phi : \mathcal{A}_\tau[\zeta_r][\sqrt[r]{c}] \cong \mathcal{A}_\tau[\zeta_r][x] = \mathcal{A}[\zeta_r]$ which fixes $\mathcal{A}_\tau[\zeta_r]$. Also, ϕ commutes with the action of Δ_r , therefore inducing an isomorphism $(\mathcal{A}_\tau[\zeta_r][\sqrt[r]{c}])_{\Delta_r} \cong \mathcal{A}$.

Proof. The idea of the proof is first to apply Lemma 4.5 to find a nonzero $x \in \mathcal{A}[\zeta_r]$ such that $x^\tau = \zeta_r x$ (note: x may not be in \mathcal{A}). Note that this x may be a zero divisor of $\mathcal{A}[\zeta_r]$, in that case we intend to decompose $\mathcal{A}[\zeta_r]$ as much as possible and apply Lemma 4.5 to each of these components. This process is repeated till it yields a $y \in \mathcal{A}[\zeta_r]^*$ such that $y^\tau = \zeta_r y$. Second, this y is used to form the x and ϕ as promised in the claim.

We maintain: a decomposition of the identity element $1 = 1_{\mathcal{A}[\zeta_r]} = 1_{\mathcal{A}}$ into orthogonal idempotents e, f that are fixed by τ ; and an element $y \in (f\mathcal{A}[\zeta_r])^*$ such that $y^\tau = \zeta_r y$ (for $f = 0$ we define $(f\mathcal{A}[\zeta_r])^*$ as (0)). Initially, we take $e = 1, f = 0, y = 0$. Since τ is semiregular its restriction to $e\mathcal{A}[\zeta_r]$ has to be nontrivial (as long as $e \neq 0$) and hence of prime order r . Therefore we can apply Lemma 4.5 with $\xi = e\zeta_r$ to find a nonzero $x \in e\mathcal{A}[\zeta_r]$ such that $x^\tau = (e\zeta_r)x = \zeta_r x$. Now compute the identity element e_1 of $x\mathcal{A}[\zeta_r]$ (which is an ideal of $e\mathcal{A}[\zeta_r]$). Note that $x\mathcal{A}[\zeta_r]$ is invariant under τ since for all $z \in \mathcal{A}[\zeta_r]$, $(xz)^\tau = x^\tau z^\tau = \zeta_r x z^\tau \in x\mathcal{A}[\zeta_r]$. This makes τ an automorphism of $x\mathcal{A}[\zeta_r]$ and so τ fixes the identity element e_1 . We could now replace e with $(e - e_1)$, f with $(f + e_1)$, y with $(x + y)$ and repeat the above steps. Note that the above one iteration decomposed $e\mathcal{A}[\zeta_r]$ into orthogonal components $(e - e_1)\mathcal{A}[\zeta_r]$ and $e_1\mathcal{A}[\zeta_r]$ and thus the procedure has to stop in at most $\dim_k \mathcal{A}[\zeta_r]$ rounds with $e = 0$.

So far we have found an element $y \in \mathcal{A}[\zeta_r]^*$ with $y^\tau = \zeta_r y$. Define $|\mathcal{A}[\zeta_r]^*| =: r^t$, $\ell := |\mathcal{A}[\zeta_r]^*|/r^t$ and $m := (-\ell)^{-1} \pmod{r}$. Since $\mathcal{A}[\zeta_r]$ is semisimple we can compute $|\mathcal{A}[\zeta_r]^*|$ by Lemma 3.4, then we can compute the highest power of r dividing this number, which gives us t . Thus, ℓ can be calculated in deterministic polynomial time. So we can compute the element $z := y^{\ell m}$. By the definition of ℓ and y , $z \in \mathcal{A}[\zeta_r]^*$ and $z^\tau = \zeta_r^{\ell m} z = \zeta_r^{-1} z$. Next compute the element $x = \prod_{b=1}^{r-1} (z^{\omega(b)})^{\rho_b^{-1}}$. Note that for all $\rho_a \in \Delta_r$,

$$x^{\rho_a} = \prod_{b=1}^{r-1} (z^{\omega(a^{-1}b)\omega(a)})^{\rho_{a^{-1}b}^{-1}} = x^{\omega(a)},$$

whence $x \in T_{\mathcal{A},r}$. Also, as τ commutes with Δ_r we have

$$x^\tau = \prod_{b=1}^{r-1} ((\zeta_r^{-1} z)^{\omega(b)})^{\rho_b^{-1}} = x \cdot \prod_{b=1}^{r-1} ((\zeta_r^{-1})^{\omega(b)})^{\rho_b^{-1}} = (\zeta_r^{-1})^{r-1} x = \zeta_r x.$$

Finally, we define c to be x^r . From the properties of x , $c \in \mathcal{A}[\zeta_r]_\tau = \mathcal{A}_\tau[\zeta_r]$ and hence $c \in T_{\mathcal{A},r}$.

Let us define the map ϕ from $\mathcal{A}_\tau[\zeta_r][\sqrt[r]{c}]$ to $\mathcal{A}[\zeta_r]$ as the one that sends $\sqrt[r]{c}$ to x and fixes $\mathcal{A}_\tau[\zeta_r]$ (formally, ϕ maps $\mathcal{A}_\tau[X, Y]/(\sum_{i=0}^{r-1} X^i, Y^r - c)$ to $\mathcal{A}[X]/(\sum_{i=0}^{r-1} X^i)$ by mapping X to X and Y to x). It is obvious from $c = x^r$ that ϕ is a homomorphism. We will show it to be an isomorphism completing the proof.

If ϕ maps an element $\sum_{i=0}^{r-1} a_i (\sqrt[r]{c})^i$ to zero, then $\sum_{i=0}^{r-1} a_i x^i = 0$. Applying τ on this j times gives $\sum_{i=0}^{r-1} a_i \zeta_r^{ij} x^i = 0$ (remember τ fixes $\mathcal{A}_\tau[\zeta_r]$ and hence the a_i 's). Summing these equations for all $0 \leq j \leq (r-1)$ we get $a_0 = 0$. As x is invertible this means that ϕ maps $\sum_{i=1}^{r-1} a_i (\sqrt[r]{c})^{i-1}$ to zero. We can now repeat the argument and deduce that the a_i 's are all zero, thus ϕ is injective. Using that $x \in T_{\mathcal{A},r}$, it is also straightforward to verify that ϕ commutes with Δ_r (viewed as automorphisms of $\mathcal{A}[\zeta_r][\sqrt[r]{c}]$). Thus it remains to show that ϕ is surjective. We will use dimension arguments. Observe that the left ring $\mathcal{A}_\tau[X, Y]/(\sum_{i=0}^{r-1} X^i, Y^r - c)$ is obviously a free module of rank $r(r-1)$ over \mathcal{A}_τ . Also, the right ring $\mathcal{A}[X]/(\sum_{i=0}^{r-1} X^i)$ is a free \mathcal{A} -module of rank $(r-1)$, hence a free \mathcal{A}_τ -module of rank $r(r-1)$ (since τ is semiregular by the hypothesis). So both the dimensions are equal, proving that ϕ is indeed surjective. \square

4.3. Zero divisors using noncyclic groups: Proof of Theorem 1.1. Let G be the group generated by the given automorphisms. We assume that \mathcal{A} is semisimple as otherwise we can compute its radical using the deterministic polynomial time algorithm of [R690, CIW96] and take an arbitrary nonzero element of the radical as a zero divisor. Notice that since G is noncyclic, the algebra \mathcal{A} is certainly not a field and zero divisors do exist. We also assume that G is semiregular, otherwise we can efficiently find a zero divisor in \mathcal{A} by Proposition 3.3. We can also assume that $|G|$ is not divisible by $\text{char } k$, otherwise $\text{char } k \leq |G| \leq \dim_k \mathcal{A}$ and Berlekamp's deterministic algorithm for polynomial factoring can be used to find all the simple components of \mathcal{A} .

As G is a small group of size $\dim_k \mathcal{A}$, we can list all its elements of prime order. The proof now proceeds by analyzing the Sylow subgroups of G and showing them all cyclic unless they yield a zero divisor of \mathcal{A} . For every prime divisor r of $|G|$ let Π_r be the set of elements of G of order r and let P_r be an r -Sylow subgroup of

G . For every $\sigma \in \Pi_r$ we can use Proposition 4.6 to compute an element $x_\sigma \in T_{\mathcal{A},r}$ with $x_\sigma^\sigma = \zeta_r x_\sigma$. Let H_r be the subgroup of $T_{\mathcal{A},r}$ generated by $\{x_\sigma | \sigma \in \Pi_r\}$.

We can assume H_r to be cyclic or else we can find a zero divisor in \mathcal{A} by Lemma 4.2. So choose an element $x \in \{x_\sigma | \sigma \in \Pi_r\}$ such that x is a generator of H_r . Now for any $\sigma \in G$, as x^σ is again in $T_{\mathcal{A},r}$, we can assume $x^\sigma \in H_r$ for otherwise we can find a zero divisor by Lemma 4.2. Thus, H_r is G -invariant and G acts as a group of automorphisms of H_r . As every element of P_r of order r moves some element in H_r , there is no nontrivial element of P_r acting trivially on H_r ; thus P_r intersects trivially with the kernel K_r of the restriction homomorphism $G \rightarrow \text{Aut}(H_r)$ (i.e. the map $\sigma \mapsto \sigma|_{H_r}$). Since H_r is cyclic, its automorphism group is Abelian. The last two observations imply that G/K_r is an Abelian group with a natural embedding of P_r into G/K_r . Thus the *normal series* $K_r \triangleleft G$ can be refined to $K_r \trianglelefteq N_r \triangleleft G$ such that $|P_r| = |G/N_r|$. With this prime r fixed, consider the intersection of “other” normal subgroups, i.e.,

$$N'_r := \bigcap_{\substack{\text{prime } q || |G| \\ q \neq r}} N_q.$$

Since N'_r is an intersection of normal subgroups, it is normal. Also, its size divides the size of each of the subgroups N_q , thus $\gcd_q\{|N_q|\}$. It means, by the definition of N_q , that $|N'_r|$ divides $|P_r|$. Note that the factor G/N'_r has a natural embedding in the direct product of the factors (G/N_q) . Thus, $|G/N'_r|$ divides $\prod_q |G/N_q| = \prod_q |P_q|$. Thus, $|G/N'_r|$ is coprime to r . This fact together with $|N'_r| \mid |P_r|$ implies that $|N'_r| = |P_r|$. Hence, N'_r is a normal r -Sylow subgroup of G . Since we have this for every r dividing $|G|$, it follows that G is a direct product of its Sylow subgroups (see Exercise 13, page 76 of [La02]). Also, as each P_r is Abelian, G is Abelian. Moreover, since the automorphism group of a cyclic group of odd prime-power order is cyclic, $\text{Aut}(H_r)$ is cyclic and finally P_r is cyclic, for every odd prime $r \mid |G|$.

It remains to show that we can find a zero divisor efficiently if the 2-Sylow subgroup P_2 of G is not cyclic. To this end we take a closer look at the subgroup H_2 constructed for the prime $r = 2$ by the method outlined above. It is generated by an element x , contains -1 , and P_2 acts faithfully as a group of automorphisms of H_2 . If $|H_2| = 2^k$, then $\text{Aut}(H_2) \cong \mathbb{Z}_{2^k}^*$. As P_2 injectively embeds in $\text{Aut}(H_2)$ and P_2 is noncyclic we get that $\mathbb{Z}_{2^k}^*$ is noncyclic, implying that $k > 2$ and structurally $\mathbb{Z}_{2^k}^*$ is the direct product of the cyclic groups generated by (-1) and (5) modulo 2^k , respectively. Now any noncyclic subgroup of such a $\mathbb{Z}_{2^k}^*$ will have the order 2 elements: (-1) and $5^{2^{k-3}} \equiv (2^{k-1} + 1)$. Thus, P_2 has the maps $\sigma_1 : x \mapsto x^{-1}$ and $\sigma_2 : x \mapsto x^{2^{k-1}+1} = -x$. Since σ_1 and σ_2 commute, \mathcal{A}_{σ_1} is σ_2 -invariant. As the group (σ_1, σ_2) is of size 4 while the group (σ_1) is only of size 2 we get by the semiregularity of G that the restriction of σ_2 to \mathcal{A}_{σ_1} is not the identity map. Hence, by Proposition 4.6 we can find an element $y \in T_{\mathcal{A}_{\sigma_1},2}$ such that $y^{\sigma_2} = -y$. We can assume that the subgroup of \mathcal{A}^* generated by x and y is cyclic as otherwise we find a zero divisor by Lemma 4.2. However, as $x \notin \mathcal{A}_{\sigma_1}$ while $y \in \mathcal{A}_{\sigma_1}$, it can be seen that (x, y) is a cyclic group only if $y \in H_2^2$ (i.e., y is square of an element in H_2). But this is a contradiction because σ_2 fixes H_2^2 . This finishes the proof of Theorem 1.1. \square

4.4. Relation with Lenstra’s result. We show how Theorem 1.1 implies Lenstra’s result on constructing isomorphisms between finite fields.

Assume that we are given two algebras \mathcal{A}_1 and \mathcal{A}_2 , both of them isomorphic to the degree r extension field of the finite field k . We consider the tensor product $\mathcal{A} = \mathcal{A}_1 \otimes_k \mathcal{A}_2$. If $|k| = q$, then the map $x \mapsto x^q$ gives a k -automorphism of order r of \mathcal{A}_1 and the bilinear map $(x, y) \mapsto (x^q, y)$ extends to an automorphism of order r of \mathcal{A} . The same holds for the map $(x, y) \mapsto (x, y^q)$. These two automorphisms generate a group G of automorphisms of \mathcal{A} isomorphic to $\mathbb{Z}_r \times \mathbb{Z}_r$. By Theorem 1.1 we can efficiently find a zero divisor $z \in \mathcal{A}$. Let $I' = z\mathcal{A}$. By Lemma 2.5 we find a nonzero ideal $I \subseteq I'$ such that $I^\sigma \cap I$ is either zero or I for every $\sigma \in G$. Put $G_0 = \{\sigma \in G \mid I^\sigma = I\}$. If G_0 is not cyclic, then we can efficiently replace I and G_0 by a smaller ideal and a smaller subgroup using Theorem 1.1 and Lemma 2.5. After at most $2 \log_2 r$ rounds we achieve the situation where G_0 is cyclic. Then $|G_0| \leq r$ because the largest cyclic subgroup of $\mathbb{Z}_r \times \mathbb{Z}_r$ has size r . We know that \mathcal{A} is the direct sum of $|G : G_0|$ copies of I , whence $\dim_k I = \dim_k \mathcal{A} / |G : G_0| = |G_0| \leq r$. Let e be the identity element of I . Then the map $\phi_1 : x \mapsto (x \otimes 1)e$ is an algebra homomorphism from \mathcal{A}_1 into I . Using the fact that \mathcal{A}_1 is a field extension of k of degree r and comparing dimensions we obtain that ϕ_1 is actually an isomorphism. Similarly, $\phi_2 : y \mapsto (1 \otimes y)e$ is an isomorphism of \mathcal{A}_2 to I and finally $\phi_2^{-1} \circ \phi_1$ is an isomorphism between \mathcal{A}_1 and \mathcal{A}_2 .

4.5. Zero divisors using Galois Groups: Proof of Theorem 1.4. If the input polynomial $f(x) \in \mathbb{Q}[x]$ has a “small” Galois group, then can we factor $f(x)$ modulo a prime p ? This question was studied in [R689b] and an algorithm was given assuming GRH. In this subsection we give a GRH-free version. We start with the following unconditional and generalized version of Theorem 3.1 in [R689b]:

Theorem 4.7. *Assume that we are given a semiregular group G of automorphisms of a commutative semisimple algebra \mathcal{A} over a finite field k with $\mathcal{A}_G = k$ and a nonzero ideal \mathcal{B} of a subalgebra \mathcal{C} of \mathcal{A} . Then in deterministic $\text{poly}(\log |\mathcal{A}|)$ time we can either find a zero divisor in \mathcal{B} or a semiregular k -automorphism σ of \mathcal{B} of order $\dim_k \mathcal{B}$.*

Remark. Here \mathcal{B} is an ideal of a subalgebra of \mathcal{A} , thus it is not assumed that $1_{\mathcal{A}} \in \mathcal{B}$.

Proof. The idea of the algorithm is to find a nontrivial ideal I of \mathcal{A} and then reduce the problem to the smaller instance I . If \mathcal{B} is a nontrivial ideal of \mathcal{C} , then $I := \mathcal{B}\mathcal{A}$ is a nontrivial ideal of \mathcal{A} , otherwise $\mathcal{B} = \mathcal{C}$. Thus, we now consider the case of \mathcal{B} being a subalgebra of \mathcal{A} .

If G is noncyclic, then using Theorem 1.1 we can find a nontrivial ideal I of \mathcal{A} . If G is cyclic, then using Proposition 3.5 we can find either a nontrivial ideal I of \mathcal{A} or a subgroup H of G with $\mathcal{B} = \mathcal{A}_H$. In the latter case $H = G_{\mathcal{B}}$ is trivially a normal subgroup of G and, by Corollary 3.6, the restriction of any generator σ of G will generate a semiregular group, of k -automorphisms of \mathcal{B} , isomorphic to G/H . Thus, we get a semiregular k -automorphism of \mathcal{B} of order $|G/H| = \dim_k \mathcal{B}$.

Let us assume we have a nontrivial ideal I of \mathcal{A} . Then, using the method of Lemma 2.5, we find an ideal J of \mathcal{A} such that the ideals $\{J^\sigma \mid \sigma \in G\}$ are pairwise orthogonal or equal. By the hypothesis $\mathcal{A}_G = k$, G acts transitively on the minimal ideals of \mathcal{A} , thus the group $G_1 := \{\sigma \in G \mid J^\sigma = J\}$ acts semiregularly on J and for coset representatives C of G/G_1 : $\mathcal{A} = \bigoplus_{\sigma \in C} J^\sigma$. Also, note that for all $\sigma \in C$ the conjugate subgroup $G_1^\sigma := \sigma^{-1}G_1\sigma$ acts semiregularly on J^σ . We can find a zero divisor in \mathcal{B} if the projection of \mathcal{B} to some J^σ is neither the zero map nor injective. Thus we assume that there is an ideal J^σ such that the projection of \mathcal{A} onto J^σ

injectively embeds \mathcal{B} . In that case we reduce our original problem to the smaller instance — J^σ instead of \mathcal{A} , G_1^σ instead of G and the embedding of \mathcal{B} instead of \mathcal{B} — and apply the steps of the preceding paragraph. \square

The following corollary gives the proof of a slightly stronger version of Theorem 1.4:

Corollary 4.8. *Let $F(X) \in \mathbb{Z}[X]$ be a polynomial irreducible over \mathbb{Q} with Galois group of size m ; let L be the maximum length of the coefficients of $F(X)$; let p be a prime not dividing the discriminant of $F(X)$; let $f(x) := F(x) \pmod{p}$; and let $g(x)$ be a nonconstant divisor of $f(x)$ in $\mathbb{F}_p[x]$. Given $F(X), p$ and $g(x)$, by a deterministic $\text{poly}(m, L, \log p)$ time algorithm we can find either a nontrivial factor of $g(x)$ or an automorphism of order $\deg g$ of the algebra $\mathbb{F}_p[x]/(g(x))$.*

Proof. The assumption on the discriminant implies that the leading coefficient of $F(X)$ is not divisible by p , and wlog we can assume $F(X)$ to be monic. Also, assume that $p > m^4$ as otherwise we can use Berlekamp’s deterministic algorithm for factoring $f(x)$ completely. Now using the algorithm of Theorem 5.3 of [R689b], we compute an algebraic integer $\alpha := X \pmod{H(X)}$ generating the splitting field $\mathbb{Q}[X]/(H(X))$ of $F(X)$ such that the discriminant of the minimal polynomial $H(X)$ of α is not divisible by p (note: $H(X)$ has integer coefficients). Define $\mathcal{A} := \mathbb{Z}[\alpha]/(p)$; then using the method described in Section 4 of [R689b], we efficiently compute a group G of automorphisms of \mathcal{A} which is isomorphic to the Galois group of α over the rationals.

Let $\beta \in \mathbb{Q}[X]/(H(X))$ be a root of $F(X)$. Then $\beta = \sum_{i=0}^{m-1} a_i \alpha^i$ for some $a_i \in \mathbb{Q}$. From Proposition 13 of Chapter 3 in [La80], for every $0 \leq i < m$, a_i can be written in the form $a_i = r_i/q_i$, where $r_i, q_i \in \mathbb{Z}$ and q_i is coprime to p . Compute $t_i \in \mathbb{Z}$ with $t_i q_i \equiv 1 \pmod{p}$. Then $\beta' := \sum_{i=0}^{m-1} r_i t_i \alpha^i$ is in $\mathbb{Z}[\alpha]$ and the minimal polynomial of the element $\bar{\beta} := \beta' \pmod{p} \in \mathcal{A}$ is $f(x)$. Let \mathcal{C} be the subalgebra $\mathbb{F}_p[\bar{\beta}]$ contained in \mathcal{A} . Notice that \mathcal{C} is isomorphic to the algebra $\mathbb{F}_p[x]/(f(x))$. Let \mathcal{B} be the ideal of \mathcal{C} generated by $f(\bar{\beta})/g(\bar{\beta})$. Then \mathcal{B} is isomorphic to the algebra $\mathbb{F}_p[x]/(g(x))$ and hence a zero divisor of \mathcal{B} will give us a factor of $g(x)$. So we run the algorithm described in Theorem 4.7 on $G, \mathcal{A}, \mathcal{B}$ and get either a factor of $g(x)$ or an automorphism of \mathcal{B} of order $\dim_{\mathbb{F}_p} \mathcal{B}$, thus finishing the proof. \square

4.6. Gluing: Extending automorphisms of \mathcal{A}_τ to \mathcal{A} , where $\tau \in \text{Aut}_k(\mathcal{A})$.

In the proof of Theorem 1.3 (Section 5.3) we need the following gluing process. It allows us to design a recursive algorithm for finding a semiregular automorphism of a given \mathcal{B} -algebra \mathcal{A} . The *gluing* refers to the process that combines the results of several recursive calls to produce a single automorphism.

Lemma 4.9 (Gluing). *Given a commutative semisimple algebra \mathcal{A} over a finite field k , a k -automorphism τ of \mathcal{A} and a k -automorphism μ of \mathcal{A}_τ . Assume that the order of τ is coprime to $\text{char } k$. Then in deterministic $\text{poly}(\log |\mathcal{A}|)$ time we can compute either a zero divisor in \mathcal{A} or a semiregular k -automorphism μ' of \mathcal{A} that extends μ such that $\mathcal{A}_{\mu'} = (\mathcal{A}_\tau)_\mu$.*

Proof. Suppose that the order of τ is $r_1 \cdots r_t$, where the r_i ’s are primes (not necessarily distinct). Clearly, it is sufficient to show how to extend μ from $\mathcal{A}_{\tau_{r_1 \cdots r_{i-1}}}$ to $\mathcal{A}_{\tau_{r_1 \cdots r_i}}$ (or find a zero divisor during the process). We can therefore assume that the order of τ is a prime r . We may also assume that both τ and μ are semiregular since otherwise we can find a zero divisor in \mathcal{A} by Proposition 3.3.

We work in the algebra $\mathcal{A}[\zeta_r]$. We extend τ to $\mathcal{A}[\zeta_r]$ and μ to $\mathcal{A}_\tau[\zeta_r]$ in the natural way. By Proposition 4.6, we can efficiently find $x \in T_{\mathcal{A},r}$ such that $x^\tau = \zeta_r x$. Clearly, $c := x^r \in T_{\mathcal{A},r}$ and $c^\mu \in T_{\mathcal{A},r}$ (μ commutes with Δ_r). The elements c and c^μ have the same order. If c^μ is not in the cyclic group generated by c , then by Lemma 4.2, we can find a zero divisor in \mathcal{A} . So assume that c^μ is in the cyclic group of c , in which case find an integer j coprime to r such that $c^\mu = c^j$ using Lemma 2.3. Note that by Lemma 4.2, we can also find a zero divisor in \mathcal{A} in the case when ζ_r is not a power of c , so assume that $\zeta_r = c^\ell$ and compute this integer ℓ . Then $\zeta_r = \zeta_r^\mu = (c^\ell)^\mu = (c^\mu)^\ell = c^{j\ell} = \zeta_r^j$, and hence $j \equiv 1 \pmod{r}$. We set $x' := x^j$. As $x^\tau = \zeta_r x$ and $x'^\tau = \zeta_r x'$, by the proof of Proposition 4.6, there are isomorphisms $\phi : \mathcal{A}_\tau[\zeta_r][\sqrt[r]{c}] \rightarrow \mathcal{A}[\zeta_r]$ and $\phi' : \mathcal{A}_\tau[\zeta_r][\sqrt[r]{c^\mu}] \rightarrow \mathcal{A}[\zeta_r]$ sending $\sqrt[r]{c}$ to x and $\sqrt[r]{c^\mu}$ to x' , respectively, and both fixing $\mathcal{A}_\tau[\zeta_r]$. We can naturally extend μ to an isomorphism $\mu'' : \mathcal{A}_\tau[\zeta_r][\sqrt[r]{c}] \rightarrow \mathcal{A}_\tau[\zeta_r][\sqrt[r]{c^\mu}]$. Then the composition map $\mu' := \phi' \circ \mu'' \circ \phi^{-1}$ is an automorphism of $\mathcal{A}[\zeta_r]$ whose restriction to $\mathcal{A}_\tau[\zeta_r]$ is μ . As μ'' , ϕ and ϕ' commute with Δ_r , so does μ' . Therefore, $\mathcal{A} = \mathcal{A}[\zeta_r]_{\Delta_r}$ is μ' -invariant and we have the promised k -automorphism of \mathcal{A} . We can assume it to be semiregular as otherwise we can compute a zero divisor in \mathcal{A} . \square

5. FINDING AUTOMORPHISMS OF ALGEBRAS VIA KUMMER EXTENSIONS

In this section we complete the proof of Theorem 1.3, i.e., given a commutative semisimple algebra \mathcal{A} over a finite field k we can unconditionally find a nontrivial k -automorphism of \mathcal{A} in deterministic quasipolynomial time. To be precise, we find a decomposition $\mathcal{A} = \bigoplus_{i=1}^t \mathcal{A}_i$ ($t \geq 1$) together with a semiregular k -automorphism of \mathcal{A}_i . This gives a k -automorphism σ of \mathcal{A} of order $\text{lcm}_i \{\dim_k \mathcal{A}_i\}$. If this lcm is 1, then $\mathcal{A}_i \cong k$, in which case we can define σ to be a nontrivial permutation of the \mathcal{A}_i 's. In this way, we can always get a nontrivial automorphism σ of \mathcal{A} . (Note: when \mathcal{A}_i is not a field, its automorphism of order $\dim_k \mathcal{A}_i$ cannot be the Frobenius.)

The proof involves computing tensor powers of \mathcal{A} , whose automorphisms we know, and then *bringing down* those automorphisms to \mathcal{A} . Before embarking on the proof we need to first see how to bring down automorphisms using Kummer extensions, and define notions related to tensor powers of \mathcal{A} .

5.1. Bringing down automorphisms of \mathcal{A} to $\mathcal{A} \leq \mathcal{D}$. We do this by using Kummer extensions, so we first show how to embed a Kummer extension of \mathcal{A} into the cyclotomic extension of \mathcal{D} .

Lemma 5.1. *Let $\mathcal{A} \leq \mathcal{D}$ be commutative semisimple algebras over a finite field k and let $r \neq \text{char } k$ be a prime. Then for any $x \in T_{\mathcal{D},r} \setminus \mathcal{A}[\zeta_r]$ satisfying $c := x^r \in \mathcal{A}[\zeta_r]$, there is a unique algebra homomorphism $\phi : \mathcal{A}[\zeta_r][\sqrt[r]{c}] \rightarrow \mathcal{D}[\zeta_r]$ that fixes $\mathcal{A}[\zeta_r]$, maps $\sqrt[r]{c}$ to x and:*

- (1) ϕ commutes with the action of Δ_r , thus $\phi(\mathcal{A}[\zeta_r][\sqrt[r]{c}]_{\Delta_r}) \subseteq \mathcal{D}$.
- (2) ϕ is injective if and only if its restriction to $\mathcal{A}[\zeta_r][\sqrt[r]{c}]_{\Delta_r}$ is injective.
- (3) If ϕ is not injective, then we can find a zero divisor of \mathcal{D} in deterministic polynomial time.

Proof. The existence and uniqueness of the homomorphism ϕ are obvious: the map from $\mathcal{A}[\zeta_r][X]$ to $\mathcal{D}[\zeta_r]$ which sends X to x factors through $\mathcal{A}[\zeta_r][\sqrt[r]{c}]$.

As $x \in T_{\mathcal{D},r}$, for every $\rho_a \in \Delta_r$ we have

$$\phi((\sqrt[r]{c})^{\rho_a}) = \phi((\sqrt[r]{c})^{\omega(a)}) = x^{\omega(a)} = (\phi(\sqrt[r]{c}))^{\rho_a}.$$

On the other hand, for every $u \in \mathcal{A}[\zeta_r]$ we have $\phi(u)^{\rho_a} = u^{\rho_a} = \phi(u^{\rho_a})$. As $\mathcal{A}[\zeta_r]$ and $(\sqrt[r]{c})$ generate $\mathcal{A}[\zeta_r][\sqrt[r]{c}]$, the two equalities above prove that ϕ commutes with the action of Δ_r . As a consequence, $\phi(\mathcal{A}[\zeta_r][\sqrt[r]{c}]_{\Delta_r}) \subseteq \mathcal{D}[\zeta_r]_{\Delta_r} = \mathcal{D}$.

Since the elements $\zeta_r^0, \dots, \zeta_r^{r-2}$ form a free basis of $\mathcal{D}[\zeta_r]$ as a \mathcal{D} -module, the subspaces $\zeta_r^i \mathcal{D}$ of $\mathcal{D}[\zeta_r]$ ($i = 0, \dots, r-2$) are independent over k . This means the images $\phi(\zeta_r^i(\mathcal{A}[\zeta_r][\sqrt[r]{c}]_{\Delta_r}))$ are independent as well, thus, $\dim_k \phi(\mathcal{A}[\zeta_r][\sqrt[r]{c}]) = (r-1) \dim_k \phi(\mathcal{A}[\zeta_r][\sqrt[r]{c}]_{\Delta_r})$. This together with the equality $\dim_k \mathcal{A}[\zeta_r][\sqrt[r]{c}] = (r-1) \dim_k \mathcal{A}[\zeta_r][\sqrt[r]{c}]_{\Delta_r}$ means that ϕ is injective if and only if its restriction to $\mathcal{A}[\zeta_r][\sqrt[r]{c}]_{\Delta_r}$ is as well.

To see the last assertion assume that ϕ , and hence its restriction to $\mathcal{C} := \mathcal{A}[\zeta_r][\sqrt[r]{c}]_{\Delta_r}$, is not injective. We compute the kernel I of $\phi|_{\mathcal{C}}$; clearly, I is a nonzero ideal of \mathcal{C} . Let σ be the semiregular k -automorphism of \mathcal{C} investigated in Proposition 4.4, which also tells us that $\dim_k \mathcal{C} = r \dim_k \mathcal{A}$. Assume that $\phi(\mathcal{C}) =: \mathcal{D}'$. We compute $J := \{u \in \mathcal{C} | uI = 0\}$, the ideal complementary to I so that $\mathcal{C} = I \oplus J$. Note that by the definition of I , the restriction of ϕ to J yields an isomorphism $J \cong \mathcal{D}'$. Hence finding a zero divisor in J implies finding a zero divisor in \mathcal{D} . Let e_J be the identity element of J , then as ϕ fixes \mathcal{A} , for all $a \in \mathcal{A}$, $a = \phi(a) = \phi(e_J a)$; in other words, ϕ induces an isomorphism $e_J \mathcal{A} \cong \mathcal{A}$. Using this we now show that the action of σ on J yields a zero divisor in J .

First, we claim that for all $1 \leq i \leq (r-1)$, we have $J \neq J^{\sigma^i}$. Suppose for some $1 \leq i \leq (r-1)$, $J^{\sigma^i} = J$ and σ^i fixes J , then $J \subseteq \mathcal{C}_{\sigma^i} = \mathcal{A}$. This together with the fact that ϕ^{-1} injectively embeds \mathcal{A} in J gives $J = \mathcal{A}$, which implies that $\phi(\mathcal{C}) = \mathcal{A}$, thus $\phi(\mathcal{A}[\zeta_r][\sqrt[r]{c}]) = \phi(\mathcal{C}[\zeta_r]) = \mathcal{A}[\zeta_r]$ contradicting $x \notin \mathcal{A}[\zeta_r]$. The other case then is: for some $1 \leq i \leq (r-1)$, $J^{\sigma^i} = J$ and the restriction of σ^i to J is a semiregular automorphism of order r of J , therefore, $\dim_k J = r \dim_k J_{\sigma^i} \geq r \dim_k e_J \mathcal{A} = r \dim_k \mathcal{A}$ (as σ^i fixes \mathcal{A} it has to fix $e_J \mathcal{A}$), which contradicts $\dim_k J < \dim_k \mathcal{C} = r \dim_k \mathcal{A}$. Second, we claim that for some $i \in \{1, \dots, r-1\}$, $J \cap J^{\sigma^i} \neq 0$. Indeed, assuming the contrary, we would have $J^{\sigma^j} \cap J^{\sigma^i} = (J \cap J^{\sigma^{i-j}})^{\sigma^j} = 0$ whenever $i \not\equiv j \pmod{r}$, whence the J^{σ^i} would be pairwise orthogonal ideals, whence $\dim_k J = \frac{1}{r} \dim_k \sum_{t=0}^{r-1} J^{\sigma^t} \leq \frac{1}{r} \dim_k \mathcal{C} = \dim_k \mathcal{A}$. This together with the fact that ϕ^{-1} injectively embeds \mathcal{A} in J gives $J = \mathcal{A}$, which implies that $\phi(\mathcal{C}) = \mathcal{A}$, thus $\phi(\mathcal{A}[\zeta_r][\sqrt[r]{c}]) = \phi(\mathcal{C}[\zeta_r]) = \mathcal{A}[\zeta_r]$ contradicting $x \notin \mathcal{A}[\zeta_r]$.

From the above two claims we get an $i \in \{1, \dots, r-1\}$, for which $J \neq J^{\sigma^i}$ and $J \cap J^{\sigma^i} \neq 0$, whence by the method of Lemma 2.5 we get a zero divisor of J , thus finishing the proof. \square

Now we show the main result of this subsection: bringing down automorphisms of \mathcal{D} to a subalgebra $\mathcal{A} \leq \mathcal{D}$.

Proposition 5.2. *Given a commutative semisimple algebra \mathcal{D} over a finite field k , a semiregular k -automorphism τ of \mathcal{D} of prime order $r \neq \text{char } k$, and a subalgebra $\mathcal{A} \supset k$ of \mathcal{D} such that $\frac{\dim_k \mathcal{D}}{\dim_k \mathcal{A}}$ is an integer not divisible by r . Then we can find in deterministic $\text{poly}(\log |\mathcal{D}|)$ time either a zero divisor in \mathcal{A} or a subalgebra $\mathcal{C} \leq \mathcal{A}$ together with a semiregular automorphism τ' of \mathcal{C} of order r such that $\mathcal{C}_{\tau'} \geq \mathcal{A}_{\tau} (:= \mathcal{A} \cap \mathcal{D}_{\tau})$.*

Proof. We use the method of Proposition 4.6 to find an element $x \in T_{\mathcal{D}, r}$ such that $x^r = \zeta_r x$. If $x \in \mathcal{A}[\zeta_r]$, then we define $\mathcal{C} := \mathcal{A}_{\tau}[\zeta_r][x]_{\Delta_r}$. As τ fixes ζ_r while Δ_r

fixes \mathcal{D} , τ commutes with Δ_r . Thus, $\mathcal{C}_\tau = (\mathcal{A}_\tau[\zeta_r][x]_\tau)_{\Delta_r} = \mathcal{A}_\tau[\zeta_r]_{\Delta_r} = \mathcal{A}_\tau$. This means that we have \mathcal{C} and $\tau' := \tau|_{\mathcal{C}}$ as promised.

Suppose now that $x \notin \mathcal{A}[\zeta_r]$. Then, since $x^{r^t} = 1_{\mathcal{D}} \in \mathcal{A}$ for some integer $t > 0$, we can choose a $y \in \{x, x^r, x^{r^2}, \dots\}$ such that $y \notin \mathcal{A}[\zeta_r]$ but $c' := y^r \in \mathcal{A}[\zeta_r]$. By Lemma 5.1, we can find a zero divisor in \mathcal{D} unless $\mathcal{A}[\zeta_r][\sqrt[r]{c'}]$ is isomorphic to the subalgebra $\mathcal{A}[\zeta_r][y]$. In the latter case $\mathcal{D}_0 := \mathcal{A}[\zeta_r][y]_{\Delta_r} \leq \mathcal{D}$ is a free \mathcal{A} -module of rank r , by Proposition 4.4. Comparing dimensions it follows that \mathcal{D} cannot be a free \mathcal{D}_0 -module, therefore we can find a zero divisor z in \mathcal{D}_0 by Lemma 2.4. Thus, whenever $x \notin \mathcal{A}[\zeta_r]$, we can find a zero divisor z in \mathcal{D} .

We proceed with computing the ideal of \mathcal{D} generated by z and using Lemma 2.5, obtain a τ -invariant decomposition of \mathcal{D} into the orthogonal ideals I_1, \dots, I_t . For $1 \leq j \leq t$, we denote by ϕ_j the projection $\mathcal{D} \rightarrow I_j$. We can assume that for all j , $\phi_j|_{\mathcal{A}}$ is injective (as otherwise we find a zero divisor in \mathcal{A}) and let $E \subseteq \{I_1, \dots, I_t\}$ be a set of representatives of all the r -sized orbits of τ . We have

$$\frac{\dim_k \mathcal{D}}{\dim_k \mathcal{A}} = \sum_{j=1}^t \frac{\dim_k I_j}{\dim_k \mathcal{A}} = \sum_{I_j^\tau = I_j} \frac{\dim_k I_j}{\dim_k \mathcal{A}} + r \sum_{I_j \in E} \frac{\dim_k I_j}{\dim_k \mathcal{A}},$$

from which we infer that the first sum is nonempty and includes at least one term not divisible by r , therefore we can choose an index j such that I_j is τ -invariant and $r \nmid \frac{\dim_k I_j}{\dim_k \mathcal{A}}$. So we can proceed with I_j and $\phi_j \mathcal{A} \cong \mathcal{A}$ in place of \mathcal{D} and \mathcal{A} , respectively, in the algorithm described above. (Note: τ remains a semiregular automorphism of I_j .)

The process described above stops when either we find a zero divisor in \mathcal{A} or an element $x \in T_{\mathcal{A}', r}$ with $x^r = \zeta_r x$, where $\mathcal{A}' \cong \mathcal{A}$ is the image of \mathcal{A} under the projection ϕ of \mathcal{D} to some τ -invariant ideal I . In the latter case we compute the subalgebra $\mathcal{C}' := \mathcal{A}'_\tau[\zeta_r][x]_{\Delta_r}$. Finally, put $\mathcal{C} := \phi^{-1}(\mathcal{C}')$ and $\tau' := \phi^{-1} \circ \tau \circ \phi$. Notice that, if e_I is the identity element of I , then τ will fix e_I and $\phi : \mathcal{D} \rightarrow I$ will just be the homomorphism $d \mapsto e_I d$, thus τ commutes with ϕ . Consequently, $\mathcal{C}_{\tau'} = \phi^{-1}(\mathcal{C}'_\tau) = \phi^{-1}(\mathcal{A}'_\tau) \geq \mathcal{A}_\tau$. We can assume τ' to be semiregular, otherwise we find a zero divisor in \mathcal{C} , hence one in \mathcal{A} . \square

5.2. Essential part of the tensor power. It was shown by Rónyai [Ró87] that, under GRH, a zero divisor in a commutative semisimple algebra \mathcal{A} over a finite field k can be found in time $\text{poly}((\dim_k \mathcal{A})^r, \log |k|)$ if r is a prime divisor of $\dim_k \mathcal{A}$. In this section we extend the method of [Ró87] and obtain a GRH-free version that will be crucial in the proof of Theorem 1.3.

A key idea in [Ró87] was to work in the *essential part* of the tensor powers of \mathcal{A} . Before going to the formal definition of it we give a motivating definition assuming $\mathcal{A} = k[X_1]/(f(X_1))$, the essential part of $\mathcal{A}^{\otimes_k 2} := \mathcal{A} \otimes_k \mathcal{A}$ is its ideal isomorphic to the algebra:

$$k[X_1, X_2]/(f(X_1), f_2(X_1, X_2)), \quad \text{where } f_2(X_1, X_2) := \frac{f(X_2)}{X_2 - X_1} \in \mathcal{A}[X_2].$$

Similarly, we can write down an expression for the essential part of $\mathcal{A}^{\otimes_k r}$ inductively, as a factor algebra of $k[X_1, \dots, X_r]$.

In a more general setting, let \mathcal{A} be a commutative semisimple algebra over a finite field k . Let \mathcal{B} be a subalgebra of \mathcal{A} such that $1_{\mathcal{A}} \in \mathcal{B}$, \mathcal{A} be a free module

over \mathcal{B} of rank $m > 1$ and let r be an integer with $1 < r \leq m$. We will concisely denote the r th tensor power of \mathcal{A} with respect to \mathcal{B} by $\mathcal{A}^{\otimes_{\mathcal{B}} r}$.

Essential part of tensor powers. For $i \neq j \in \{1, \dots, r\}$ we denote by μ_{ij} the unique \mathcal{B} -module homomorphism from $\mathcal{A}^{\otimes_{\mathcal{B}} r}$ to $\mathcal{A}^{\otimes_{\mathcal{B}}(r-1)}$ which maps

$$x_1 \otimes \dots \otimes x_i \otimes \dots \otimes x_j \otimes \dots \otimes x_r \text{ to } x_1 \otimes \dots \otimes x_{i-1} \otimes x_{i+1} \otimes \dots \otimes x_i x_j \otimes \dots \otimes x_r$$

for every $(x_1, \dots, x_r) \in \mathcal{A}^r$. The *essential part* $\widetilde{\mathcal{A}^{\otimes_{\mathcal{B}} r}}$ of $\mathcal{A}^{\otimes_{\mathcal{B}} r}$ is defined as

$$\widetilde{\mathcal{A}^{\otimes_{\mathcal{B}} r}} := \bigcap_{1 \leq i < j \leq r} \ker \mu_{ij}.$$

Obviously, the maps μ_{ij} are k -algebra homomorphisms. Therefore, $\widetilde{\mathcal{A}^{\otimes_{\mathcal{B}} r}}$ is an ideal of $\mathcal{A}^{\otimes_{\mathcal{B}} r}$ which, being the intersection of the kernels of linear maps, can be efficiently computed:

Lemma 5.3. *A basis for $\widetilde{\mathcal{A}^{\otimes_{\mathcal{B}} r}}$ over k can be computed by a deterministic algorithm in time $\text{poly}(m^r, \log |\mathcal{A}|)$.*

Example 5.4. *For $r = 2$, the essential part has a especially nice description. As we have the following exact sequence:*

$$0 \rightarrow \widetilde{\mathcal{A}^{\otimes_{\mathcal{B}} 2}} \rightarrow \mathcal{A}^{\otimes_{\mathcal{B}} 2} \xrightarrow{\mu_{12}} \mathcal{A} \rightarrow 0.$$

Dimension and automorphisms of the essential part. The symmetric group S_r acts as a group of automorphisms of $\mathcal{A}^{\otimes_{\mathcal{B}} r}$. The action of $\pi \in S_r$ is the \mathcal{B} -linear extension of the map $h_1 \otimes \dots \otimes h_r \mapsto h_{\pi(1)} \otimes \dots \otimes h_{\pi(r)}$. Unfortunately, this action is not semiregular on the tensor power algebra. However, the ideal $\widetilde{\mathcal{A}^{\otimes_{\mathcal{B}} r}}$ is obviously invariant under this action and on it S_r acts semiregularly:

Lemma 5.5. (1) $\dim_k \widetilde{\mathcal{A}^{\otimes_{\mathcal{B}} r}} = m(m-1) \dots (m-r+1) \dim_k \mathcal{B}$.

(2) *The restriction of the action of S_r defined above is semiregular on $\widetilde{\mathcal{A}^{\otimes_{\mathcal{B}} r}}$.*

Proof. Let k' be a finite extension of k which splits \mathcal{A} , i.e., $\mathcal{A}' = k' \otimes_k \mathcal{A}$ is isomorphic to a direct sum of copies of k' . Then with $\mathcal{B}' = k' \otimes_k \mathcal{B}$ we have $\mathcal{A}'^{\otimes_{\mathcal{B}' r}} = k' \otimes_k \mathcal{A}^{\otimes_{\mathcal{B}} r}$. Furthermore, for $1 \leq i < j \leq r$, the map $\mu'_{ij} : \mathcal{A}'^{\otimes_{\mathcal{B}' r}} \rightarrow \mathcal{A}'$ analogous to μ_{ij} are just the k' -linear extensions of the maps μ_{ij} , therefore, $\widetilde{\mathcal{A}'^{\otimes_{\mathcal{B}' r}}} = k' \otimes_k \widetilde{\mathcal{A}^{\otimes_{\mathcal{B}} r}}$. This implies $\dim_{k'} \widetilde{\mathcal{A}'^{\otimes_{\mathcal{B}' r}}} = \dim_k \widetilde{\mathcal{A}^{\otimes_{\mathcal{B}} r}}$. Also, if the action of S_r is not semiregular on $\widetilde{\mathcal{A}^{\otimes_{\mathcal{B}} r}}$, then neither is its extension on $\widetilde{\mathcal{A}'^{\otimes_{\mathcal{B}' r}}$. Indeed, if there is an ideal I of $\widetilde{\mathcal{A}^{\otimes_{\mathcal{B}} r}}$ on which some nonidentity permutation acts identically, so does its extension on the ideal $k' \otimes_k I$ of $\widetilde{\mathcal{A}'^{\otimes_{\mathcal{B}' r}}$. Furthermore, \mathcal{A}' is a free \mathcal{B}' -module of rank m . Thus it is sufficient to show both statements for k' , \mathcal{A}' and \mathcal{B}' in place of k , \mathcal{A} and \mathcal{B} , respectively.

Thus we may assume that \mathcal{A} is split. Then so is \mathcal{B} . Let f_1, \dots, f_s be the primitive idempotents of \mathcal{B} . Then $f_i = \sum_{j=1}^m e_{ij}$ where e_{ij} are primitive idempotents of \mathcal{A} and the idempotents $e_1 := \sum_{i=1}^s e_{i1}, \dots, e_m := \sum_{i=1}^s e_{im}$ form a free basis of \mathcal{A} over \mathcal{B} . Let $x = \sum_{i_1, \dots, i_r=1}^m x_{i_1, \dots, i_r} e_{i_1} \otimes \dots \otimes e_{i_r} \in \mathcal{A}^{\otimes_{\mathcal{B}} r}$ where $x_{i_1, \dots, i_r} \in \mathcal{B}$. Then $x \in \ker \mu_{jj'}$ if and only if $x_{i_1, \dots, i_r} = 0$ for every tuple $(i_1, \dots, i_r) \in [m]^r$ such that $i_j = i_{j'}$ and hence the elements $e_{i_1} \otimes \dots \otimes e_{i_r}$ with i_1, \dots, i_r are pairwise distinct form a free \mathcal{B} -basis for $\widetilde{\mathcal{A}^{\otimes_{\mathcal{B}} r}}$. This proves the first statement. To see the second

statement observe that the element $x = \sum x_{i_1, \dots, i_r} e_{i_1} \otimes \dots \otimes e_{i_r}$ is fixed by S_r if and only if $x_{i_1, \dots, i_r} = x_{\pi(i_1), \dots, \pi(i_r)}$ for every $\pi \in S_r$. Therefore the elements $\sum_{\pi \in S_r} e_{\pi(t_1)} \otimes \dots \otimes e_{\pi(t_r)}$ with $1 \leq t_1 < \dots < t_r \leq m$ form a free \mathcal{B} -basis of $\widetilde{\mathcal{A}^{\otimes_{\mathcal{B}} r}}_{S_r}$ whence $\dim_k \mathcal{A}^{\otimes_{\mathcal{B}} r}_{S_r} = \dim_k \widetilde{\mathcal{A}^{\otimes_{\mathcal{B}} r}}/r!$. Semiregularity now follows from Lemma 3.2. \square

Embedding \mathcal{A} in the essential part. \mathcal{A} can be embedded into $\mathcal{A}^{\otimes_{\mathcal{B}} r}$ by sending $h \in \mathcal{A}$ to $h \otimes 1_{\mathcal{A}} \otimes \dots \otimes 1_{\mathcal{A}}$. Composing this embedding with the projection onto the ideal $\widetilde{\mathcal{A}^{\otimes_{\mathcal{B}} r}}$ (which exists by the semisimplicity of the tensor power) we obtain an embedding of \mathcal{A} in $\widetilde{\mathcal{A}^{\otimes_{\mathcal{B}} r}}$. (Remark: Here the choice of $h \otimes 1_{\mathcal{A}} \otimes \dots \otimes 1_{\mathcal{A}}$ is arbitrary. There are analogous $(r - 1)$ other ways to embed \mathcal{A} into $\widetilde{\mathcal{A}^{\otimes_{\mathcal{B}} r}}$, and they would be exploited later.)

By Lemma 5.5 the ideal $\widetilde{\mathcal{A}^{\otimes_{\mathcal{B}} r}}$ has dimension $m \cdots (m - r + 1) \dim_k \mathcal{B}$ over k . Denoting the above embedded image of \mathcal{A} also by \mathcal{A} , if r is a prime divisor of m , then $m \cdots (m - r + 1)/m = \dim_k \widetilde{\mathcal{A}^{\otimes_{\mathcal{B}} r}}/\dim_k \mathcal{A}$ is not divisible by r and we can apply Proposition 5.2 with $\widetilde{\mathcal{A}^{\otimes_{\mathcal{B}} r}}$ as \mathcal{D} and the cyclic permutation $(1 \dots r)$ as τ . This immediately gives us the following GRH-free version of the result of [Ró87]:

Theorem 5.6. *Let \mathcal{B} be a subalgebra of a commutative semisimple algebra \mathcal{A} over a finite field k such that $k \subseteq \mathcal{B}$; let \mathcal{A} be a free \mathcal{B} -module of rank m ; and let r be a prime divisor of m . Then in deterministic $\text{poly}(m^r, \log |\mathcal{A}|)$ time one can either find a zero divisor in \mathcal{A} or compute a subalgebra \mathcal{C} of \mathcal{A} together with a semiregular automorphism τ of order r such that $\mathcal{C}_{\tau} \geq \mathcal{B}$.*

In the proof of Theorem 1.3 we will need one more property of the essential part of the tensor square.

Left and right mappings. Note that there are two ways to map \mathcal{A} into an ideal $I \trianglelefteq \widetilde{\mathcal{A} \otimes_{\mathcal{B}} \mathcal{A}}$: either by first embedding \mathcal{A} into $\mathcal{A} \otimes_{\mathcal{B}} \mathcal{A}$ by $h \mapsto h \otimes 1$ or by first embedding \mathcal{A} into $\mathcal{A} \otimes_{\mathcal{B}} \mathcal{A}$ by $h \mapsto 1 \otimes h$, and then projecting to the ideal I (which is also an ideal of $\mathcal{A} \otimes_{\mathcal{B}} \mathcal{A}$). The former we call the *left* mapping while the latter the *right* mapping (of \mathcal{A} into I).

We will now show that these two mappings of \mathcal{A} into $I \trianglelefteq \widetilde{\mathcal{A} \otimes_{\mathcal{B}} \mathcal{A}}$ are quite different if I is large enough.

Lemma 5.7 (Left is not right). *Let $m := \text{rk}_{\mathcal{B}} \mathcal{A}$ and I be a nonzero ideal of $\widetilde{\mathcal{A} \otimes_{\mathcal{B}} \mathcal{A}}$. Let $\tau_1 : \mathcal{A} \rightarrow I$ be the left mapping of \mathcal{A} while τ_2 be the right mapping of \mathcal{A} into I . Then there exists an element $x \in \mathcal{A}$ such that $\tau_1(x) \neq \tau_2(x)$. Furthermore, if $\dim_k I/\dim_k \mathcal{B} > m$, then $\tau_1(\mathcal{A}) \neq \tau_2(\mathcal{A})$.*

Proof. To see the first statement observe that $\widetilde{\mathcal{A} \otimes_{\mathcal{B}} \mathcal{A}}$ is the ideal of $\mathcal{A} \otimes_{\mathcal{B}} \mathcal{A}$ generated by the set $\{x \otimes 1 - 1 \otimes x | x \in \mathcal{A}\}$. (This can be immediately seen from $x \otimes y - 1 \otimes \mu_{12}(x \otimes y) = x \otimes y - 1 \otimes xy = (x \otimes 1 - 1 \otimes x)1 \otimes y$.) It follows that I (as an ideal) is generated by the elements $\{\tau_1(x) - \tau_2(x) | x \in \mathcal{A}\}$. Consequently, if $\tau_1(x) - \tau_2(x) = 0$ for all $x \in \mathcal{A}$, then $I = 0$.

To see the second assertion, note that I is an ideal of the essential part of the semisimple algebra $\mathcal{A} \otimes_{\mathcal{B}} \mathcal{A}$, which itself is an ideal of $\mathcal{A} \otimes_{\mathcal{B}} \mathcal{A}$. Thus, I is also an ideal of $\mathcal{A} \otimes_{\mathcal{B}} \mathcal{A}$, and there exists a natural projection $\phi : \mathcal{A} \otimes_{\mathcal{B}} \mathcal{A} \rightarrow I$. Then $\tau_1(\mathcal{A}) = \phi(\mathcal{A} \otimes_{\mathcal{B}} 1)$ and $\tau_2(\mathcal{A}) = \phi(1 \otimes \mathcal{A})$. From this and from the fact that

$\mathcal{A} \otimes_{\mathcal{B}} 1$ and $1 \otimes_{\mathcal{B}} \mathcal{A}$ generate $\mathcal{A} \otimes_{\mathcal{B}} \mathcal{A}$ we infer that $\tau_1(\mathcal{A})$ and $\tau_2(\mathcal{A})$ generate I . As $\dim_k \tau_i(\mathcal{A}) \leq \dim_k \mathcal{A} = m \dim_k \mathcal{B} < \dim_k I$, this excludes the possibility of $\tau_1(\mathcal{A}) = \tau_2(\mathcal{A})$. \square

5.3. Proof of Theorem 1.3. We now prove the following slightly stronger version of Theorem 1.3.

Theorem 5.8. *Given a commutative semisimple algebra \mathcal{A} over a finite field k and a subalgebra $\mathcal{B} \supseteq k$ of \mathcal{A} such that \mathcal{A} is a free \mathcal{B} -module of rank m . Then in deterministic $\text{poly}(m^{\log m}, \log |\mathcal{A}|)$ time one can either find a zero divisor in \mathcal{A} or a semiregular automorphism σ of \mathcal{A} of order m with $\mathcal{A}_\sigma = \mathcal{B}$.*

Proof. We will give an algorithm that will recurse on smaller instances. For clarity we denote the algorithm by $\mathcal{F}(\mathcal{C}, \mathcal{D})$ for given commutative semisimple finite algebras \mathcal{C} and \mathcal{D} (with k embedded). It begins by checking whether \mathcal{C} is a free \mathcal{D} -module (otherwise \mathcal{F} outputs a zero divisor in \mathcal{C}) and continues to find a semiregular \mathcal{D} -automorphism of \mathcal{C} . We now describe the behavior of $\mathcal{F}(\mathcal{A}, \mathcal{B})$.

We may assume that $\text{char } k > m^2$ as otherwise using Berlekamp's factoring algorithm we can completely decompose \mathcal{A} into simple components.

If m is even, then using the algorithm of Theorem 5.6 we either find a zero divisor in \mathcal{A} or a subalgebra $\mathcal{C} \leq \mathcal{A}$ together with a semiregular automorphism σ_0 of \mathcal{C} of order 2 with $\mathcal{C}_{\sigma_0} \geq \mathcal{B}$ in deterministic polynomial time. In the former case we are done while in the latter case we make two recursive calls: compute $\mathcal{F}(\mathcal{A}, \mathcal{C})$ and $\mathcal{F}(\mathcal{C}_{\sigma_0}, \mathcal{B})$. This way we either find a zero divisor in \mathcal{A} or we find a semiregular automorphism σ_1 of \mathcal{A} satisfying $\mathcal{A}_{\sigma_1} = \mathcal{C}$ as well as a semiregular automorphism σ_2 of \mathcal{C}_{σ_0} satisfying $(\mathcal{C}_{\sigma_0})_{\sigma_2} = \mathcal{B}$. In the former case we are done while in the latter case we apply the algorithm of Lemma 4.9 two times to construct σ from $\sigma_0, \sigma_1, \sigma_2$. This finishes the case when m is even.

Assume for the rest of the proof that m is odd. We outline here the overall flow of the algorithm. We work in the algebra $\mathcal{A}' := \widetilde{\mathcal{A} \otimes_{\mathcal{B}} \mathcal{A}}$ and $\mathcal{B}' := \phi_1(\mathcal{A})$ where, ϕ_1 and ϕ_2 are, respectively, the left and right embeddings of \mathcal{A} into \mathcal{A}' . During the course of the algorithm we maintain a nonzero ideal $I \trianglelefteq \mathcal{A}'$ with \mathcal{B}' embedded in it. Any time we find a zero divisor in I we replace I with either the ideal generated by the zero divisor or its complement, depending on which has smaller dimension. We can assume the new ideal to be a free module over an embedded \mathcal{B}' as otherwise we can find a zero divisor in \mathcal{B}' (equivalently in \mathcal{A}). Note that the rank of the new ideal over the embedded \mathcal{B}' is at most half of the original one. Initially $I = \mathcal{A}'$ and it is a free \mathcal{B}' -module of even rank $(m - 1)$ and so we can apply the recursion outlined in the second paragraph of this proof. In this way at any stage we either find a smaller ideal of I or a semiregular automorphism σ of I such that $I_\sigma = e_I \mathcal{B}' \cong \mathcal{B}'$, where e_I is the identity element of I . In the former case we replace I by the smaller ideal (with an embedded \mathcal{B}') and apply recursion which again either finds a zero divisor (and hence a smaller ideal) or a \mathcal{B}' -automorphism of the new ideal.

The recursion outlined above halts either with a zero divisor found in \mathcal{B}' (equivalently in \mathcal{A}) or with a semiregular automorphism σ of an $I \trianglelefteq \mathcal{A}'$ such that $I_\sigma = e_I \mathcal{B}' \cong \mathcal{B}'$. In the former case we are done while the latter case is what we handle now. Let $\tau_1 : \mathcal{A} \rightarrow I$ mapping a to $e_I \phi_1(a)$ be the embedding of \mathcal{A} into I . Look at the homomorphism $\tau_2 : \mathcal{A} \rightarrow I$ that maps $a \mapsto e_I \phi_2(a)$. It is a nonzero homomorphism as $\tau_2(1) = e_I \neq 0$. So we can assume τ_2 to be an embedding of \mathcal{A} in I as well or else we get a zero divisor in \mathcal{A} .

If σ is trivial, i.e., $I = e_I \mathcal{B}' \cong \mathcal{B}' \cong \mathcal{A}$, then $\mu := \tau_2^{-1} \tau_1$ is a nontrivial \mathcal{B} -automorphism of \mathcal{A} by the first part of Lemma 5.7. If μ is not semiregular, then we can find a zero divisor by Proposition 3.3 while if μ is semiregular, then we can recursively compute $\mathcal{F}(\mathcal{A}_\mu, \mathcal{B})$, find an automorphism of \mathcal{A}_μ and finally extend it to a promised automorphism of \mathcal{A} by Lemma 4.9.

So let us assume that σ is nontrivial, i.e., $I > I_\sigma = \tau_1(\mathcal{A})$, thus $\text{rk}_{\tau_1(\mathcal{B})} I > m$. We intend to apply the second part of Lemma 5.7 here. We define $\mathcal{B}'' := \tau_2(\mathcal{A})$ and recursively compute $\mathcal{F}(I, \mathcal{B}'')$. We either find a zero divisor of I or obtain a semiregular automorphism σ' of I with $I_{\sigma'} = \mathcal{B}''$. In the former case we can proceed with a smaller ideal of I or finish with a zero divisor of \mathcal{B}'' and hence of \mathcal{A} , so the latter case of having a σ' is what we think about now. We can assume that σ and σ' commute as otherwise we can find a zero divisor of I by the algorithm of Theorem 1.1 and proceed with recursion. Thus, $I_{\sigma'}$ is σ -invariant and I_σ is σ' -invariant. Thus σ and σ' can be viewed as automorphisms of $\tau_2(\mathcal{A})$ and $\tau_1(\mathcal{A})$, respectively. If both of these actions are trivial, then $\tau_1(\mathcal{A}) = I_\sigma = (I_\sigma)_{\sigma'} = (I_{\sigma'})_\sigma = I_{\sigma'} = \tau_2(\mathcal{A})$, which contradicts the second statement of Lemma 5.7. Thus one of them is nontrivial, wlog, say, σ is a nontrivial automorphism of $\tau_2(\mathcal{A})$. Then $\mu := \tau_2^{-1} \sigma \tau_2$ is a nontrivial automorphism of \mathcal{A} . Again we can either find a zero divisor of \mathcal{A} or recursively compute $\mathcal{F}(\mathcal{A}_\mu, \mathcal{B})$, getting a promised automorphism of \mathcal{A} by the algorithm of Lemma 4.9.

To see the dominating term in the time complexity observe that in any recursive call on some pair, say $\mathcal{F}(\mathcal{C}, \mathcal{D})$ with $d := \text{rk}_{\mathcal{D}} \mathcal{C}$, if d is odd then we need to recurse to the tensor square $\mathcal{C} \otimes_{\mathcal{D}} \mathcal{C}$. Thus $\text{rk}_{\mathcal{D}} \mathcal{C}$ does not increase while $\text{rk}_{\mathcal{B}} \mathcal{C}$ increases by a factor of d . As we start with rank m , we have $d \leq m$ and as the rank d is at least halved in the subsequent recursive call (if there is one), we deduce that the algorithm works at all times in an algebra of rank (over \mathcal{B}) at most $m^{\log m}$. It is then routine to verify that the algorithm requires in all just $\text{poly}(m^{\log m})$ many \mathcal{B} -operations, which proves the time complexity as promised. \square

Below is a version obtained by iterating the process of the above theorem.

Theorem 5.9 (Computing semiregular automorphisms). *Given a commutative semisimple algebra \mathcal{A} of dimension n over a finite field k . Then there is a deterministic algorithm which in quasipolynomial time $\text{poly}(n^{\log n}, \log |k|)$ computes a decomposition of \mathcal{A} into a direct sum $\mathcal{A}_1 \oplus \dots \oplus \mathcal{A}_t$ and finds an automorphism of order $\dim_k \mathcal{A}_i$ of the algebra \mathcal{A}_i , for each $1 \leq i \leq t$.*

Proof. Apply the process described in the proof of Theorem 5.8 to $\mathcal{B} = k$. If it yields a zero divisor z of \mathcal{A} , then the ideal $I := \mathcal{A}z$ and its complementary ideal I^\perp give a decomposition of $\mathcal{A} = I \oplus I^\perp$. If e_I is the identity element of I , then we can repeat the process now with \mathcal{A} replaced by $e_I \mathcal{A} = I$ and \mathcal{B} replaced by $e_I k \cong k$. Thus after several iterations based on Theorem 5.8 we get the direct sum decomposition of \mathcal{A} together with automorphisms as promised. \square

6. SPECIAL FINITE FIELDS: PROOF OF THEOREM 1.5

In this section we assume that $k = \mathbb{F}_p$ for a prime $p > 2$ and the prime factors of $(p - 1)$ are bounded by S . We also assume that all the algebras that appear in the section are completely *split* semisimple algebras over k , i.e., isomorphic to direct sums of copies of k .

We first show an algorithm that constructs an r -th Kummer extension of an algebra given a prime $r|(p-1)$. We basically generalize Lemma 2.3 of [R689a] to the following form:

Lemma 6.1. *Assume that \mathcal{A} is a free module over its subalgebra \mathcal{B} of rank d . Then in time $\text{poly}(\log |\mathcal{A}|, S)$ we can find either a zero divisor in \mathcal{A} or an element $x \in \mathcal{A}^*$ with a power of r order, for a prime $r|(p-1)$, satisfying one of the following conditions:*

- (1) $r \neq d$, $x \notin \mathcal{B}$ and $x^r \in \mathcal{B}$,
- (2) $r = d$, $x^r \notin \mathcal{B}$ and $x^{r^2} \in \mathcal{B}$,

Proof. As \mathcal{B} is a completely split semisimple algebra, say, of dimension n over k , there are orthogonal primitive idempotents f_1, \dots, f_n such that $f_i \mathcal{B} \cong k$ for all i . For an $i \in \{1, \dots, n\}$, we can project the hypothesis to the f_i component, thus $\dim_k f_i \mathcal{A} = d$ and there are orthogonal primitive idempotents $e_{i,1}, \dots, e_{i,d}$ of \mathcal{A} such that $f_i \mathcal{A} = e_{i,1} \mathcal{A} \oplus \dots \oplus e_{i,d} \mathcal{A}$. As f_i is an identity element of $f_i \mathcal{A}$ we further get that $f_i = (e_{i,1} + \dots + e_{i,d})$.

Now pick an $y \in \mathcal{A} \setminus \mathcal{B}$. Suppose (for contradiction) that for all $1 \leq i \leq n$ there is a single $y_i^* \in k$ that satisfies $ye_{i,j} = y_i^* e_{i,j}$ for all $1 \leq j \leq d$. Then their sum gives us that $y = \sum_{i=1}^n y_i^* f_i$, as each $y_i^* f_i \in \mathcal{B}$ we further get that $y \in \mathcal{B}$. This contradiction shows that there is an $i \in \{1, \dots, n\}$ and distinct $j, j' \in \{1, \dots, d\}$ such that $ye_{i,j} = y_1 e_{i,j}$ and $ye_{i,j'} = y_2 e_{i,j'}$ for some $y_1 \neq y_2 \in k$. Let us fix these i, j, j', y_1, y_2 for the rest of the proof, we do not compute them but use their existence for the correctness of the algorithm. We can assume $y \in \mathcal{A}^*$, otherwise we have a zero divisor and we are done.

Let r_1, \dots, r_t be the prime divisors of $(p-1)$. Let us assume $p \geq (S \log p + 1)$ as otherwise we can just invoke Berlekamp's polynomial factoring algorithm to find a complete split of \mathcal{A} , and we are done. As $p \geq (S \log p + 1)$, then there is an integer $0 \leq a < (S \log p + 1)$ such that $(y_1 + a)^{r_\ell} \neq (y_2 + a)^{r_\ell}$ for all $\ell \in \{1, \dots, t\}$ (since there can be at most tS elements in \mathbb{F}_p satisfying at least one of these equations). We could also assume $(y + a)$ to be invertible as otherwise we are done. Note that $(y + a)^{r_\ell} e_{i,j} = (y_1 + a)^{r_\ell} e_{i,j}$ and $(y + a)^{r_\ell} e_{i,j'} = (y_2 + a)^{r_\ell} e_{i,j'}$ which together with $(y_1 + a)^{r_\ell} \neq (y_2 + a)^{r_\ell}$ imply that $(y + a)^{r_\ell} \notin \mathcal{B}$. Thus $z := (y + a)$ is an element in \mathcal{A}^* for which $z^{r_\ell} \notin \mathcal{B}$ for $\ell \in \{1, \dots, t\}$.

Note that $z^{p-1} = 1$, in particular, $z^{p-1} \in \mathcal{B}$. Thus we can find two, not necessarily distinct, prime divisors r_1 and r_2 of $(p-1)$ such that replacing z with an appropriate power of it we have $z^{r_1}, z^{r_2} \notin \mathcal{B}$ but $z^{r_1 r_2} \in \mathcal{B}$. Either $r_1 = r_2 = d$ and we take $(x, r) = (z, d)$, or $(\text{wlog}) r_1 \neq d$ and we take $(x, r) = (z^{r_2}, r_1)$. Finally we can raise x by a suitable power (coprime to r) so that x has order power of r together with the other properties. \square

For an integer m we denote by $\Phi_m(X)$ the m -th cyclotomic polynomial in $k[X]$. Let r_1, \dots, r_t be the prime divisors of $(p-1)$. Then for a subset I of $\{1, \dots, t\}$ we denote the product $\prod_{i \in I} r_i$ by r_I . We now give an algorithm that either finds a zero divisor in \mathcal{A} or a homomorphism from an r_I -th cyclotomic extension onto \mathcal{A} .

Lemma 6.2. *Let $\mathcal{B} < \mathcal{A}$. Assume that we are also given a surjective k -algebra homomorphism from $k[X]/(\Phi_{r_I}(X))$ onto \mathcal{B} for some subset I of $\{1, \dots, t\}$. Then in time $\text{poly}(\log |\mathcal{A}|, S)$ we can compute either a zero divisor in \mathcal{A} or a subalgebra $\mathcal{B}' > \mathcal{B}$ of \mathcal{A} together with a surjective homomorphism from $k[X]/(\Phi_{r_{I'}}(X))$ onto \mathcal{B}' for some subset $I' \subseteq \{1, \dots, t\}$.*

Proof. We may clearly assume that \mathcal{A} is a free module (of rank d) over \mathcal{B} . Let the prime r and the element $x \in \mathcal{A}^*$ be the result of an application of the algorithm of Lemma 6.1. If $\mathcal{B}[x]$ is a proper subalgebra of \mathcal{A} , then we can solve the problem by two recursive calls: first on $(\mathcal{B}[x], \mathcal{B})$ and then on $(\mathcal{A}, \mathcal{B}[x])$. Thus the base case of the recursion is when $\mathcal{A} = \mathcal{B}[x]$. We handle this case now. In this case clearly $d \leq r$.

Assume case (2), i.e., $d = r$. We can assume $\mathcal{A} = \mathcal{B}[x^r]$ as otherwise the subalgebra $\mathcal{B}[x^r]$ is a proper subalgebra of \mathcal{A} and we can find a zero divisor because \mathcal{A} cannot be a free module over this subalgebra (as $\text{rk}_{\mathcal{B}} \mathcal{A} = r$ is a prime). It follows that $\Phi_r(x^r) \neq 0$ because otherwise the rank of \mathcal{A} as a \mathcal{B} -module would be at most $\phi(r) < r$, a contradiction. So we can assume $x^{r^2} \neq 1$ as otherwise $\Phi_r(x^r)|(x^{r^2} - 1)$ is a zero divisor and we are done. Thus we can find a power $\zeta \neq 1$ of x^{r^2} for which $\zeta^r = 1$. This means, in particular, that a primitive r -th root of unity is in \mathcal{B} , and we have $\mathcal{A} \cong \mathcal{B}[X]/(X^r - x^{r^2})$. So we get a \mathcal{B} -automorphism σ of \mathcal{A} that sends $x^r \mapsto \zeta x^r$. The automorphism σ is of order r , is semiregular and satisfies $\mathcal{A}_\sigma = \mathcal{B}$. We compute the element $z := \prod_{i=0}^{r-1} x^{\sigma^i}$. Then $z^\sigma = z$, therefore $z \in \mathcal{B}$. Also, $z^r = \prod_{i=0}^{r-1} (x^r)^{\sigma^i} = \zeta^{r(r-1)/2} x^{r^2}$. If r is odd, then $z^r = x^{r^2}$ while $z \neq \zeta^i x^r$ for all i ($z, \zeta^i \in \mathcal{B}$ but $x^r \notin \mathcal{B}$), thus $(z - \zeta^i x^r)$ is a zero divisor of \mathcal{A} , for some i , and we are done. If $r = 2$, then $z^2 = -x^4$. This means $4|(p-1)$ as otherwise $z^{p-1} = (-1)^{(p-1)/2} \cdot x^{2(p-1)}$, so $1 = -1$, which contradicts $p > 2$. We use the algorithm of [Sch85] for finding a square root w of -1 in k , observe that $(wz)^2 = x^4$. Again as $wz \neq \pm x^2$ ($z, w \in \mathcal{B}$ but $x^2 \notin \mathcal{B}$), thus $(wz - x^2)$ is a zero divisor of \mathcal{A} and we are done.

Assume case (1), i.e., $d < r$, with $x^r \neq 1$. We may assume $\mathcal{A} = \mathcal{B}[x]$ to be a free \mathcal{B} -module with the free basis $\{1, x, \dots, x^{d-1}\}$, as otherwise we can find a zero divisor in \mathcal{B} by Lemma 2.4. Also we can find a power $\zeta \neq 1$ of x^r for which $\zeta^r = 1$. These two facts mean that there is a well-defined endomorphism ϕ of \mathcal{A} that maps x to ζx and fixes \mathcal{B} . Compute the kernel $J \subsetneq \mathcal{A}$ of this endomorphism. If J is nonzero, then the elements of J are zero divisors of \mathcal{A} (as ϕ cannot send a unit to zero), and we are done. If J is zero, then ϕ is a \mathcal{B} -automorphism of \mathcal{A} , clearly of order r . As $\text{rk}_{\mathcal{B}} \mathcal{A} < r$, ϕ cannot be semiregular, so we get a zero divisor by Proposition 3.3 and we are done.

Finally, assume again case (1), i.e., $d < r$, with $x^r = 1$. Let ψ denote the given map $k[X]/(\Phi_{r_I}(X))$ onto \mathcal{B} . If $r \in I$, then put $y := \psi(X^{r_I/r})$. Then $y \in \mathcal{B}^* \setminus \{1\}$ because $X^{r_I/r}, (X^{r_I/r} - 1)$ are coprime to $\Phi_{r_I}(X)$ and are thus units. As $x^r = y^r$ but $x \neq x^i y$ for all i ($y \in \mathcal{B}$ while $x \notin \mathcal{B}$), we deduce that $(x - x^i y)$ is a zero divisor for some i , and we are done. Assume that $r \notin I$. Let $I' := I \cup \{r\}$ and let $\mathcal{C} = k[X]/(\Phi_{r_{I'}}(X))$. We now break \mathcal{C} using Chinese Remaindering. Let q_1 be a multiple of r which is congruent to 1 modulo r_I and let q_2 be a multiple of r_I congruent 1 modulo r . Let $X_1 := X^{q_1}$, $X_2 := X^{q_2}$ and let \mathcal{C}_1 resp. \mathcal{C}_2 be the subalgebras of \mathcal{C} generated by X_1 resp. X_2 . Then $\mathcal{C}_1 \cong k[X_1]/(\Phi_{r_I}(X_1))$ and $\mathcal{C}_2 \cong k[X_2]/(\Phi_r(X_2))$. Let ψ_1 be the given surjective map from \mathcal{C}_1 onto \mathcal{B} and let ψ_2 be the map from \mathcal{C}_2 sending X_2 to x . Let ψ' be the map from $\mathcal{C} \cong \mathcal{C}_1 \oplus \mathcal{C}_2$ into \mathcal{A} that is the linear extension of the map sending $X^i = (X_1^i, X_2^i)$ to $\psi_1(X_1^i)\psi_2(X_2^i)$. Clearly, ψ' is a homomorphism from \mathcal{C} to \mathcal{A} and is onto (as $\mathcal{A} = \mathcal{B}[x]$). This finishes the proof. \square

Using Lemma 6.2 as an induction tool, we obtain the following.

Theorem 6.3. *Let $f(X)$ be a polynomial of degree n which completely splits into linear factors over \mathbb{F}_p . Let $r_1 < \dots < r_t$ be the prime factors of $(p-1)$. Then by a deterministic algorithm of running time $\text{poly}(r_t, n, \log p)$, we can either find a non-trivial factor of $f(X)$ or compute a surjective homomorphism ψ from $\mathbb{F}_p[X]/(\Phi_{r_I}[X])$ to $\mathbb{F}_p[X]/(f(X))$, where $r_I = \prod_{i \in I} r_i$ for some subset I of $\{1, \dots, t\}$ and $\Phi_{r_I}(X)$ is the cyclotomic polynomial of degree $\prod_{i \in I} (r_i - 1)$.*

Note that if ψ is not an isomorphism, then we can break the cyclotomic ring above and find its invariant decomposition into ideals by Lemma 2.5. As we know the automorphism group of cyclotomic extension rings over \mathbb{F}_p (and of their ideals as well), this theorem immediately implies the statement of Theorem 1.5.

7. NONCOMMUTATIVE APPLICATIONS

In this section we show that given a noncommutative algebra \mathcal{A} over a finite field we can unconditionally find zero divisors of \mathcal{A} in deterministic quasipolynomial time. The idea is to compute a commutative subalgebra \mathcal{D} of \mathcal{A} , find an automorphism of \mathcal{D} using the algorithm described in Theorem 5.8, and finally construct a zero divisor of \mathcal{A} using this automorphism.

Preprocessing. Let \mathcal{A} be a finite dimensional noncommutative algebra over a finite field k . If \mathcal{A} is not semisimple, then we can compute the radical of \mathcal{A} , by the deterministic polynomial time algorithm of [R690, CIW96], and get several zero divisors. So we can assume that \mathcal{A} is semisimple. We can efficiently compute the center \mathcal{C} of \mathcal{A} (\mathcal{C} is the subalgebra having elements that commute with all elements in \mathcal{A}) by solving a system of linear equations. By the Artin-Wedderburn Theorem (see Fact 2.2) we know that if $\mathcal{C}_1, \dots, \mathcal{C}_r$ are the simple components of \mathcal{C} then, structurally, $\mathcal{A} = \bigoplus_{i=1}^r M_{m_i}(\mathcal{C}_i)$, where $M_m(R)$ stands for the algebra of all $m \times m$ matrices over the k -algebra R . Note that if the m_i 's are not all the same, then \mathcal{A} would not be a free module over \mathcal{C} and hence we can find a zero divisor in \mathcal{C} by Lemma 2.4. So we can assume $\mathcal{A} = \bigoplus_{i=1}^r M_m(\mathcal{C}_i) = M_m(\bigoplus_{i=1}^r \mathcal{C}_i) = M_m(\mathcal{C})$. Thus the hard case is to find a zero divisor in an algebra isomorphic to $M_m(\mathcal{C})$, this is what we focus on in the remaining section. We identify \mathcal{C} with the scalar matrices in $M_m(\mathcal{C})$.

7.1. Automorphisms of a commutative semisimple subalgebra of $M_m(\mathcal{C})$.

Note that for any invertible matrix A there is a natural automorphism of the full matrix algebra that maps x to $A^{-1}xA$, we call this a *conjugation* automorphism. We show in the first lemma that, under a convenient condition, an automorphism of a commutative semisimple subalgebra of the full matrix algebra corresponds to a conjugation automorphism.

Recall that every maximal commutative semisimple algebra of the full matrix algebra $M_m(F)$ over a perfect field F has dimension m over F . This follows from the fact that over an algebraic closed field (e.g., the algebraic closure of the original base field) commuting semisimple matrices can be simultaneously diagonalized. In other words, if F is algebraically closed, then every commutative semisimple subalgebra of $M_m(F)$ is in fact (up to a conjugation isomorphism) a subalgebra of the diagonal matrices. Of course, such a subalgebra will have rather special automorphisms. We characterize them in the following lemma by using standard techniques.

Lemma 7.1 (Skolem-Noether). *Let \mathcal{C} be a commutative semisimple algebra over a perfect field F , let $\mathcal{B} \leq M_m(\mathcal{C})$ be a commutative semisimple \mathcal{C} -algebra and let σ be a \mathcal{C} -automorphism of \mathcal{B} . Let there be a maximal commutative semisimple subalgebra $\mathcal{D} \leq M_m(\mathcal{C})$ containing \mathcal{B} such that \mathcal{D} is a free \mathcal{B} -module. Then there exists an invertible $y \in M_m(\mathcal{C})$ such that $\forall x \in \mathcal{B}, x^\sigma = y^{-1}xy$.*

Proof. We get hold of this element y by reducing the question to the case of \mathcal{C} being an algebraically closed field, when \mathcal{D} becomes a direct sum of m copies of \mathcal{C} , and \mathcal{B} becomes a direct sum of $r|m$ copies of \mathcal{C} . In that case we can find a basis of 0-1 diagonal matrices for \mathcal{B} that is permuted by σ and hence construct the desired y as a permutation matrix.

First, we can assume \mathcal{C} to be a field because if I_1, \dots, I_c are the simple components of \mathcal{C} , then clearly the I_i 's are all perfect fields, and we can try finding the promised y_i for the instance of $(\mathcal{D}I_i, \mathcal{B}I_i, I_i)$. Note that since σ was fixing I_i , σ is still a (I_i) -automorphism of $\mathcal{B}I_i$ and by the freeness condition, $\mathcal{D}I_i$ is still a free $(\mathcal{B}I_i)$ -module and it is a maximal commutative semisimple subalgebra of $M_m(I_i)$. Also, once we have the y_i , for all $1 \leq i \leq c$, satisfying $y_i x^\sigma = x y_i$ for all $x \in I_i$; it is easy to see that $(y_1 + \dots + y_r)$ is a suitable y . So for the rest of the proof we assume that \mathcal{C} is a finite field extension of F .

Second, notice that the condition $yx^\sigma = xy$ is equivalent to the system of equations: $yx_1^\sigma = x_1y, \dots, yx_r^\sigma = x_r y$ for a \mathcal{C} -basis x_1, \dots, x_r of \mathcal{B} . In terms of the entries of the matrix y this is a system of homogeneous linear equations in the field \mathcal{C} . This system has a nonzero solution over \mathcal{C} iff the same system has a nonzero solution over the algebraic closure $\bar{\mathcal{C}}$ of \mathcal{C} . A solution over $\bar{\mathcal{C}}$ gives a matrix $y \in M_m(\bar{\mathcal{C}})$ such that $yx^\sigma = xy$ for every $x \in \bar{\mathcal{B}}$ where $\bar{\mathcal{B}} := \bar{\mathcal{C}} \otimes_{\mathcal{C}} \mathcal{B}$ and we extend σ $\bar{\mathcal{C}}$ -linearly to an algebra automorphism of $\bar{\mathcal{B}}$. Because F was a perfect field, $\bar{\mathcal{B}} \leq M_m(\bar{\mathcal{C}})$ is a commutative semisimple algebra over $\bar{\mathcal{C}}$. Similarly, $\bar{\mathcal{D}} := \bar{\mathcal{C}} \otimes_{\mathcal{C}} \mathcal{D}$ is a maximal commutative semisimple subalgebra of $M_m(\bar{\mathcal{C}})$, and is also a free $\bar{\mathcal{B}}$ -module. By the former condition $\text{rk}_{\bar{\mathcal{C}}} \bar{\mathcal{D}} = m$ and by the latter condition $r|m$. We will now focus on the instance of $(\bar{\mathcal{D}}, \bar{\mathcal{B}}, \bar{\mathcal{C}})$ and try to construct the desired y .

As $\bar{\mathcal{D}}$ is a sum of m copies of $\bar{\mathcal{C}}$, by an appropriate basis change we can make $\bar{\mathcal{D}}$ the algebra of all diagonal matrices in $M_m(\bar{\mathcal{C}})$. Also, as $\bar{\mathcal{D}}$ is a free $\bar{\mathcal{B}}$ -module, a further basis change makes $\bar{\mathcal{B}}$ the algebra generated by the matrices e_1, \dots, e_r where each e_j is a diagonal 0-1 matrix having m/r consecutive 1-s (note: e_i 's are primitive idempotents). In that case the automorphism σ has a simple action, namely it permutes the matrices $\{e_1, \dots, e_r\}$. Let y be a block $r \times r$ -matrix whose blocks are all $m/r \times m/r$ zero matrices except at positions i, i^σ (i^σ is defined by $e_i^\sigma = e_{i^\sigma}$), where the block is the $m/r \times m/r$ identity matrix. Clearly then, $e_{i^\sigma} = y^{-1}e_i y$ for all $1 \leq i \leq r$ and hence $x^\sigma = y^{-1}xy$ for every $x \in \bar{\mathcal{B}}$ by extending the equalities linearly to $\bar{\mathcal{B}}$. □

In the next lemma we show that a conjugation automorphism of prime order of a commutative semisimple subalgebra corresponds to a zero divisor of the original algebra.

Lemma 7.2. *Let \mathcal{A} be a finite dimensional algebra over the perfect field F and let $\mathcal{B} \leq \mathcal{A}$ be a commutative semisimple algebra containing $F1_{\mathcal{A}}$. Let r be a prime different from $\text{char } F$ and let $y \in \mathcal{A}$ be of order r such that $y^{-1}\mathcal{B}y = \mathcal{B}$, but there is an element $x \in \mathcal{B}$ with $y^{-1}xy \neq x$. Then the minimal polynomial of y over F is*

in fact $(X^r - 1)$. As a consequence, $(y - 1)$ and $(1 + y + \dots + y^{r-1})$ is a pair of zero divisors in \mathcal{A} .

Proof. Let \overline{F} be the algebraic closure of F . Note that in $\overline{\mathcal{A}} := \overline{F} \otimes_F \mathcal{A}$, the minimal polynomial of $1 \otimes y$ is the same as that of y in \mathcal{A} , $\overline{\mathcal{B}} := \overline{F} \otimes \mathcal{B}$ remains commutative semisimple and conjugation by $1 \otimes y$ acts on it as an automorphism of order r . Thus for the rest of the proof we can assume F to be algebraically closed.

As conjugation by y does not fix \mathcal{B} , there exists a primitive idempotent e of \mathcal{B} for which the elements $e_j = y^{-j} e y^j$ ($j = 0, \dots, r-1$) are pairwise orthogonal primitive idempotents of \mathcal{B} . This means that the corresponding left ideals $L_j := \mathcal{A} e_j$ are linearly independent over F . Assume now that the minimal polynomial of y has degree less than r . So there are elements $\alpha_0, \dots, \alpha_{r-1} \in F$, not all zero, such that $\sum_{j=0}^{r-1} \alpha_j y^j = 0$. Implying that $e \sum_{j=0}^{r-1} \alpha_j y^j = \sum_{j=0}^{r-1} \alpha_j y^j e_j = 0$, this together with the fact that $y^j e_j$'s are all nonzero, contradicts the linear independence of L_1, \dots, L_r . \square

7.2. Proof of Theorem 1.6. In this subsection we give a proof of Theorem 1.6: given a noncommutative algebra \mathcal{A} over a finite field k , a commutative subalgebra \mathcal{B} of \mathcal{A} and an automorphism σ of \mathcal{B} such that $\mathcal{B}_\sigma \geq \mathcal{B} \cap \mathcal{C}$, we find a zero divisor in \mathcal{A} in deterministic polynomial time, where \mathcal{C} is the center of \mathcal{A} . We wish to apply Lemma 7.1 on σ to get a y and then apply Lemma 7.2 to get a zero divisor. But we cannot do this directly as the y may not be of a prime order, so we have to work more. The basic idea in the algorithm is to define a subalgebra of \mathcal{A} which is a so-called *cyclic algebra*, and then find a zero divisor in this cyclic algebra by the method of [W05]. The *cyclic algebras* \mathcal{A}' we encounter have two generators x, y over a certain commutative semisimple algebra \mathcal{C}' such that for a prime r : $xy = \zeta_r yx$ and the multiplicative orders of x, y are powers of r . These algebras have the *ring of quaternions* as their classic special case, when $x^2 = y^2 = -1$ and $xy = -yx$.

By the *preprocessing* discussed in the beginning of the section it is clear that we need to only handle the case of $\mathcal{A} \cong M_m(\mathcal{C})$, where \mathcal{C} is a commutative semisimple algebra over k . We compute the unique linear extension σ' of σ to $\mathcal{B}\mathcal{C}$ which acts as the identity of \mathcal{C} . Then σ' is an automorphism of the commutative subalgebra $\mathcal{B}\mathcal{C}$ of \mathcal{A} . By replacing \mathcal{B} with $\mathcal{B}\mathcal{C}$ and σ with σ' in this way we achieve that $\mathcal{B} \geq \mathcal{C}$.

If \mathcal{B} is not semisimple, then we find a zero divisor by computing its radical. Therefore we can assume that \mathcal{B} is semisimple. We can compute a maximal commutative semisimple subalgebra \mathcal{D} of \mathcal{A} containing \mathcal{B} by the deterministic polynomial time algorithm of [GI00]. We can find a zero divisor in \mathcal{B} if \mathcal{D} is not a free module over \mathcal{B} . Otherwise by Lemma 7.1, there certainly exists a $y \in \mathcal{A}$ such that $b^\sigma = y^{-1} b y$ for every $b \in \mathcal{B}$, so by picking a nonzero solution of system of linear equations corresponding to $y b^\sigma = b y$ (where b runs over a set of algebra generators for \mathcal{B}) we either find a zero divisor of \mathcal{A} or we find a y with $b^\sigma = y^{-1} b y$ for every $b \in \mathcal{B}$. Suppose the second case. We may assume that σ is semiregular since otherwise we can find a zero divisor. Therefore the order of σ is a small number which we can compute.

We can efficiently obtain a multiple M of the multiplicative order of y , $\text{ord}(y)$, just by looking at the degrees of the irreducible factors of the minimal polynomial of y over k (this can be done deterministically without actually computing the factorization). Fix a prime divisor r of the order of σ , and, using M , replace y and σ by an appropriate power such that $\text{ord}(y)$ is a power of r while $\text{ord}(\sigma) = r$. By

this construction, conjugation by y is now a \mathcal{C} -automorphism σ of \mathcal{B} of order r . Put $z := y^r$, thus $b = b^{\sigma} = z^{-1}bz$ for every $b \in \mathcal{B}$. Note that we can assume $z \neq 1$ as otherwise $(y - 1)$ is a zero divisor of \mathcal{A} by Lemma 7.2. Thus an appropriate power, say, ζ_r , of z has order r . Consider the subalgebra $\mathcal{B}[z]$, it is commutative by the action of z on \mathcal{B} as seen before; it can also be assumed to be semisimple as otherwise we can find many zero divisors by just computing its radical. So $\mathcal{B}[z]$ is a commutative semisimple algebra and we can replace \mathcal{B} with $\mathcal{B}[z]$ and σ with the conjugation action of y on $\mathcal{B}[z]$. In this way we achieved the situation where $z \in \mathcal{B}$ and hence an appropriate power ζ_r of z is also in \mathcal{B} . So by Lemma 4.5 we can find efficiently either a zero divisor in \mathcal{B} or an $x \in \mathcal{B}^*$ such that $x^{\sigma} = \zeta_r x$. We assume the latter case and we replace x by an appropriate power so that $\text{ord}(x)$ is an r -power. (Remark: This replaces ζ_r by some power which is still a primitive r -th root of unity, which we now call ζ_r .) Let $w := x^r$. Then $w \in \mathcal{B}_{\sigma}$.

Let $\mathcal{C}' := \mathcal{B}_{\sigma}$, $\mathcal{A}' := \mathcal{C}'[x, y]$, $\mathcal{B}_x := \mathcal{C}'[x] \leq \mathcal{A}'$, $\mathcal{B}_y := \mathcal{C}'[y] \leq \mathcal{A}'$. Note that by the definitions of w, z it is easy to deduce that \mathcal{C}' is in the center of \mathcal{A}' and $x, y \notin \mathcal{C}'$. Furthermore, by $xy = \zeta_r yx$ it follows that the set $\{x^i y^j \mid 1 \leq i, j \leq (r - 1)\}$ is a system of generators for \mathcal{A}' as a \mathcal{C}' -module. The relation $xy = \zeta_r yx$ also implies, that conjugation by y acts on \mathcal{B}_x as an automorphism of order r and that the conjugation by x acts on \mathcal{B}_y as an automorphism of order r . We can assume that both of these \mathcal{C}' -automorphisms are semiregular as otherwise we can find a zero divisor by Proposition 3.3. Thus both \mathcal{B}_x and \mathcal{B}_y are free modules over \mathcal{C}' of rank r , furthermore we may assume \mathcal{A}' to be a free \mathcal{C}' -module or else we find a zero divisor in \mathcal{C}' by Lemma 2.4.

We can assume that w, z generate a cyclic subgroup of \mathcal{C}' , otherwise by Lemma 2.3 we can find a zero divisor in \mathcal{C}' . If the order of z is larger than the order of w , then there is a $u \in \mathcal{C}'$ with $u^r = w$. Put $x' := u^{-1}x$, then $x'^r = 1$ and $x'y = \zeta_r yx'$, thus conjugation by x' gives an automorphism of \mathcal{B}_y (of order r), whence $(x' - 1)$ is a zero divisor by Lemma 7.2. Similarly, we find a zero divisor if the order of w is larger than the order of z . Thus we can assume that w and z have equal orders, say r^t . By looking at the elements $w^{r^{t-1}}$ and $z^{r^{t-1}}$, both of which have order r and they generate a cyclic group, we can find a unique $0 < j < r$ such that $\text{ord}(w^j z) < r^t$. We now follow the method of the proof of Theorem 5.1 of [W05] to find a zero divisor in \mathcal{A}' .

Define $y' := x^j y$, and using $(yx = \zeta_r^{-1}xy)$ repeatedly we get,

$$\begin{aligned} y'^r &= (x^j y)^{r-2} (x^j y) (x^j y) = (x^j y)^{r-3} (x^j y) (\zeta_r^{-j} x^{2j} y^2) \\ &= \zeta_r^{-jr(r-1)/2} x^{rj} y^r = \zeta_r^{-jr(r-1)/2} w^j z. \end{aligned}$$

Thus, if r is odd, then $y'^r = w^j z$, and replacing y with y' leads to the case discussed above where the order of the new z (i.e. $w^j z$) is less than that of w (remember that $xy' = \zeta_r y'x$ still holds), and we already get a zero divisor. If $r = 2$, then $y'^2 = -wz$ ($j = 1$), and the argument of the odd case can be repeated except when $\text{ord}(-wz)$ does not fall, i.e., orders are such that $\text{ord}(wz) < \text{ord}(w) = \text{ord}(z) = \text{ord}(-wz)$. This case is only possible (recall $z \neq 1$) when $w = z = -1$, so $x^2 = y^2 = -1$ and $y^{-1}xy = -x$. Notice that in this case \mathcal{A}' is like a ring of quaternions and we handle this case next in a standard way.

To treat this case, by Theorem 6.1 of [W05], one can efficiently find $\alpha, \beta \in k$ such that $\alpha^2 + \beta^2 = -1$. Put $u := (\alpha y + \beta) \in \mathcal{B}_y$ and $x' := ux$. (Remark: Since $u(\alpha y - \beta) = 1$, u is invertible.) If $x' \in \mathcal{B}_y$, then $x \in u^{-1}\mathcal{B}_y = \mathcal{B}_y$, which is a

contradiction. Thus, $x' \notin \mathcal{B}_y$, in particular, $x' \neq \pm 1$. While using $xy = -yx$ we can deduce that $x'^2 = (\alpha y + \beta)x(\alpha y + \beta)x = (\alpha y + \beta)(-\alpha y + \beta)x^2 = (\alpha^2 + \beta^2)(-1) = 1$. Thus $(x' - 1)$ is a zero divisor. This finishes the proof of Theorem 1.6 in all cases. \square

7.3. Proof of Theorems 1.7 and 1.8. We begin this subsection with the proof of Theorem 1.7: given a noncommutative algebra \mathcal{A} over a finite field k , one can unconditionally find zero divisors of \mathcal{A} in deterministic quasipolynomial time. By the *preprocessing* discussed in the beginning of the section it is clear that we need only handle the case of $\mathcal{A} \cong M_m(\mathcal{C})$, where \mathcal{C} is a commutative semisimple algebra over k . We can compute easily the center of \mathcal{A} , and it will be \mathcal{C} . We can also compute a maximal commutative semisimple subalgebra \mathcal{D} of \mathcal{A} by the deterministic polynomial time algorithm of [GI00] (\mathcal{D} has an unknown isomorphism to the subalgebra of diagonal matrices of $M_m(\mathcal{C})$). Being maximal, \mathcal{D} is a free module over \mathcal{C} of rank m . By Theorem 5.8 we can, in deterministic $\text{poly}(m^{\log m}, \log |\mathcal{A}|)$ time, either find a zero divisor in \mathcal{D} or compute a semiregular automorphism σ of \mathcal{D} such that $\mathcal{D}_\sigma = \mathcal{C}$. In the former case we are done, while in the latter case we can apply Theorem 1.6 to finish the proof of Theorem 1.7.

In the algorithm outlined above we can use Theorem 5.6 instead of Theorem 5.8 in cases when it gives a faster algorithm for finding an automorphism of a subalgebra. We obtain the following.

Corollary 7.3. *Given an algebra \mathcal{A} which is isomorphic to $M_m(\mathcal{C})$ for a commutative semisimple algebra \mathcal{C} over the finite field k , one can find a zero divisor in time $\text{poly}(m^{\min(r, \log m)}, \dim_k \mathcal{A}, \log |k|)$, where r is the smallest prime factor of m .*

Now we move to Theorem 1.8. Given an algebra \mathcal{A} isomorphic to $M_m(K)$, we can explicitly construct the isomorphism in deterministic quasipolynomial time. We prove it by iterative applications of Corollary 7.3. We first find an element $z \in \mathcal{A}$ which, considered as an m by m matrix, has rank one. We use Corollary 7.3 to find a zero divisor z in \mathcal{A} . Using the methods of [BR90, CIK97], from z one can construct another zero divisor z' such that, if the rank of z as an m by m matrix is d , then the rank of z' is $d' = \gcd(d, m)$. The algorithm is as follows: let ℓ be a positive integer such that $\ell d \equiv \gcd(d, m) \pmod{m}$. Form the direct sum M of ℓ copies of a left \mathcal{A} -module isomorphic to the left ideal $\mathcal{A}z$ and apply FINDFREE from [BR90] to find a maximal dimensional free submodule F . FINDFREE is proved to work over characteristic zero fields in [BR90], while over finite fields we can use Theorem 10 of [CIK97]. The factor \mathcal{A} -module M/F is isomorphic to a left ideal which is the sum of $\gcd(d, m)$ minimal left ideals. By Theorem 9 of [CIK97] we can find a generator of the module M/F . The annihilator of this generator will be a left ideal of rank $(m - \gcd(d, m))$, a right identity element e of the annihilator will be an idempotent of rank $(m - \gcd(d, m))$, and $(1 - e) =: z'$ will be an idempotent of rank $\gcd(d, m)$. We can compute e by a solution of a system of inhomogeneous linear equations.

We consider the left ideal $L = \mathcal{A}z'$ and the right ideal $R = z'\mathcal{A}$ generated by z' . Then $L \cap R$ is a subalgebra of \mathcal{A} isomorphic to $M_{d'}(k)$. If $d' > 1$, we recurse into $L \cap R$. This recursion will give us a zero divisor z of rank one in time $\text{poly}(m^{\min(r, \log m)}, \dim_k \mathcal{A}, \log |k|)$. Then the left ideal $L = \mathcal{A}z$ is a vector space of dimension m and multiplication of the element of L by elements of \mathcal{A} gives a representation of \mathcal{A} as m by m matrices, finishing the proof of Theorem 1.8.

ACKNOWLEDGEMENTS

The authors thank Hausdorff Research Institute for Mathematics, Bonn for its kind support and hospitality. The research of the first and third authors was also supported by the Hungarian Research Fund (OTKA), grants 72845 and 77476. We are grateful to the anonymous referee for giving a very detailed review and a host of useful suggestions that greatly improved the writeup.

REFERENCES

- [BR90] L. Babai, L. Rónyai, Computing irreducible representations of finite groups, *Proc. 30th IEEE FOCS (1989)* pp. 93-98; journal version appeared in *Mathematics of Computation* 55, 192 (1990), 705-722. MR1035925 (91d:68057)
- [BGL01] E. Bach, J. von zur Gathen, H. W. Lenstra, Jr., Factoring polynomials over special finite fields; *Finite Fields and Their Applications* 7(2001), 5-28. MR1803933 (2001k:11252)
- [Be67] E. R. Berlekamp, Factoring polynomials over finite fields, *Bell System Technical Journal* 46(1967), 1853-1859. MR0219231 (36:2314)
- [Cam83] P. Camion, A deterministic algorithm for factorizing polynomials of $\mathbb{F}_q[x]$, *Ann. Discr. Math.*, 17, (1983), 149-157. MR841289 (87i:11177)
- [CH00] Q. Cheng, M. A. Huang, Factoring Polynomials over Finite Fields and Stable Colorings of Tournaments, *Algorithmic Number Theory Symposium (ANTS) IV, LNCS 1838, (2000)*, 233-245. MR1850608 (2002g:11178)
- [CIK97] A. Chistov, G. Ivanyos, M. Karpinski, Polynomial time algorithms for modules over finite dimensional algebras, *Proc. ISSAC 1997*, 68-74. MR1809971
- [CIW96] A. M. Cohen, G. Ivanyos, D. B. Wales, Finding the radical of an algebra of linear transformations, *Journal of Pure and Applied Algebra*, 117-118 (1997), 177-193. (*Proc. MEGA '96.*) MR1457838 (98h:16026)
- [CZ81] D. G. Cantor, H. Zassenhaus, A new algorithm for factoring polynomials over finite fields, *Mathematics of Computation*, 36(154), 1981, 587-592. MR606517 (82e:12020)
- [Ev89] S. A. Evdokimov, Factorization of a solvable polynomial over finite fields and the generalized Riemann Hypothesis, *Zapiski Nauchnykh Seminarov LOMI*, 176(1989), 104-117. MR1023599 (91a:11063)
- [Ev94] S. Evdokimov, Factorization of polynomials over finite fields in subexponential time under GRH, *Proc. 1st ANTS, Lecture Notes In Computer Science 877, Springer-Verlag 1994*. MR1322724 (95m:11145)
- [FR85] K. Friedl, L. Rónyai, Polynomial time solutions of some problems of computational algebra; *Proc. 17th ACM STOC (1985)*, pp. 153-162.
- [Gao01] S. Gao, On the deterministic complexity of factoring polynomials, *J. of Symbolic Computation*, 31(1-2), 2001, 19-36. MR1806204 (2002c:11175)
- [G87] J. von zur Gathen, Factoring polynomials and primitive elements for special primes, *Theoretical Computer Science*, 52, 1987, 77-89. MR918114 (89a:11126)
- [GHPS06] W. A. de Graaf, M. Harrison, J. Pilnikova, J. Schicho, A Lie algebra method for rational parametrization of Severi-Brauer surfaces, *J. Algebra* 303, 2006, 514-529. MR2255120 (2007e:14058)
- [GI00] W. A. de Graaf, G. Ivanyos, Finding splitting elements and maximal tori in matrix algebras, In: *F. Van Oystaeyen, M. Saorin (eds), Interactions between Ring Theory and Representations of Algebras, (Proc. Euroconference in Murcia, 1998), Lecture Notes in Pure and Applied Mathematics 210, Marcel Dekker 2000, 95-105*. MR1758404 (2001f:16089)
- [GS92] J. von zur Gathen, V. Shoup, Computing Frobenius maps and factoring polynomials, *Comput. Complexity*, 2(1992), 187-224. MR1220071 (94d:12011)
- [Hua85] M. A. Huang, Riemann hypothesis and finding roots over finite fields, *Proc. 17th ACM STOC (1985)* pp. 121-130; journal version appeared in *J. Algorithms*, 12 (1991), 464-481. MR1114921 (92j:68057)

- [Hu86] D. Husemöller, Elliptic curves; *Springer, 1986*. MR868861 (88h:11039)
- [IKS08] G. Ivanyos, M. Karpinski, N. Saxena, Schemes for Deterministic Polynomial Factoring, *Proc. 34th ISSAC (2009)*, 191-198.
- [KI03] V. Kabanets, R. Impagliazzo, Derandomizing polynomial identity tests means proving circuit lower bounds, *Proc. 35th ACM STOC (2003)*, 355-364; journal version appeared in *Computational Complexity*, 13 (2004), 1-46. MR2105971 (2005i:68025)
- [KS98] E. Kaltofen, V. Shoup, Subquadratic-time factoring of polynomials over finite fields, *Math. Comp.*, 67 (1998), 1179-1197. MR1459389 (99m:68097)
- [KS05] N. Kayal, N. Saxena, On the Ring Isomorphism and Automorphism Problems, *Proc. 20th IEEE Conference on Computational Complexity (2005)* pp. 2-12; journal version appeared in *Computational Complexity* 15 (4), (2006), 342-390. MR2324421 (2009g:68074)
- [KU08] K. Kedlaya, C. Umans, Fast modular composition in any characteristic, *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (2008)* pp. 146-155.
- [La80] S. Lang, Algebraic number theory, *Springer-Verlag, 1980*. MR1282723 (95f:11085)
- [La02] S. Lang, Algebra, *Springer-Verlag, 2002*. MR1878556 (2003e:00003)
- [L91] H. W. Lenstra, Finding isomorphisms between finite fields, *Mathematics of Computation* 56 (1991), 329-347. MR1052099 (91d:11151)
- [MS88] M. Mignotte, C.-P. Schnorr, Calcul déterministe des racines d'un polynôme dans un corps fini, *Comptes Rendus Académie des Sciences (Paris)*, 306 (1988), 467-472. MR939433 (89i:11134)
- [Moe77] R. T. Moenck, On the efficiency of algorithms for polynomial factoring, *Math. Comp.*, 31 (1977), 235-250. MR0422193 (54:10185)
- [PH78] S. Pohlig, M. Hellman, An Improved Algorithm for Computing Logarithms over $GF(p)$ and its Cryptographic Significance, *IEEE Transactions on Information Theory*, 24 (1978), 106-110. MR0484737 (58:4617)
- [Rab80] M. O. Rabin, Probabilistic algorithms in finite fields, *SIAM J. Comput.*, 9 (1980), 273-280. MR568814 (81g:12002)
- [Ró87] L. Rónyai, Factoring Polynomials over finite fields, *Proc. 28th IEEE FOCS (1987)* pp. 132-137; journal version appeared in *Journal of Algorithms*, 9 (1988), 391-400. MR955147 (89k:11124)
- [Ró89a] L. Rónyai, Factoring polynomials modulo special primes, *Combinatorica*, 9 (1989), 199-206. MR1030373 (90k:11161)
- [Ró90] L. Rónyai, Computing the structure of finite algebras, *Journal of Symbolic Computation*, 9 (1990), 355-373. MR1056632 (91h:68093)
- [Ró89b] L. Rónyai, Galois Groups and Factoring Polynomials over Finite Fields, *Proc. 30th IEEE FOCS (1989)* pp. 99-104; journal version appeared in *SIAM J. on Discrete Mathematics*, 5 (1992), 345-365. MR1172743 (94a:68055)
- [Sch85] R. J. Schoof, Elliptic curves over finite fields and the computation of square roots mod p , *Mathematics of Computation*, 44 (1985), 483-494. MR777280 (86e:11122)
- [S96] G. Stein, Factoring cyclotomic polynomials over large finite fields, *Proceedings of the third international conference on finite fields and applications (1996)* pp. 349-354. MR1433158 (98b:11122)
- [S01] G. Stein, Using the theory of cyclotomy to factor cyclotomic polynomials over finite fields, *Mathematics of Computation* 70 (2001), 1237-1251. MR1709159 (2001j:11119)
- [W05] C. van de Woestijne, Deterministic equation solving over finite fields, *Proc. ISSAC 2005*, 348-353. MR2280567

COMPUTER AND AUTOMATION RESEARCH INSTITUTE OF THE HUNGARIAN ACADEMY OF SCIENCES, LÁGYMÁNYOSI U. 11, 1111 BUDAPEST, HUNGARY.

E-mail address: `Gabor.Ivanyos@sztaki.hu`

DEPARTMENT OF COMPUTER SCIENCE AND HAUSDORFF CENTER FOR MATHEMATICS, UNIVERSITY OF BONN, 53117 BONN, GERMANY

E-mail address: `marek@cs.uni-bonn.de`

COMPUTER AND AUTOMATION RESEARCH INSTITUTE OF THE HUNGARIAN ACADEMY OF SCIENCES, LÁGYMÁNYOSI U. 11, 1111 BUDAPEST, HUNGARY AND DEPARTMENT OF ALGEBRA, BUDAPEST UNIVERSITY OF TECHNOLOGY AND ECONOMICS, MŰEGYETEM RKP. 3-9, 1111 BUDAPEST, HUNGARY.

E-mail address: `lajos@ilab.sztaki.hu`

HAUSDORFF CENTER FOR MATHEMATICS, UNIVERSITY OF BONN, 53115 BONN, GERMANY

E-mail address: `ns@hcm.uni-bonn.de`