Trading GRH for Algebra: Algorithms for Factoring Polynomials and Related Structures

Gábor Ivanyos * Marek Karpinski † Lajos Rónyai ‡ Nitin Saxena §¶

Abstract

In this paper we develop techniques that eliminate the need of the Generalized Riemann Hypothesis (GRH) from various (almost all) known results about deterministic polynomial factoring over finite fields. Our main result shows that given a polynomial $f(x)$ of degree n over a finite field k, we can find in deterministic $poly(n^{\log n}, \log |k|)$ time *either* a nontrivial factor of $f(x)$ *or* a nontrivial automorphism of $k[x]/(f(x))$ of order n. This main tool leads to various new GRH-free results, most striking of which are:

1. Given a noncommutative algebra $\mathcal{A}$ of dimension n over a finite field k. There is a deterministic $poly(n^{\log n}, \log |k|)$ time algorithm to find a zero divisor in $\mathcal{A}$. This is the best known deterministic GRH-free result since Friedl and Rónyai (STOC 1985) first studied the problem of finding zero divisors in finite algebras and showed that this problem has the same complexity as factoring polynomials over finite fields.
2. Given a positive integer r such that either $8|r$ or r has at least two distinct odd prime factors. There is a deterministic polynomial time algorithm to find a nontrivial factor of the r-th cyclotomic polynomial over a finite field. This is the best known deterministic GRH-free result since Huang (STOC 1985) showed that cyclotomic polynomials can be factored over finite fields in deterministic polynomial time assuming GRH.

In this paper, following the seminal work of Lenstra (1991) on constructing isomorphisms between finite fields, we further generalize classical Galois theory constructs like cyclotomic extensions, Kummer extensions, Teichmüller subgroups, to the case of commutative semisimple algebras with automorphisms. These generalized constructs help eliminate the dependence on GRH.

1 Introduction

The problem of finding a nontrivial factor of a given polynomial over a finite field is a fundamental computational problem. There are many problems whose known algorithms first

*Computer and Automation Research Institute of the Hungarian Academy of Sciences (MTA SZTAKI), Lágymányosi u. 11, 1111 Budapest, Hungary. E-mail: Gabor.Ivanyos@sztaki.hu

†Department of Computer Science and Hausdorff Center for Mathematics, University of Bonn, 53117 Bonn, Germany. E-mail: marek@cs.uni-bonn.de

‡MTA SZTAKI and Department of Algebra, Budapest University of Technology and Economics, Műegyetem rkp. 3-9, 1111 Budapest, Hungary. E-mail: lajos@ilab.sztaki.hu

§Hausdorff Center for Mathematics, Endenicher Allee 62, 53115 Bonn, Germany. E-mail: ns@hcm.uni-bonn.de

¶The authors thank the Hausdorff Research Institute for Mathematics for its kind support.

require factoring polynomials. Thus, polynomial factoring is an intensely studied question and various randomized polynomial time algorithms are known – Berlekamp [Be67], Rabin [Rab80], Cantor and Zassenhaus [CZ81], von zur Gathen and Shoup [GS92], Kaltofen and Shoup [KS98] – but its deterministic complexity is a longstanding open problem. There are although several partial results known about the deterministic complexity of polynomial factoring based on the conjectured truth of the generalized Riemann Hypothesis (GRH). 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 primitive 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 [Ev89]).

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

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 *subexponential* time $poly(n^{\log n}, \log |k|)$. 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].

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, except for finitely many primes p. Other results of a similar flavor are: Evdokimov [Ev89] showed under GRH that $f(x)$ can be factored in deterministic polynomial time if it has a *solvable* Galois group while Huang [Hua85] showed under GRH that $f(x)$ can be factored in deterministic polynomial time if it has an *Abelian* Galois group.

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].

Polynomial factoring has several applications both in the real world - coding theory and cryptography - and in fundamental computational algebra problems. The latter kind of applications are 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 subexponential time algorithms. In this work we give an unconditional version of this result. We show that if the given algebra is noncommutative then in fact we can find a zero divisor in deterministic subexponential time *without* needing GRH.

1.1 Our Results and Techniques

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. As a first step in that direction we give in this work GRH free versions of all the results mentioned above. In these versions the basic tool is that we either successfully find a nontrivial factor of a polynomial $f(x)$ over a finite field k *or* we find a nontrivial automorphism of the algebra $k[x]/(f(x))$. Formally speaking the main result of the paper is:

Main Theorem: *Let $\mathcal{A}$ be a commutative semisimple algebra of dimension n over a finite field k and let $\mathcal{A}$ be given in the input in terms of basis elements over k. Then there is a deterministic algorithm which in subexponential time poly$(n^{\log n}, \log|k|)$ computes a decomposition of $\mathcal{A}$ into a direct sum $\mathcal{A}_1 \oplus \ldots \oplus \mathcal{A}_t$ and finds an automorphism of order $\dim_k \mathcal{A}_i$ of the algebra $\mathcal{A}_i$, for each $1 \leq i \leq t$.*

This main theorem can be considered as a GRH-free version of Evdokimov's factoring result [Ev94], but its proof leads us to significantly generalize standard notions and develop novel algebraic techniques that suggest a general paradigm for GRH elimination. We are going to use it as a tool for more important applications but first let us explain the importance of this result itself. It is the first deterministic subexponential 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 special case when $\mathcal{A} = k[x]/(f(x))$ with $f(x)$ splitting over k as $\prod_{j=1}^{n}(x-\alpha_j)$, with $\alpha_1, \ldots, \alpha_n$ all distinct, we have $\mathcal{A} \cong \oplus_{j=1}^{n} k[x]/(x-\alpha_j)$. The above algorithm either gives $t > 1$ components of $\mathcal{A}$ – in which case it effectively yields a nontrivial factor of $f(x)$ – or $t = 1$ and it gives an automorphism σ of $\mathcal{A}$ of order n, thus yielding n distinct "roots" of $f(x)$ – x, $\sigma(x), \ldots, \sigma^{n-1}(x)$ – 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.

The key ideas in finding a nontrivial automorphism of a given commutative semisimple $\mathcal{B}$-algebra $\mathcal{A}$ over a finite field $k \subseteq \mathcal{B}$ are as follows. We 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 a standard homomorphism of $\mathcal{A} \otimes_{\mathcal{B}} \mathcal{A}$ onto $\mathcal{A}$ and has rank ("dimension") $\mathrm{rk}_{\mathcal{B}}\mathcal{A}(\mathrm{rk}_{\mathcal{B}}\mathcal{A} - 1)$ over $\mathcal{B}$. The algebra $\mathcal{A}$ gets 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 $\phi(\mathcal{B}) \cong \mathcal{B}$. Also, 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 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 $\mathrm{rk}_{\mathcal{A}}\mathcal{A}' = \mathrm{rk}_{\mathcal{B}}\mathcal{A}'/\mathrm{rk}_{\mathcal{B}}\mathcal{A}$ is odd or even.

(1) If the rank $\mathrm{rk}_{\mathcal{B}}\mathcal{A}$ is even then $\mathrm{rk}_{\mathcal{A}}\mathcal{A}'$ 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 $\mathcal{B}$-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 $\mathcal{A}'$ 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 $\mathrm{rk}_{\mathcal{A}}\mathcal{I}$ 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. **(2)** If the rank $\mathrm{rk}_{\mathcal{B}}\mathcal{A}$ is odd then $\mathrm{rk}_{\mathcal{A}}\mathcal{A}'$ 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 $\mathrm{rk}_{\mathcal{A}}\mathcal{A}'/2 = (\mathrm{rk}_{\mathcal{B}}\mathcal{A} - 1)/2$. 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 2 gets handled in 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 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 depth of the recursion is at most $\log \mathrm{rk}_{\mathcal{B}}\mathcal{A}$. This gives an $n^{\log n}$ term in the time complexity analysis.

Roots of unity play a significant role in gluing automorphisms (i.e. in extending an automorphism of a subalgebra, of elements fixed by another automorphism, to the whole algebra). The gluing process is described in Section 4.4. As we do not know roots of unity in k we resort to attaching virtual r-th roots of unity for a suitable prime r, i.e. working in the cyclotomic extension $k[\zeta_r] := k[x]/(\sum_{i=1}^{r-1} x^i)$ and $\mathcal{A}'[\zeta_r] := k[\zeta_r] \otimes_k \mathcal{A}'$. We then need to generalize standard algebraic constructions, like *Kummer extensions* and *Teichmüller subgroups* which were first used in a context similar to ours by Lenstra [L91] to find isomorphisms between fields, to our situation of commutative semisimple algebras.

The above theorem and its proof techniques have important applications. The first one is in finding zero divisors in a noncommutative algebra.

Application 1: *Let $\mathcal{A}$ be an algebra of dimension n over a finite field k and let $\mathcal{A}$ be given in the input in terms of basis elements over k. Assume that $\mathcal{A}$ is noncommutative. Then there is a deterministic algorithm which finds a zero divisor in $\mathcal{A}$ in time $poly(n^{\log n}, \log |k|)$.*

The previous best result was due to Rónyai [Ró90] who gave an algorithm invoking polynomial factorization over finite fields and hence taking subexponential 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 subexponential time without needing GRH.

If $\mathcal{A}$ is a simple algebra over the finite field k then 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 Application 1 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. Thus, for constant m, Application 1 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, algebra isomorphism problem over finite fields 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 of constructing isomorphism with the algebra of matrices over the rationals has a surprising application to rational parametrization of certain curves, see [GHPS06].

The techniques used to prove Main Theorem can be applied to find a nontrivial factor of an r-th cyclotomic polynomial over a finite field k, for almost all r's, in deterministic polynomial time.

Application 2: *Let r be a positive integer such that the multiplicative group $\mathbb{Z}_r^*$ is noncyclic and let $\phi_r(x)$ be the r-th cyclotomic polynomial. Then we can find a nontrivial factor of $\phi_r(x)$ over a finite field k in deterministic poly$(r, \log |k|)$ time.*

Roots of an r-th cyclotomic polynomial over k are the r-th roots of unity and thus naturally related to all polynomial factoring algorithms. Assuming GRH several algorithms are known to factor these important polynomials (see [Ev89]). The above result gives the first deterministic polynomial time algorithm to nontrivially factor "most" of the cyclotomic polynomials without assuming GRH.

The third application of the techniques used to prove Main Theorem is in 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].

Application 3: *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 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* $\deg f$.

The fourth application of the techniques used to prove Main Theorem 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].

Application 4: *Let $f(x)$ be a polynomial of degree n, that splits into linear factors over $\mathbb{F}_p$. Let $r_1 < \ldots < r_t$ be the prime factors of $(p-1)$. Then by a deterministic algorithm of running time 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 case $n \not| \ lcm\{r_i - 1 | 1 \le i \le t\}$.*

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

1.2 Organization

In Section 2 we collect various standard objects 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 log for elements of prime-power order, finding a free base 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 such 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* wrt 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* over algebras, that we are going to use. We investigate there the action of the automorphisms of an algebra $\mathcal{A}$ on a certain subgroup, *Teichmüller subgroup*, of the multiplicative group of a Kummer-type extension of $\mathcal{A}$. The proofs of Applications 2 and 3 get completed in this section.

In Section 5 we apply the machinery of Section 4 to the tensor power algebras and complete the proof of Main Theorem.

In Section 6 we find suitable subalgebras of a given noncommutative algebra to invoke Main Theorem and complete the proof of Application 1.

In Section 7 we use the techniques developed for the Main Theorem in the case of special finite fields and complete the proof of Application 4.

2 Preliminaries

In this section we list some algebraic notions that we use in this work and that can be found in standard algebra texts, for example [La80].

Rings, Units and Zero-divisors: A *ring with identity* (or ring, for short) R is a set of elements together with two operations – denoted by *addition* $+$ and *multiplication* $\cdot$ – such that $(R,+)$ is an Abelian group, $\cdot$ is associative, distributes over $+$ and has an *identity* element 1_R. Note that the set R^*, containing all the elements of R that have a multiplicative inverse, is a multiplicative group called the group of *units*. For a prime integer r we call a unit x an *r-element* if the multiplicative order of x is a power of r. An element x is called a *zero divisor* if $x \neq 0$ and there exist nonzero $y, y' \in \mathcal{A}$ such that $yx = xy' = 0$.

Modules: Let $(R,+,\cdot)$ be a commutative ring and $(M,+)$ be an Abelian group. We call M an *R-module* wrt an operation $R \times M \to M$ (called *scalar multiplication* and denoted as rx for $r \in R$ and $x \in M$) if for all $r,s \in R; x,y \in M$, we have: $r(x+y) = rx + ry$; $(r+s)x = rx + sx$; $(rs)x = r(sx)$ and $1x = x$. Note that a vector space V over a field $\mathbb{F}$ is also an $\mathbb{F}$-module.

Free and Cyclic: For an R-module M, a set $E \subset M$ is called a *free basis* of M if: E is a *generating set* for M, i.e. every element of M is a finite sum of elements of E multiplied by coefficients in R, and E is a *free set*, i.e. for all $r_1, \ldots, r_n \in R$; $e_1, \ldots, e_n \in E$, $r_1e_1 + \cdots + r_ne_n = 0$ implies that $r_1 = \cdots = r_n = 0$. A *free module* is a module with a free basis. $|E|$ is called the *rank* or *dimension* 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.

Algebras: Let $(R,+,\cdot)$ be a commutative ring and $(\mathcal{A},+,\cdot)$ be a ring which is also an R-module, where the additive operation of $\mathcal{A}$ as a module coincides with $+$. We say that $\mathcal{A}$ is an associative R-algebra with identity (or just an *R-algebra* for short) if multiplication

by elements of R commutes with multiplication by elements of $\mathcal{A}$: for every $r \in R$ and for every $a, b \in \mathcal{A}$ we have $r(ab) = (ra)b = a(rb)$.

Subalgebras: A *subalgebra* $\mathcal{B}$ of an R-algebra $(\mathcal{A}, +, \cdot)$ is just a submodule of $\mathcal{A}$ closed under multiplication. In this paper unless otherwise stated, by a subalgebra of $\mathcal{A}$ we mean a subalgebra containing the identity element $1_{\mathcal{A}}$. 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.

Presentation: In this work we will consider only k-algebras $\mathcal{A}$ that are *finite dimensional* over a finite field k. So we can assume that an algebra $\mathcal{A}$ is always presented in the input-output in terms of an additive basis of $(\mathcal{A}, +)$ over k, i.e. there are *basis elements* $b_1, \ldots, b_n \in \mathcal{A}$ such that $\mathcal{A} = kb_1 + \cdots + kb_n$ and furthermore $a_{i,j,\ell} \in k$ are given such that $b_i \cdot b_j = \sum_\ell a_{i,j,\ell} b_\ell$. Such an n is called the *dimension*, $dim_k \mathcal{A}$, of $\mathcal{A}$ over k.

Extension: If $\mathcal{B}$ is a commutative k-algebra and a $\mathcal{B}$-algebra $\mathcal{A}$ is also a free module over $\mathcal{B}$ then we call $\mathcal{A}$ an *algebra extension* or an *extension algebra* over $\mathcal{B}$. This terminology is justified by the fact that $\mathcal{B}$ is embedded into (the center of) $\mathcal{A}$ by the map $b \mapsto b1_{\mathcal{A}}$. We denote the rank ("dimension") of $\mathcal{A}$ as a $\mathcal{B}$-module by $\mathrm{rk}_{\mathcal{B}}\mathcal{A}$ or $[\mathcal{A} : \mathcal{B}]$. We sometimes use this notation also when there is an implicit embedding of $\mathcal{B}$ in $\mathcal{A}$.

Primitive Element: We call an algebra extension $\mathcal{A}$ over $\mathcal{B}$ *simple* if there is an $\alpha \in \mathcal{A}$ such that $\{1, \alpha, \ldots, \alpha^{n-1}\}$ forms a free basis of $\mathcal{A}$ over $\mathcal{B}$. We call α a *primitive element* and write $\mathcal{A} = \mathcal{B}[\alpha]$.

Following is a version of the standard *Primitive Element Theorem*.

Fact 1. *If $K \supseteq F$ are fields such that char F is 0 or $> [K : F]^2$, then K has a primitive element over F.*

There are two natural operations defined on algebras – the *direct sum* and the *tensor product* – each constructs a bigger algebra.

Direct Sum: Let $(\mathcal{A}_1, +, \cdot)$ and $(\mathcal{A}_2, +, \cdot)$ be two algebras. Then the *direct sum algebra*, $\mathcal{A}_1 \oplus \mathcal{A}_2$, is the set $\{(a_1, a_2) \mid a_1 \in \mathcal{A}_1, a_2 \in \mathcal{A}_2\}$ together with component-wise addition and multiplication operations. In a similar vein, for subalgebras $\mathcal{A}_1, \mathcal{A}_2$ of an algebra $\mathcal{A}$ we write $\mathcal{A} = \mathcal{A}_1 \oplus \mathcal{A}_2$, if $\mathcal{A} = \mathcal{A}_1 + \mathcal{A}_2$ and $\mathcal{A}_1, \mathcal{A}_2$ are *orthogonal* i.e. $\forall\ a_1 \in \mathcal{A}_1$, $a_2 \in \mathcal{A}_2$, $a_1 a_2 = a_2 a_1 = 0$.

Tensor Product: Furthermore, if $\mathcal{B}$ is a commutative algebra such that $\mathcal{A}_1, \mathcal{A}_2$ are $\mathcal{B}$-algebras of dimensions n_1, n_2 respectively over $\mathcal{B}$ then their *tensor product algebra wrt* $\mathcal{B}$, $\mathcal{A}_1 \otimes_{\mathcal{B}} \mathcal{A}_2$, is the set $\{a_1 \otimes a_2 \mid a_1 \in \mathcal{A}_1, a_2 \in \mathcal{A}_2\}$ naturally viewed as a $\mathcal{B}$-module having the multiplication operation: $(a_1 \otimes a_2) \cdot (a_1' \otimes a_2') = (a_1 a_1' \otimes a_2 a_2')$ for all $a_1, a_1' \in \mathcal{A}_1$ and $a_2, a_2' \in \mathcal{A}_2$. Note that the tensor product algebra has dimension $n_1 n_2$ over $\mathcal{B}$. Thus, if $\mathcal{B}$ is finite then $|\mathcal{A}_1 \oplus \mathcal{A}_2| = |\mathcal{B}|^{n_1+n_2}$, while $|\mathcal{A}_1 \otimes_{\mathcal{B}} \mathcal{A}_2| = |\mathcal{B}|^{n_1 n_2}$.

Nilpotent and Idempotent: In an algebra $\mathcal{A}$ we call an element $x \in \mathcal{A}$ *nilpotent* if $x^m = 0$ for some $m \in \mathbb{Z}$, while we call x *idempotent* if $x^2 = x$. 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 0 or 1.

Decomposability: An algebra $\mathcal{A}$ is called *indecomposable* if there are no nonzero algebras R, S such that $\mathcal{A} \cong R \oplus S$.

Following are some standard facts relating decomposability to idempotents in commutative algebras.

Fact 2. *Let $\mathcal{A}$ be a commutative algebra then:*
(1) $\mathcal{A}$ decomposes iff $\mathcal{A}$ has a nontrivial idempotent.
(2) If e is an idempotent in $\mathcal{A}$ then $\mathcal{A} \cong e\mathcal{A} \oplus (1-e)\mathcal{A}$.
(3) If e is a primitive idempotent in $\mathcal{A}$ then $e\mathcal{A}$ is indecomposable.

Ideal: An *ideal* I of an algebra $\mathcal{A}$ is a subset that is an additive subgroup of $\mathcal{A}$, is closed under multiplication and it contains both $aI := \{a \cdot i \mid i \in I\}$; $Ia := \{i \cdot a \mid i \in I\}$ for all $a \in \mathcal{A}$. 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.

Semisimplicity: An algebra $\mathcal{A}$ is called *simple* if it has no nontrivial ideal. An algebra is called *semisimple* if it is a direct sum of simple algebras.

Following are some standard facts about commutative semisimple algebras.

Fact 3. *Let $\mathcal{A}$ be a commutative semisimple algebra then:*
(1) $\mathcal{A}$ is a direct sum of fields.
(2) If I is an ideal of $\mathcal{A}$ and $I^{\perp} := \{a \in \mathcal{A} \mid aI = 0\}$ (called the complement *of I) then $\mathcal{A} = I \oplus I^{\perp}$. Furthermore, there exists an idempotent e of $\mathcal{A}$ such that $I = e\mathcal{A}$ thus giving an explicit projection from $\mathcal{A}$ to I.*

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

Fact 4. *Any semisimple algebra $\mathcal{A}$ is isomorphic to a direct sum of $n_i \times n_i$ matrix algebras over division rings D_i (i.e. D_i satisfies all field axioms except commutative multiplication). Both the n_i's and D_i's are uniquely determined up to permutation of the indices i.*

Morphisms: Let ϕ be a map between two algebras $\mathcal{A}$, $\mathcal{B}$. If ϕ preserves the addition and multiplication operations of the algebras then we call it a *homomorphism*. If the homomorphism ϕ is injective then we call it an *embedding*. If the homomorphism ϕ is both injective and surjective then we call it an *isomorphism*. A homomorphism from an algebra to itself is called an *endomorphism*. An isomorphism from an algebra to itself is called an *automorphism*. A set S is said to be *invariant* under the automorphism ϕ of $\mathcal{A}$ if for all $s \in S$, $\phi(s) \in S$. ϕ is said to *fix* S if ϕ fixes each element of S, i.e. for all $s \in S$, $\phi(s) = s$. The group of S-automorphisms of $\mathcal{A}$, $Aut_S(\mathcal{A})$, is the set of all automorphisms of $\mathcal{A}$ that fix S.

Throughout this paper all algebras are algebras with identity elements. Unless otherwise stated explicitly, by a subalgebra we mean a subalgebra *containing the identity element*. Thus, in this strict sense a proper ideal is *not* considered as a subalgebra. In the rest of this section $\mathcal{A}$ stands for a commutative semisimple algebra over the finite field k.

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 2.1. *Given a prime r distinct from the characteristic of a finite field k, a finite dimensional 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. There is a deterministic algorithm which computes in time $poly(r, \log|\mathcal{A}|)$:*
(1) either two non-negative integers s, s' such that $a^s - b^{s'}$ is a zero divisor in $\mathcal{A}$,
(2) or an integer $s \geq 0$ with $a^s = b$.

Proof. Let t_a be the smallest non negative 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 $a^{r^0} - 1, a^{r^1} - 1, \ldots, a^{r^{t_a}} - 1$ in $poly(\log|\mathcal{A}|)$ 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 non-negative 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 assure 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 $s_0 = 0, s_1, s_2, \ldots, s_t$ together with idempotents $e_1, \ldots, e_t$ 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 exhaustive search, in at most r rounds, an $s_1 \in \{1, \ldots, r-1\}$ 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 exhaustive search, in at most r rounds, an integer $d_{j+1} \in \{0, \ldots, r-1\}$ 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 $poly(r, \log|\mathcal{A}|)$ time. □

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. Following Lemma states the basic tool to do that.

Lemma 2.2. *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 $poly(\dim_{\mathcal{A}} V, \log|\mathcal{A}|)$.*

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. □

2.3 Automorphisms and Invariant Ideal Decompositions

Automorphisms of $\mathcal{A}$ 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 $\mathcal{A}_\Gamma$ denotes the set of the elements of $\mathcal{A}$ fixed by every $\sigma \in \Gamma$. It is obvious that $\mathcal{A}_\Gamma$ is a subalgebra of $\mathcal{A}$. For a single automorphism σ we use $\mathcal{A}_\sigma$ 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 3, $\mathcal{A} = I \oplus I^\perp$. 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 2.3. *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 $J_1, \ldots, J_s$, there is a deterministic algorithm of time complexity $poly(|\Gamma|, \log|\mathcal{A}|)$ that computes a decomposition of $\mathcal{A}$ into a sum of pairwise orthogonal ideals $I_1, \ldots, I_t$ such that:*
(1) the new decomposition is a refinement of the original one – for every $j \in \{1, \ldots, t\}$, there exists $i \in \{1, \ldots, s\}$ such that $I_j \subseteq J_i$, and
(2) the new decomposition is invariant under Γ – the group generated by Γ permutes the ideals $I_1, \ldots, I_t$, i.e. for every $\sigma \in \Gamma$ and for every index $j \in \{1, \ldots, t\}$, we have $I_j^\sigma = I_{j^\sigma}$ for some index $j^\sigma \in \{1, \ldots, t\}$.

3 Semiregularity

In this section we continue to assume that $\mathcal{A}$ is a commutative semisimple algebra over a finite field k. Given $\Gamma \subseteq \mathrm{Aut}_k(\mathcal{A})$, a basis of $\mathcal{A}_\Gamma$ can be computed by solving a system of linear equations in $\mathcal{A}$. Thus, we can apply the method of Lemma 2.2 considering $\mathcal{A}$ as a $\mathcal{A}_\Gamma$-module wrt 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 2.2 is ineffective.

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

We have the following characterization of semiregularity.

Lemma 3.1. *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.

Let e be a primitive idempotent of $\mathcal{A}$. We denote the stabilizer of e in G by G_e, i.e, $G_e = \{\sigma \in G | e^\sigma = e\}$. Let C be a complete set of right coset representatives modulo G_e in G. The orbit of e under G is $\{e^\gamma | \gamma \in C\}$ and they are $|G : G_e|$ many pairwise orthogonal primitive idempotents in $\mathcal{A}$. This means that the component field $e\mathcal{A}$ is sent to the other component fields $\{e^\gamma \mathcal{A} | \gamma \in C\}$ by G. Thus, the element $f := \sum_{\gamma \in C} e^\gamma \in \mathcal{A}_G$ 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 $\lambda : G_e \to Aut_k(e\mathcal{A})$ whose kernel say is N_e, so $N_e = \{\sigma \in G | \sigma \text{ fixes } e\mathcal{A}\}$ is a normal subgroup of G_e, thus 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 onto 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 dimension $|G : G_e|$ 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 non-identity automorphism in G that fixes $e\mathcal{A}$, thus equivalently for every nontrivial ideal I of $\mathcal{A}$ there is no non-identity automorphism in G that fixes I. 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 dimension $|G|$. □

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

Proposition 3.2. *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 $poly(|\Gamma|, \log |\mathcal{A}|)$ 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 $\mathcal{A}_\Gamma$ by linear algebra over k. We can assume that $\mathcal{A}$ is a free $\mathcal{A}_\Gamma$-module otherwise the algorithm in Lemma 2.2 finds a zero divisor. By Lemma 3.1

$|G| \geq \dim_{\mathcal{A}_\Gamma} \mathcal{A} =: m$ so try to enumerate $(m+1)$ different elements in the group G. If we are unable to get that many elements then, by Lemma 3.1, 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 from the proof of Lemma 3.1 we obtain $|G : N_e| \leq m$ which means, by 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 $\mathcal{A}_\sigma$ 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 $\mathcal{A}_\sigma$. We have from the proof of Lemma 3.1 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 $\mathcal{A}_\sigma$ and hence we can find a zero divisor of $\mathcal{A}$ using the method of Lemma 2.2. $\square$

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 $\mathcal{B}$. We give below a Galois theory-like characterization of $G_\mathcal{B}$.

Proposition 3.3. *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 if $\mathcal{B} \neq \mathcal{A}_{G_\mathcal{B}}$.*

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 Fact 1 gives a deterministic polynomial time algorithm for finding a primitive element x of $\mathcal{A}$ over k, see [GI00].

Let $|G| = d$. We may assume that the elements $1, x, x^2, \ldots, x^{d-1}$ form a free basis of $\mathcal{A}$ over $\mathcal{A}_G$ since otherwise we find a zero divisor in $\mathcal{A}$ using the method of Lemma 2.2. Let $x^d = \sum_{i=0}^{d-1} a_i x^i$ with $a_i \in \mathcal{A}_G$ and let $f(X) := X^d - \sum_{i=0}^{d-1} a_i X^d \in \mathcal{A}_G[X]$. 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 $f(X)$. Again by Lemma 2.2 we may assume that $\mathcal{A}$ is a $\mathcal{B}$-module with $\{1, x, \ldots x^{m-1}\}$ 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 $\{1, \ldots, x^{m-1}\}$ 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 $\phi_\sigma : \mathcal{B}[X] \to \mathcal{A}$ 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 3.2 we can find a zero divisor in $\mathcal{A}$ if the maps μ_σ ($\sigma \in K$) generate a non-semiregular 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 3.1, that the set $\{\mu_\sigma | \sigma \in K\}$ is a group say H. 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 3.1, that H is a subgroup of G. Thus, by Lemma 3.1, $[\mathcal{A} : \mathcal{A}_H] = |H| = |K| = [\mathcal{A} : \mathcal{B}]$ 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 $[\mathcal{A} : \mathcal{A}_H] < [\mathcal{A} : \mathcal{A}_{G_{\mathcal{B}}}] \leq [\mathcal{A} : \mathcal{B}]$ which is a contradiction). Thus, if none of the above steps yield a zero divisor then $\mathcal{B} = \mathcal{A}_{G_{\mathcal{B}}}$. $\square$

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, an r-th primitive 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 algebra $\mathcal{A}$ 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 *char* k. By $k[\zeta_r]$ we denote the factor algebra $k[X]/(\sum_{i=1}^{r-1} X^i)$ and $\zeta_r := X \pmod{\sum_{i=1}^{r-1} X^i}$. Then $k[\zeta_r]$ is an $(r-1)$-dimensional k-algebra with basis $\{1, \zeta_r, \ldots, \zeta_r^{r-2}\}$ 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 $\zeta_2 = -1$, $\mathcal{A}[\zeta_2] = \mathcal{A}$ 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 $\mathcal{A} = \mathcal{A}_1 \oplus \ldots \oplus \mathcal{A}_t$ then $\mathcal{A}[\zeta_r] = \mathcal{A}_1[\zeta_r] \oplus \ldots \oplus \mathcal{A}_t[\zeta_r]$, 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 $\{\zeta_r, \ldots, \zeta_r^{r-1}\}$. 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 a_i's are equal iff $x \in \mathcal{A}$. □

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 r-elements of $\mathcal{A}[\zeta_r]^*$. Note that $\mathcal{A}[\zeta_r]_r^*$ is of an r-power size and is also 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 map $a \mapsto a^{r^{t-1}}$ corresponds just to the projection of the multiplicative group $\mathbb{Z}_{r^t}^* \cong (\mathbb{Z}_{r-1}, +) \oplus (\mathbb{Z}_{r^{t-1}}, +)$ to the first component. This together with the fact that for any $x \in \mathcal{A}[\zeta_r]_r^*$, $x^{r^t} = 1$ 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], $\omega(a) : x \mapsto x^{\omega(a)} := x^{a^{r^{t-1}}}$ from $\mathcal{A}[\zeta_r]_r^*$ to itself. Note that we use the term $\omega(a)$ for both the above map as well as the residue of $a^{r^{t-1}}$ modulo r^t.

Note that the map $\omega(a)$ is an automorphism of the group $\mathcal{A}[\zeta_r]_r^*$ and it commutes with all the endomorphisms of the group $\mathcal{A}[\zeta_r]_r^*$. Also, the map $a \mapsto \omega(a)$ is a group embedding $\mathbb{Z}_r^* \to \mathrm{Aut}(\mathcal{A}[\zeta_r]_r^*)$.

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 $\mathcal{A}[\zeta_r]^*$:

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

Note that for $r = 2$, $T_{\mathcal{A},2}$ is just the 2-Sylow 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 non-cylicness 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 poly$(r, \log|\mathcal{A}|)$ time.*

Proof. Suppose the subgroup generated by u and v is not cyclic. Then, by Lemma 2.1 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 $a \in \{1, \ldots, r-1\}$, 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 $a \in \{1, \ldots, r-1\}$, $e^{\rho_a} \in I$ which means that I is invariant under Δ_r. □

Now we are in a position to define what we call 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 $poly(r, \log|\mathcal{A}|)$ 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 $\sqrt[s]{c}$.

Clearly, $\mathcal{A}[\zeta_r][\sqrt[s]{c}]$ is a free $\mathcal{A}[\zeta_r]$-module of rank s with basis $\{1, \sqrt[s]{c}, \ldots, \sqrt[s]{c}^{s-1}\}$. If $c \in T_{\mathcal{A},r}$ then $\sqrt[s]{c}$ is an r-element of $\mathcal{A}[\zeta_r][\sqrt[s]{c}]^*$ 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 $\mathcal{A}[\zeta_r][\sqrt[s]{c}]$ that sends $\sqrt[s]{c}$ to $(\sqrt[s]{c})^{\omega(a)}$.*

Proof. For a $\rho_a \in \Delta_r$ let $\tilde{\rho}_a$ denote the map from $\mathcal{A}[\zeta_r][Y]$ to $\mathcal{A}[\zeta_r][\sqrt[s]{c}]$ that fixes $\mathcal{A}$, sends ζ_r to ζ_r^a and Y to $(\sqrt[s]{c})^{\omega(a)}$. 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 $\mathcal{A}[\zeta_r][\sqrt[s]{c}]$ that sends $\sqrt[s]{c}$ to $(\sqrt[s]{c})^{\omega(a)}$. 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 $\mathcal{A}[\zeta_r][\sqrt[s]{c}]$ we get that $\tilde{\rho}_a$ is also an automorphism of $\mathcal{A}[\zeta_r][\sqrt[s]{c}]$, completing the proof. In the rest of the paper we will use ρ_a also to refer to the automorphism $\tilde{\rho}_a$. □

We saw above automorphisms of the Kummer extension $\mathcal{A}[\zeta_r][\sqrt[s]{c}]$ 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 automorphisms of $\mathcal{A}[\zeta_r][\sqrt[s]{c}]$ 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 $\mathcal{A}[\zeta_r][Y]/(Y^r - c)$. Thus implying 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 3.1 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 3.1: $\dim_k \mathcal{A}[\zeta_r][\sqrt[r]{c}]_{\Delta_r} = \dim_k \mathcal{A}[\zeta_r][\sqrt[r]{c}]/|\Delta_r| = r \dim_k \mathcal{A} = |(\sigma)| \dim_k \mathcal{A}$. This again implies that σ is a semiregular automorphism of $\mathcal{A}[\zeta_r][\sqrt[r]{c}]_{\Delta_r}$. □

4.2 $\mathcal{A}$ and the Kummer extension of $\mathcal{A}_\tau$, where $\tau \in Aut_k(\mathcal{A})$

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

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

Proof. Observe that if $\xi \in \mathcal{A}$ is a root of $1 + X + \ldots + X^{r-1}$ then so is every power ξ^i $(i = 1, \ldots, r-1)$. Take an element $y \in \mathcal{A} \setminus \mathcal{A}_\tau$ 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 $(y, \xi^0) = y + y^\tau + \ldots + y^{\tau^{r-1}} \in \mathcal{A}_\tau$ as $\tau^r = id$, while $\sum_{j=0}^{r-1}(y, \xi^j) = ry + \sum_{i=1}^{r-1}\sum_{j=0}^{r-1} \xi^{ij} y^{\tau^i} = ry + \sum_{i=1}^{r-1} y^{\tau^i} \sum_{j=0}^{r-1}(\xi^i)^j = ry \notin \mathcal{A}_\tau$. It follows that for some $1 \leq j \leq (r-1)$, $(y, \xi^j) \notin \mathcal{A}_\tau$, fix this j. 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: $x^\tau = ((y, \xi^j)^\tau)^l = (\xi^{-j}(y, \xi^j))^l = \xi x$. □

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 $\mathcal{A}_\tau$.

Embedding $Aut_k(\mathcal{A})$ in $Aut_k(\mathcal{A}[\zeta_r])$: 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. *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 $r \neq$ char k, we can find in deterministic poly$(\log|\mathcal{A}|)$ time an element $x \in T_{\mathcal{A},r}$ such that $x^\tau = \zeta_r x$.*

Any such x satisfies $c := x^r \in T_{\mathcal{A}_\tau,r}$ and defines an isomorphism $\phi : \mathcal{A}_\tau[\zeta_r][\sqrt[r]{c}] \cong \mathcal{A}[\zeta_r]$ 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 proof idea is to first apply Lemma 4.5 to find a nonzero $x \in \mathcal{A}[\zeta_r]$ such that $x^\tau = \zeta_r x$. Note that this x maybe 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 an $y \in \mathcal{A}[\zeta_r]^*$ such that $y^\tau = \zeta_r y$. Secondly, 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 $|\mathcal{A}[\zeta_r]_r^*| =: r^t$, $\ell := |\mathcal{A}[\zeta_r]^*|/r^t$ and $m := (-\ell)^{-1} \pmod r$. Note that ℓ can be calculated from the sizes of the simple components of $\mathcal{A}[\zeta_r]$ which in turn can be easily computed by using the standard distinct degree factorization of polynomials over finite fields. Thus, we can compute the element $z := y^{\ell m}$. By the definition of ℓ and y, $z \in \mathcal{A}[\zeta_r]_r^*$ 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 the c as x^r. From the properties of x, $c \in \mathcal{A}[\zeta_r]_\tau = \mathcal{A}_\tau[\zeta_r]$ and hence $c \in T_{\mathcal{A}_\tau,r}$.

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 $\mathcal{A}_\tau[\zeta_r]$. It is obvious from $c = x^r$ that ϕ is a homomorphism. 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 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 $\sum_{i=1}^{r-1} a_i x^{i-1}$ to zero. We can now repeat the argument and deduce that 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. To this end let $\mathcal{B}$ denote the image of ϕ. Then $\mathcal{B}$ is the subalgebra of $\mathcal{A}[\zeta_r]$ generated by $\mathcal{A}_\tau[\zeta_r]$ and x, thus $\mathcal{B}$ is τ-invariant. Suppose we can show τ semiregular on $\mathcal{B}$. Then by Lemma 3.1, $\dim_k \mathcal{B} = r \dim_k \mathcal{B}_\tau$, this together with $\mathcal{B}_\tau$ containing $\mathcal{A}_\tau[\zeta_r]$ and the injectivity of ϕ means that $\dim_k \mathcal{B} \geq r \dim_k \mathcal{A}_\tau[\zeta_r] = r \dim_k \mathcal{A}[\zeta_r]_\tau$ which is further equal to $\dim_k \mathcal{A}[\zeta_r]$ as τ is semiregular on $\mathcal{A}[\zeta_r]$. Thus, $\dim_k \mathcal{B} \geq \dim_k \mathcal{A}[\zeta_r]$ which obviously means that ϕ is indeed surjective.

It remains to prove the semiregularity of τ on $\mathcal{B}$. Assume for contradiction that I is a nonzero ideal of $\mathcal{B}$ such that τ fixes I and e be the identity element of I. Then $(ex)^\tau = ex$. On the other hand, as $e^\tau = e$ and $x^\tau = \zeta_r x$, we have $(ex)^\tau = \zeta_r ex$. Combining the two equalities we obtain that $(ex)(\zeta_r - 1) = 0$. Note that if $r = 2$ then *char* $k > 2$ and $(\zeta_r - 1)$ is not a zero divisor and if $r > 2$ then $\mathcal{A}[\zeta_r]$ is a free $\mathcal{A}$-module with basis $\{1, \ldots, \zeta_r^{r-2}\}$. Thus, $x(\zeta_r - 1)$ is invertible in all cases, implying $e = 0$ which is a contradiction. Thus τ is indeed semiregular on $\mathcal{B}$ completing the proof that ϕ is an isomorphism. □

4.3 Zero Divisors using Noncyclic Groups: Proof of Application 2

In this part we prove Application 2 by proving the following stronger result.

Theorem 4.7. *Given a commutative semisimple algebra $\mathcal{A}$ over a finite field k together with a noncyclic group G of k-automorphisms of $\mathcal{A}$ (in terms of generators), one can find*

a zero divisor in $\mathcal{A}$ in deterministic polynomial time.

Proof. Notice that since G is noncyclic, the algebra $\mathcal{A}$ is certainly not a field and zero divisors do exist. We assume that G is semiregular otherwise we can efficiently find a zero divisor in $\mathcal{A}$ by Proposition 3.2. We can also assume that $|G|$ is not divisible by *char* k otherwise *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 $G \to Aut(H_r)$. 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 $P_r \to G/K_r \cong Aut(H_r)$. Thus the normal series $K_r \lhd G$ can be refined to $K_r \trianglelefteq N_r \lhd G$ such that $|P_r| = |G/N_r|$. Since we have this for every r dividing $|G|$, it follows that G is a direct product of its Sylow subgroups. 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, $Aut(H_r)$ is cyclic and finally P_r is cyclic, for every odd prime $r||G|$.

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 $Aut(H_2) \cong \mathbb{Z}_{2^k}^*$. As P_2 injectively embeds in $Aut(H_2)$ 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, $\mathcal{A}_{\sigma_1}$ 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 $\mathcal{A}_{\sigma_1}$ 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. □

Now we can give a proof of Application 2. Let r be a positive integer such that the multiplicative group $\mathbb{Z}_r^*$ is noncyclic and let $\phi_r(x)$ be the r-th cyclotomic polynomial. We can assume r to be coprime to *char* k as otherwise we factor $\phi_r(x)$ simply by using Berlekamp's algorithm for polynomial factoring. Define $\mathcal{A} := k[x]/(\phi_r(x))$, it is clearly

a commutative semisimple algebra of dimension $\phi(r)$ over k. Moreover, if $\zeta_r \in \overline{k}$ is a primitive r-th root of unity then: $\phi_r(x) = \prod_{i \in \mathbb{Z}_r^*}(x - \zeta_r^i)$. This implies that for any $i \in \mathbb{Z}_r^*$, $\phi_r(x)|\phi_r(x^i)$ and if for a $g(X) \in k[X]$, $\phi_r(x)|g(x^i)$ then $\phi_r(X)|g(X)$ as well. In other words for any i coprime to r the map $\rho_i : x \to x^i$ is a k-automorphism of $\mathcal{A}$. Consider the group $G := \{\rho_i | i \in \mathbb{Z}_r^*\}$, it is clearly isomorphic to the multiplicative group $\mathbb{Z}_r^*$, which is noncyclic for our r. Thus, G is noncyclic and we can find a zero divisor $a(x) \in \mathcal{A}$ by Theorem 4.7. Finally, the gcd of $a(x)$ and $\phi_r(x)$ gives a nontrivial factor of $\phi_r(x)$.

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 $F_j(X) = (X^2 - X + 1)^3 - \frac{j}{2^8}X^2(X-1)^2$ 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 4.7 gives a partial factorization of the polynomial $F_j(X)$ modulo p where p is odd and j is coprime to p.

4.4 Extending Automorphisms of $\mathcal{A}_\tau$ to $\mathcal{A}$, where $\tau \in Aut_k(\mathcal{A})$

Lemma 4.8. *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 $\mathcal{A}_\tau$. Assume that the order of τ is coprime to char k. Then in deterministic poly$(\log|\mathcal{A}|)$ time we can compute either a zero divisor in $\mathcal{A}$ or a 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 r_i's are primes (not necessarily distinct). Cleary 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 3.2. 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, $c := x^r \in T_{\mathcal{A}_\tau,r}$ and $c^\mu \in T_{\mathcal{A}_\tau,r}$. 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 2.1. 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 isomorphism maps $\phi : \mathcal{A}_\tau[\zeta_r][\sqrt[r]{c}] \to \mathcal{A}[\zeta_r]$ and $\phi' : \mathcal{A}_\tau[\zeta_r][\sqrt[r]{c^\mu}] \to \mathcal{A}[\zeta_r]$ sending $\sqrt[r]{c}$ to x and $\sqrt[r]{c^\mu}$ to x' respectively; both fixing $\mathcal{A}_\tau[\zeta_r]$. We can naturally extend μ to an isomorphism map $\mu'' : \mathcal{A}_\tau[\zeta_r][\sqrt[r]{c}] \to \mathcal{A}_\tau[\zeta_r][\sqrt[r]{c^\mu}]$. 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 $\mathcal{A}$. □

4.5 Zero Divisors using Galois Groups: Proof of Application 3

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 [Ró89b] 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 [Ró89b]:

Theorem 4.9. *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 $\mathcal{B}$ (with k embedded) of a subalgebra of $\mathcal{A}$. Then in deterministic 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 G is noncyclic then using Theorem 4.7 we can find a nontrivial ideal I of $\mathcal{A}$. If G is cyclic then using Proposition 3.3 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 is trivially a normal subgroup of G and 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 2.3, we find an ideal J of $\mathcal{A}$ such that the ideals $\{J^\sigma | \sigma \in G\}$ 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 $G_1 := \{\sigma \in G | J^\sigma = J\}$ acts semiregularly on J and for coset representatives C of G/G_1: $\mathcal{A} = \oplus_{\sigma \in C} J^\sigma$. 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 last paragraph. □

The following Corollary gives the proof of a slightly stronger version of Application 3.

Corollary 4.10. *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 non-constant divisor of $f(X)$ in $\mathbb{F}_p[X]$. Then by a deterministic 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 [Ró89b], 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. Define $\mathcal{A} := \mathbb{Z}[\alpha]/(p)$ and using the method described in Section 4 of [Ró89b], we efficiently compute a group G of automorphisms of $\mathcal{A}$ which is isomorphic to the Galois group of α over 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 $0 \le i < m$, 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 $\overline{\beta} := \beta' \pmod{p} \in \mathcal{A}$ is $f(X)$. Let $\mathcal{C}$ be the subalgebra $\mathbb{F}_p[\overline{\beta}]$ 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 $f(\overline{\beta})/g(\overline{\beta})$. 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 4.9 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. □

5 Finding Automorphisms of Algebras via Kummer Extensions

In this section we complete the proof of our main Theorem, 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 subexponential time. 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 $\mathcal{D}$ 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 $r \neq$ 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 ring homomorphism $\phi : \mathcal{A}[\zeta_r][\sqrt[r]{c}] \to \mathcal{D}[\zeta_r]$ 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 $\zeta_r^0, \ldots, \zeta_r^{r-2}$ form a free basis of $\mathcal{D}[\zeta_r]$ as a $\mathcal{D}$-module, the subspaces $\zeta_r^i \mathcal{D}$ of $\mathcal{D}[\zeta_r]$ $(i = 0, \ldots, r-2)$ 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 fact $\dim_k \mathcal{A}[\zeta_r][\sqrt[r]{c}] = (r-1)\dim_k \mathcal{A}[\zeta_r][\sqrt[r]{c}]_{\Delta_r}$ means that ϕ is injective if and only if its restriction to $\mathcal{A}[\zeta_r][\sqrt[r]{c}]_{\Delta_r}$ is.

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.

Firstly, we claim that for all $1 \leq i \leq (r-1)$, $J \neq J^{\sigma^i}$. 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 to $\dim_k J < \dim_k \mathcal{C} = r \dim_k \mathcal{A}$. Secondly, we claim that for some $i \in \{1, \ldots, r-1\}$, $J \cap J^{\sigma^i} \neq 0$. 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 $i \in \{1, \ldots, r-1\}$, for which $J \neq J^{\sigma^i}$ and $J \cap J^{\sigma^i} \neq 0$, whence by the method of Lemma 2.3 we get a zero divisor of J, thus finishing the proof. □

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

Proposition 5.2. *Given a commutative semisimple algebra $\mathcal{D}$ over a finite field k, its semiregular k-automorphism τ of prime order $r \neq$ char k, 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 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 $x^\tau = \zeta_r x$. 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 the $\mathcal{C}$ and the $\tau' := \tau|_{\mathcal{C}}$ as promised. On the other hand if $x \notin \mathcal{A}[\zeta_r]$ then we claim that we can find a zero divisor in $\mathcal{D}$, decompose $\mathcal{D}$ into a direct sum of orthogonal ideals and construct the $\mathcal{C}$ and the τ' in one of the ideals recursively.

Say $x \notin \mathcal{A}[\zeta_r]$, then since $x^{r^t} = 1_\mathcal{D} \in \mathcal{A}$ for some integer $t > 0$, we can choose a $y \in \{x, x^r, x^{r^2}, \ldots\}$ 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 2.2. 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 2.3, obtain a τ-invariant decomposition of $\mathcal{D}$ into the orthogonal ideals $I_1, \ldots, I_t$. For $1 \leq j \leq t$, we denote by ϕ_j the projection $\mathcal{D} \to I_j$. 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 $E \subseteq \{I_1, \ldots, I_t\}$ 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 $r \not| \frac{\dim_k I_j}{\dim_k \mathcal{A}}$. 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.

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 $x^\tau = \zeta_r x$, 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 $\phi : \mathcal{D} \to I$ will just be the homomorphism $d \mapsto e_I d$, thus τ commutes with ϕ. Consequently, $\mathcal{C}_{\tau'} = \phi^{-1}(\mathcal{C}'_\tau) = \phi^{-1}(\mathcal{A}'_\tau) \geq \mathcal{A}_\tau$. $\square$

5.2 Essential Part of the Tensor Power

Let $\mathcal{A}$ be a commutative semisimple algebra over a finite field k. Let $\mathcal{B}$ be its subalgebra such that $k \subseteq \mathcal{B}$ and $\mathcal{A}$ be a free module over $\mathcal{B}$ of rank m. If $char\ k \leq m^2$ then polynomial factorization can be done in deterministic time by Berlekamp's algorithm and consequently, all our results can be obtained easily. So we assume from now on that $char\ k > m^2$. But then we can also assume that $\mathcal{A}$ is a simple extension algebra of $\mathcal{B}$ and find a primitive element α by running an algorithmic version of Fact 1 (if this "fails" then it gives a zero divisor of $\mathcal{A}$). If $g(X) \in \mathcal{B}[X]$ is a minimal polynomial of α then we have that $\mathcal{A} = \mathcal{B}[X]/(g(X))$.

It was shown by Rónyai [Ró87] that, under GRH, a zero divisor in $\mathcal{A}$ can be found in time $poly((\dim_k \mathcal{A})^r, \log |k|)$ 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 Main Theorem. A key idea of Rónyai 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 $k[X_1, \ldots, X_r]$.

Functional interpretation of tensor powers: Let a commutative semisimple $\mathcal{A}$ be a simple extension algebra over $\mathcal{B} \supseteq k$ such that $\mathcal{A} = \mathcal{B}[X]/(g(X))$ and $g(X) \in \mathcal{B}[X]$ is a monic polynomial of degree m. Let $r \leq m$. We consider the r-th tensor power $\mathcal{A}^{\otimes_{\mathcal{B}} r}$ ($\mathcal{A}$ tensored with itself r times wrt $\mathcal{B}$). To define (and compute) the essential part of this tensor power it is convenient to interpret $\mathcal{A}$ as a collection of functions $V \to \overline{\mathcal{B}}$ that are expressible as a polynomial over $\mathcal{B}$ (called *$\mathcal{B}$-polynomial functions*), where $\overline{\mathcal{B}} := \overline{k} \otimes_k \mathcal{B}$ is the *algebraic closure* of $\mathcal{B}$ and $V \subset \overline{\mathcal{B}}$ is a set of roots of $g(X)$. If $\mathcal{B}$ is not a field then there are various possibilities for V and we need one with $\prod_{v \in V}(X - v) = g(X)$. Such a V clearly exists by the definition of the algebraic closure. This *functional interpretation* of $\mathcal{A}$ generalizes to $\mathcal{A} \otimes_{\mathcal{B}} \mathcal{A}$, which now becomes the set of all $\mathcal{B}$-polynomial functions from the set $V \times V$ to $\overline{\mathcal{B}}$ and finally $\mathcal{A}^{\otimes_{\mathcal{B}} r}$ is the set of all $\mathcal{B}$-polynomial functions from the set V^r to $\overline{\mathcal{B}}$. Note that in this interpretation a rank 1 tensor element $h_1 \otimes \cdots \otimes h_r$ in $\mathcal{A}^{\otimes_{\mathcal{B}} r}$ corresponds to the function $V^r \to \overline{\mathcal{B}}$ that maps $(v_1, \ldots, v_r) \mapsto h_1(v_1) \cdots h_r(v_r)$.

Essential part of tensor powers: The *essential part* $\widetilde{\mathcal{A}^{\otimes_{\mathcal{B}} r}}$ of $\mathcal{A}^{\otimes_{\mathcal{B}} r}$ is the subset of functions that vanish on all the r-tuples $(v_1, \ldots, v_r)$ that have $v_i = v_j$ for some $i \neq j$. It can be seen that $\widetilde{\mathcal{A}^{\otimes_{\mathcal{B}} r}}$ is an ideal of $\mathcal{A}^{\otimes_{\mathcal{B}} r}$. We show below that given a basis of $\mathcal{A}$ over $\mathcal{B}$ we can directly compute a basis for $\widetilde{\mathcal{A}^{\otimes_{\mathcal{B}} r}}$ over $\mathcal{B}$.

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

Proof. Consider embeddings μ_i of $\mathcal{A}$ into $\mathcal{A}^{\otimes_{\mathcal{B}} r}$ $(i = 1, \ldots, r)$ given as $\mu_i(a) = 1 \otimes \ldots \otimes 1 \otimes a \otimes 1 \otimes \ldots \otimes 1$ where a is in the i-th place. In the interpretation as functions, $\mu_i(\mathcal{A})$ correspond to the $\mathcal{B}$-polynomial functions on V^r which depend only on the ith element in the tuples. Observe that the set, for $1 \leq i < j \leq r$:

$$\Delta^r_{i,j} = \{b \in \mathcal{A}^{\otimes_{\mathcal{B}} r} \mid (\mu_i(a) - \mu_j(a))b = 0 \text{ for every } a \in \mathcal{A}\}$$

is the ideal of $\mathcal{A}^{\otimes_{\mathcal{B}} r}$ consisting of the $\mathcal{B}$-polynomial functions which are zero on every tuple $(v_1, \ldots, v_r)$ with $v_i \neq v_j$. Given a basis for $\mathcal{A}$, a basis for $\Delta^r_{i,j}$ can be computed by solving a system of linear equations in time (counting k-operations as unit time) polynomial in $\dim_k \mathcal{A}^{\otimes_{\mathcal{B}} r} = m^r \dim_k \mathcal{B}$. Finally, notice that $\widetilde{\mathcal{A}^{\otimes_{\mathcal{B}} r}}$ can be computed as well since it is the annihilator of $\sum_{1 \leq i < j \leq r} \Delta^r_{i,j}$. □

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 $h_1 \otimes \cdots \otimes h_r \mapsto h_{\pi(1)} \otimes \cdots \otimes h_{\pi(r)}$. This action is not semiregular on the tensor power algebra as it fixes the set I_0 of $\mathcal{B}$-polynomial functions on V^r that are zero on all the points $V^r \setminus \{(v, \ldots, v) | v \in V\}$, where I_0 can be seen to be an ideal of $\mathcal{A}^{\otimes_{\mathcal{B}} r}$. However, the ideal $\widetilde{\mathcal{A}^{\otimes_{\mathcal{B}} r}}$ is invariant under this action and on it S_r acts semiregularly.

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 $h \otimes 1_{\mathcal{A}} \otimes \cdots \otimes 1_{\mathcal{A}}$. Composing this embedding with the projection onto 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 $\widetilde{\mathcal{A}^{\otimes_{\mathcal{B}} r}}$.

Note that the ideal $\widetilde{\mathcal{A}^{\otimes_{\mathcal{B}} r}}$ is a free $\mathcal{B}$-module of rank $m \cdots (m - r + 1)$. Denoting the above embedding 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 $(1 \ldots r)$ as τ. This immediately gives us the following GRH-free version of the result of [Ró87]:

Theorem 5.4. *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 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 $\mathcal{C}$ of order r such that $\mathcal{C}_\tau \geq \mathcal{B}$.*

In the proof of Main Theorem 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 5.5. *Let $m := dim_{\mathcal{B}}\mathcal{A}$ and I be a nonzero ideal of $\widetilde{\mathcal{A} \otimes_{\mathcal{B}} \mathcal{A}}$. Let $\tau_1 : \mathcal{A} \to I$ 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 of elements $\{x \otimes 1 - 1 \otimes x | x \in \mathcal{A}\}$, see Lemma 5.3. 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 as I is an ideal of the essential part of the semisimple $\mathcal{A} \otimes_{\mathcal{B}} \mathcal{A}$, there is a natural projection $\phi : \mathcal{A} \otimes_{\mathcal{B}} \mathcal{A} \to I$. Then $\tau_1(\mathcal{A}) = \phi(\mathcal{A} \otimes_{\mathcal{B}} 1)$ and $\tau_2(\mathcal{A}) = \phi(1 \otimes_{\mathcal{B}} \mathcal{A})$. 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 $\tau_1(\mathcal{A}) = \tau_2(\mathcal{A})$. □

5.3 Proof of Main Theorem

We now prove the following slightly stronger version of Main Theorem.

Theorem 5.6. *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 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 may assume that *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 5.4 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: one on the pair $(\mathcal{A}, \mathcal{C})$ and the other on the pair $(\mathcal{C}_{\sigma_0}, \mathcal{B})$. 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 $\mathcal{C}_{\sigma_0}$ 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 4.8 two times to construct σ from $\sigma_0, \sigma_1, \sigma_2$. This finishes the even m case.

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 $\tau_1 : \mathcal{A} \to I$ mapping $a \mapsto e_I\phi_1(a)$ be the embedding of $\mathcal{A}$ into I. Look at the homomorphism $\tau_2 : \mathcal{A} \to I$ 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 5.5. If μ is not semiregular then we can find a zero divisor by Proposition 3.2 while if μ is semiregular then we can apply recursion to the pair $(\mathcal{A}_\mu, \mathcal{B})$, find an automorphism of $\mathcal{A}_\mu$ and finally extend it to a promised automorphism of $\mathcal{A}$ by Lemma 4.8.

So let us assume that σ is nontrivial, i.e. $I > I_\sigma = \tau_1(\mathcal{A})$, thus $\mathrm{rk}_{\tau_1(\mathcal{B})}I > m$. Then we define $\mathcal{B}'' := \tau_2(\mathcal{A})$ and apply recursion to the pair $(I, \mathcal{B}'')$. 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 4.7 and proceed with recursion. Thus, $I_{\sigma'}$ is σ-invariant and I_σ is σ'-invariant. Thus both σ and σ' can be viewed as automorphisms of $\tau_2(\mathcal{A})$ and $\tau_1(\mathcal{A})$ respectively. If both 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 statment of Lemma 5.5. 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 proceed with a recursion to the pair $(\mathcal{A}_\mu, \mathcal{B})$, getting a promised automorphism of $\mathcal{A}$ by the algorithm of Lemma 4.8.

To see the dominating term in the time complexity observe that in any recursive call on some pair, say $\mathcal{C}, \mathcal{D}$ with $d := \mathrm{rk}_{\mathcal{D}}\mathcal{C}$, if d is odd then we need to go to the tensor square of $\mathcal{C}$ wrt $\mathcal{D}$. Thus we need to then work in an algebra of rank d times the original rank. 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 $poly(m^{\log m})$ many $\mathcal{B}$-operations, which proves the time complexity as promised. □

To finish the proof of Main Theorem, apply the process described in the above Theorem 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 $e_Ik \cong k$. Thus after several iterations based on Theorem 5.6 we get the direct sum decomposition of $\mathcal{A}$

together with automorphisms as promised in Main Theorem.

6 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 subexponential 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 5.6, 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 [Ró90, 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 4) we know that if $\mathcal{C}_1, \ldots, \mathcal{C}_r$ 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 2.2. So we can assume $\mathcal{A} = \bigoplus_{i=1}^{r} M_m(\mathcal{C}_i) = M_m(\oplus_{i=1}^{r}\mathcal{C}_i) = M_m(\mathcal{C})$. 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 $M_m(\mathcal{C})$.

6.1 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 certain mild 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. If F is algebraically closed then every commutative semisimple subalgebra of $M_m(F)$ is in fact (upto a conjugation isomorphism) a subalgebra of the diagonal matrices.

Lemma 6.1. *Let $\mathcal{C}$ be a commutative semisimple algebra over a finite field k, 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 a nonzero $y \in M_m(\mathcal{C})$ 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 promised y as a permutation matrix.

Firstly, we can assume $\mathcal{C}$ to be a field because if $I_1, \ldots, I_c$ are the simple components of $\mathcal{C}$ then clearly the I_i's are all finite 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 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 $(y_1 + \ldots + y_r)$ is the promised y. So for the rest of the proof we assume that $\mathcal{C}$ is a finite field extension of k. Secondly, notice that the condition $yx^\sigma = xy$ is equivalent to the system of equations: $yx_1^\sigma = x_1 y, \ldots, yx_r^\sigma = x_r y$ for a $\mathcal{C}$-basis $x_1, \ldots, x_r$ 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 $\overline{\mathcal{C}}$ of $\mathcal{C}$. A solution over $\overline{\mathcal{C}}$ gives a matrix $y \in M_m(\overline{\mathcal{C}})$ such that $yx^\sigma = xy$ for every $x \in \overline{\mathcal{B}}$ where $\overline{\mathcal{B}} := \overline{\mathcal{C}} \otimes_{\mathcal{C}} \mathcal{B}$ and we extend σ $\overline{\mathcal{C}}$-linearly to an algebra automorphism of $\overline{\mathcal{B}}$. Because k was a finite field, $\overline{\mathcal{B}} \leq M_m(\overline{\mathcal{C}})$ is a commutative semisimple algebra over $\overline{\mathcal{C}}$. Similarly, $\overline{\mathcal{D}} := \overline{\mathcal{C}} \otimes_{\mathcal{C}} \mathcal{D}$ is a maximal commutative semisimple subalgebra of $M_m(\overline{\mathcal{C}})$, and is also a free $\overline{\mathcal{B}}$-module. By the former condition $\dim_{\overline{\mathcal{C}}} \overline{\mathcal{D}} = m$ and by the latter condition $r|m$. We will now focus on the instance of $(\overline{\mathcal{D}}, \overline{\mathcal{B}}, \overline{\mathcal{C}})$ and try to construct the promised y.

As $\overline{\mathcal{D}}$ is a sum of m copies of $\overline{\mathcal{C}}$, by an appropriate basis change we can make $\overline{\mathcal{D}}$ the algebra of all diagonal matrices in $M_m(\overline{\mathcal{C}})$. Also, as $\overline{\mathcal{D}}$ is a free $\overline{\mathcal{B}}$-module, a further basis change makes $\overline{\mathcal{B}}$ the algebra generated by the matrices $e_1, \ldots e_r$ where each e_j is a diagonal 0-1 matrix having m/r consecutive 1's. In that case the automorphism σ has a simple action, namely it permutes the matrices $\{e_1, \ldots, e_r\}$. 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 $x^\sigma = y^{-1} x y$ for every $x \in \overline{\mathcal{B}}$ by extending the equalities linearly to $\overline{\mathcal{B}}$. $\square$

In the second 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 6.2. *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* 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 $(1 + y + \ldots + y^{r-1})$ 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 $e_j = y^{-j} e y^j$ $(j = 1, \ldots, r)$ 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 $\alpha_0, \ldots, \alpha_{r-1} \in F$, 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 $L_1, \ldots, L_r$. $\square$

6.2 Proof of Application 1

In this subsection we give the proof of Application 1: given a noncommutative algebra $\mathcal{A}$ over a finite field k, one can unconditionally find zero divisors of $\mathcal{A}$ in deterministic subexponential time. 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. The basic idea in the algorithm then is to find a maximal commutative semisimple subalgebra $\mathcal{D} \leq \mathcal{A}$, find a $\mathcal{C}$-automorphism σ of $\mathcal{D}$, use it 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* $\mathcal{A}'$ over $\mathcal{C}$ we encounter have two generators x, y 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 $xy = -yx$.

Given the algebra $\mathcal{A}$ (with an unknown isomorphism to $M_m(\mathcal{C})$) in basis form over the finite field 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 5.6 we can, in deterministic $poly(m^{\log m}, \log|\mathcal{A}|)$ 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, so it is the latter case that we now assume. By Lemma 6.1, there certainly exists a $y \in \mathcal{A}$ such that $d^\sigma = y^{-1}dy$ for every $d \in \mathcal{D}$, so by picking a nonzero solution of the corresponding system of linear equations we either find a zero divisor of $\mathcal{A}$ or we find such a y. So suppose we find a y such that $d^\sigma = y^{-1}dy \neq d$ for every $d \in \mathcal{D} \setminus \mathcal{C}$.

We can efficiently obtain a multiple M of the multiplicative order of y, $ord(y)$, 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 factor $r|m$, as σ is a semiregular $\mathcal{C}$-automorphism of $\mathcal{D}$, σ is of order m, hence using M we can replace y and σ by an appropriate power such that $ord(y)$ is a power of r while $ord(\sigma) = r$. By this construction, conjugation by y is now a $\mathcal{C}$-automorphism σ of $\mathcal{D}$ of order r. Put $z := y^r$, thus $d = d^{\sigma^r} = z^{-1}dz$ for every $d \in \mathcal{D}$. Note that we can assume $z \neq 1$ as otherwise $(y-1)$ is a zero divisor of $\mathcal{A}$ by Lemma 6.2. Thus an appropriate power, say ζ_r, of z has order r. Consider the subalgebra $\mathcal{D}[z]$, it is commutative by the action of z on $\mathcal{D}$ 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 $\mathcal{D}[z]$ is a commutative semisimple algebra. By the maximality of $\mathcal{D}$ we deduce that $\mathcal{D}[z] = \mathcal{D}$, hence $z \in \mathcal{D}$ and $\zeta_r \in \mathcal{D}$. So by Lemma 4.5 we can find efficiently either a zero divisor in $\mathcal{D}$ or an $x \in \mathcal{D}^*$ such that $x^\sigma = \zeta_r x$. We assume the latter case and we replace x by an appropriate power so that $ord(x)$ is an r-power. Let $w := x^r$, as σ fixes w, it has to be in $\mathcal{C}$.

Let $\mathcal{A}' := \mathcal{C}[x, y]$, $\mathcal{D}_x := \mathcal{C}[x] \leq \mathcal{A}'$, $\mathcal{D}_y := \mathcal{C}[y] \leq \mathcal{A}'$ and $\mathcal{C}' := \mathcal{C}[w, z] \leq \mathcal{A}'$. 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 $\{x^i y^j | 1 \leq i, j \leq (r-1)\}$ 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 $\mathcal{D}_x$ as an automorphism of order r and that the conjugation by x acts on $\mathcal{D}_y$ as an automorphism of order r. We can assume that both these $\mathcal{C}'$-automorphisms are semiregular as otherwise we can find a zero divisor by Proposition 3.2. Thus both $\mathcal{D}_x$ and

$\mathcal{D}_y$ are free modules over $\mathcal{C}$ of rank r, furthermore assume $\mathcal{A}'$ to be a free $\mathcal{C}$-module (also free $\mathcal{C}'$-module) or else we find a zero divisor in $\mathcal{C}$ (or $\mathcal{C}'$) by Lemma 2.2.

We can assume that w, z generate a cyclic subgroup of $\mathcal{C}'$ otherwise by Lemma 2.1 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 $x'y = \zeta_r y x'$, thus conjugation by x' gives an automorphism of $\mathcal{D}_y$, whence $(x' - 1)$ is a zero divisor by Lemma 6.2. 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 $ord(w^j z) < r^t$. 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 $(yxy^{-1} = \zeta_r^{-1}x)$ repeatedly we get, $y'^r = (x^j y)^{r-2}(x^j y)(x^j y)$ $= (x^j y)^{r-3}(x^j y)(\zeta_r^{-j} x^{2j} y^2) = \cdots = \zeta_r^{-jr(r-1)/2} x^{rj} y^r = \zeta_r^{-jr(r-1)/2} w^j z$. 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 r case can be repeated except when $ord(-wz)$ does not fall, i.e. orders are such that $ord(wz) < ord(w) = ord(z) = ord(-wz)$. 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 $u := (\alpha y + \beta) \in \mathcal{D}_y$ and $x' := ux$. If $x' \in \mathcal{D}_y$ then $x \in u^{-1}\mathcal{D}_y = \mathcal{D}_y$ which is a contradiction. Thus, $x' \notin \mathcal{D}_y$, 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 Application 1 in all cases.

6.3 Further Results on Finding Zero Divisors in $M_m(\mathcal{C})$

In this part we briefly outline an alternative of the approach of Application 1. Formal statements and details of proofs will be subject of a subsequent paper.

Assume that $\mathcal{A} \cong M_m(\mathcal{C})$ for some commutative semisimple algebra $\mathcal{C}$ over the finite field k. As in the proof of Application 1, we use the method of [GI00] to find a maximal semisimple subalgebra $\mathcal{D}$ of $\mathcal{A}$. Note that $\mathcal{D}$ is a free module over $\mathcal{C}$ of rank m. Let r be a prime divisor of m. Then we can use the algorithm of Theorem 5.4 to find an automorphism of a *subalgebra* $\mathcal{B}$ of order r in time $poly(m^r, \log|\mathcal{A}|)$. The remaining part of the proof of Application 1 can be modified so that an automorphism of prime order of a subalgebra of $\mathcal{D}$ rather than one of the whole $\mathcal{D}$ can be used to find a zero divisor in $\mathcal{A}$ in polynomial time. This way we obtain a deterministic algorithm of complexity $poly(m^r, \log|\mathcal{A}|)$ for finding a zero divisor in an algebra $\mathcal{A}$ isomorphic to $M_m(\mathcal{C})$, where r is the smallest prime divisor of m.

Using a generalization [CIK97] of a method of [BR90] we can use the zero divisor obtained above to compute a subalgebra of $\mathcal{A}$ (in the broader sense, thus a subalgebra of a one-sided ideal of $\mathcal{A}$) isomorphic to $M_{m'}(\mathcal{C})$, where m' is a certain divisor of m. Iterating this method we ultimately find a zero divisor z of $\mathcal{A}$ which is equivalent to an elementary matrix (a matrix having just one nonzero entry) wrt an isomorphism $\mathcal{A} \cong M_m(\mathcal{C})$. Then the left ideal $\mathcal{A}z$ is isomorphic to the standard module for $M_m(\mathcal{C})$ (the module of column

vectors of length m over $\mathcal{C}$). Finding such a module is equivalent to constructing an explicit isomorphism with $M_m(\mathcal{C})$. The time complexity is $poly(m^r, \log|\mathcal{A}|)$, where r is the *largest* prime divisor of m. In particular, if $\mathcal{A} \cong M_{2^\ell}(\mathcal{C})$, our method computes such an isomorphism in deterministic *polynomial time*.

7 Special Finite Fields: Proof of Application 4

In this section we assume that $k = \mathbb{F}_p$ for a prime $p > 3$ 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 [Ró89a] to the following form:

Lemma 7.1. *Assume that $\mathcal{A}$ is a free module over its subalgebra $\mathcal{B}$ of rank d. Then in time $poly(\log|\mathcal{A}|, S)$ 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 $f_1, \ldots, f_n$ such that $f_i\mathcal{B} \cong k$ for all i. For an $i \in \{1, \ldots, n\}$, we can project the hypothesis to the f_i component, thus $\dim_k f_i\mathcal{A} = d$ and there are orthogonal primitive idempotents $e_{i,1}, \ldots, e_{i,d}$ of $\mathcal{A}$ such that $f_i\mathcal{A} = e_{i,1}\mathcal{A} \oplus \cdots \oplus e_{i,d}\mathcal{A}$. As f_i is an identity element of $f_i\mathcal{A}$ we further get that $f_i = (e_{i,1} + \cdots + e_{i,d})$.

Now pick an $y \in \mathcal{A} \setminus \mathcal{B}$. Suppose (for the sake of contradiction) for all $1 \leq i \leq n$ there is a single $y_i^* \in k$ that satisfies for all $1 \leq j \leq d$, $ye_{i,j} = y_i^* e_{i,j}$. 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 $i \in \{1, \ldots, n\}$ and distinct $j, j' \in \{1, \ldots, d\}$ 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 $r_1, \ldots, r_t$ 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 $\ell \in \{1, \ldots, t\}$ (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 $(y_1 + a)^{r_\ell} \neq (y_2 + a)^{r_\ell}$ implies 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 $\ell \in \{1, \ldots, t\}$.

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 $r_1 \neq r_2$ in which case say 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 a power of r order together with the other properties. □

For an integer m we denote by $\Phi_m(X)$ the mth cyclotomic polynomial in $k[X]$. Let $r_1, \ldots, r_t$ be the prime divisors of $(p-1)$. Then for a subset I of $\{1, \ldots, t\}$ 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 7.2. *Let $\mathcal{B} < \mathcal{A}$. Assume that we are also given a surjective homomorphism from $k[X]/(\Phi_{r_I}(X))$ onto $\mathcal{B}$ for some subset I of $\{1, \ldots, t\}$. Then in time 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 $I' \subseteq \{1, \ldots, t\}$.*

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 7.1. 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 $dim_{\mathcal{B}}\mathcal{A} = r$ 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 $z^2 = -x^4$. 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 could assume $\mathcal{A} = \mathcal{B}[x]$ to be a free $\mathcal{B}$-module with the free basis $\{1, x, \ldots, x^{d-1}\}$, as otherwise we can find a zero divisor in $\mathcal{B}$ by Lemma 2.2. 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 $\dim_{\mathcal{B}} \mathcal{A} < r$, ϕ cannot be semiregular, so we get a zero divisor by Proposition 3.2 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. □

Using Lemma 7.2 as an induction tool, we obtain the following.

Theorem 7.3. *Let $f(X)$ be a polynomial of degree n which completely splits into linear factors over $\mathbb{F}_p$. Let $r_1 < \ldots < r_t$ be the prime factors of $(p-1)$. Then by a deterministic algorithm of running time poly$(r_t, n, \log p)$, we can either find a nontrivial 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 $\{1, \ldots, t\}$ 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 2.3. 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 Application 4.

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.*

[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.*

[Be67] E. R. Berlekamp, Factoring polynomials over finite fields, *Bell System Technical Journal 46(1967), 1853-1859.*

[Cam83] P. Camion, A deterministic algorithm for factorizing polynomials of $\mathbb{F}_q[x]$, *Ann. Discr. Math., 17, (1983), 149-157.*

[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.*

[CIK97] A. Chistov, G. Ivanyos, M. Karpinski, Polynomial time algorithms for modules over finite dimensional algebras, *Proc. ISSAC 1997, 68-74.*

[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.)*

[CZ81] D. G. Cantor, H. Zassenhaus, A new algorithm for factoring polynomials over finite fields, *Mathematics of Computation, 36(154), 1981, 587-592.*

[Ev89] S. A. Evdokimov, Factorization of a solvable polynomial over finite fields and the generalized Riemann Hypothesis, *Zapiski Nauchnyck Seminarov LOMI, 176(1989), 104-117.*

[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.*

[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.*

[G87] J. von zur Gathen, Factoring polynomials and primitive elements for special primes, *Theoretical Computer Science, 52, 1987, 77-89.*

[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.

[GI00] W. A. de Graaf, G. Ivanyos, Finding maximal tori and splitting elements in matrix algebras, *In: F. van Oysteayen, M. Saorin (eds), Interaction between Ring Theory and Representations of Algebras, Lecture Notes in Pure and Applied Mathematics 210, Marcel Dekker 2000, 95-105.*

[GS92] J. von zur Gathen, V. Shoup, Computing Frobenius maps and factoring polynomials, *Comput. Complexity, 2(1992), 187-224.*

[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.*

[Hu86] D. Husemöller, Elliptic curves; *Springer, 1986.*

[IKS08] G. Ivanyos, M. Karpinski, N. Saxena, Schemes for Deterministic Polynomial Factoring, *Preprint: CoRR abs/0804.1974, (2008).*

[KS98] E. Kaltofen, V. Shoup, Subquadratic-time factoring of polynomials over finite fields, *Math. Comp., 67(1998), 1179-1197.*

[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.*

[La80] S. Lang, Algebraic number theory, *Springer-Verlag, 1980.*

[L91] H. W. Lenstra, Finding isomorphisms between finite fields, *Mathematics of Computation 56(1991), 329-347.*

[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.*

[Moe77] R. T. Moenck, On the efficiency of algorithms for polynomial factoring, *Math. Comp., 31, (1977), 235-250.*

[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.*

[Rab80] M. O. Rabin, Probabilistic algorithms in finite fields, *SIAM J. Comput, 9 (1980), 273–280.*

[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*

[Ró89a] L. Rónyai, Factoring polynomials modulo special primes, *Combinatorica, 9, (1989), 199-206.*

[Ró90] L. Rónyai, Computing the structure of finite algebras, *Journal of Symbolic Computation 9, (1990) 355-373.*

[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.*

[Sch85] R. J. Schoof, Elliptic curves over finite fields and the computation of square roots mod p, *Mathematics of Computation 44 (1985), 483-494.*

[W05] C. van de Woestijne, Deterministic equation solving over finite fields, *Proc. ISSAC 2005, 348-353.*