

Looking beyond XTR

Wieb Bosma¹, James Hutton², Eric R. Verheul³

¹ Mathematisch Instituut, Universiteit Nijmegen
Postbus 9010, 6500 GL Nijmegen, The Netherlands

`bosma@sci.kun.nl`

² Thales e-Security,

149 Preston Road, Brighton, BN1 6BN, U.K.

`jamie.hutton@thales-esecurity.com`

³ PricewaterhouseCoopers, GRMS Crypto Group,
P.O. Box 85096, 3508 AB Utrecht, The Netherlands

`eric.verheul@nl.pwcglobal.com, pobox.com`

Abstract. XTR is a general method that can be applied to discrete logarithm based cryptosystems in extension fields of degree six, providing a compact representation of the elements involved. In this paper we present a precise formulation of the Brouwer-Pellikaan-Verheul conjecture, originally posed in [4], concerning the size of XTR-like representations of elements in extension fields of arbitrary degree. If true this conjecture would provide even more compact representations of elements than XTR in extension fields of degree thirty. We test the conjecture by experiment, showing that in fact it is unlikely that such a compact representation of elements can be achieved in extension fields of degree thirty.

1 Introduction

Many public key cryptosystems are based on the assumed intractability of the Discrete Logarithm (DL) problem: given a cyclic group $G = \langle g \rangle$ and $h \in G$ find $0 \leq x < \#G$ such that $h = g^x$.

Any cryptosystem based on the DL problem requires a large cyclic group G as a parameter of the system. We require that exponentiation is efficient in G but that the DL problem is believed to be hard.

The seminal example of DL-based cryptosystems is Diffie-Hellman key exchange (see [6]), a method that enables two parties (Alice and Bob) to establish a shared secret key by exchanging messages over an open channel. Alice generates a random key $2 \leq a < \#G$ and sends $A = g^a$ to Bob. Similarly Bob generates $2 \leq b < \#G$ and sends $B = g^b$ to Alice. Alice and Bob can now both determine the common secret key $S = A^b = B^a = g^{ab}$.

The basic and original version of Diffie-Hellman key exchange uses $G = \mathbb{F}_p^*$ where the prime p and a generator g of G are public parameters. There are other choices for the group G . For example Claus Schnorr proposed using a prime order subgroup of \mathbb{F}_p^* (see [21]). Alternatively one can use the group of points on certain elliptic curves.

In this paper we explore another choice: a carefully chosen subgroup G of prime order q of the multiplicative group of an extension field \mathbb{F}_{p^k} . We can represent elements of G by their minimal polynomials over a subfield of \mathbb{F}_{p^k} and thereby for certain values of k achieve a comparatively compact representation of the group elements involved. This is the idea behind LUC ($k = 2$) and XTR ($k = 6$); see Section 2.

In Section 3 we refer to a conjecture implicitly posed by Brouwer, Pellikaan and Verheul in [4] (the ‘BPV’ conjecture) concerning the size of minimal polynomial representations of elements in field extensions of arbitrary degree. In Sections 4 and 5 we prove some general results concerning the coefficients of minimal polynomials and develop precise formulations of the BPV conjecture.

Our main objective was to investigate the possibility of obtaining a more compact representation of elements than XTR for values of k larger than 6. We used a MAGMA program (described in Section 6) to conduct our investigations. Since the BPV conjecture (if true) would provide a more compact representation than XTR in field extensions of degree thirty we considered this to be the most interesting case. However we also investigated intermediate values of k and discovered (rather to our surprise) some cases that support the conjecture, although these cases do not provide a more compact representation than XTR.

In Section 7 we present the experimental results of our investigations. We show that if the conjectured relations exist in the degree 30 case then they are most likely too complicated to be of practical value.

2 Representing elements by their minimal polynomials

A standard method for representing elements of an extension field \mathbb{F}_{p^k} is as vectors over a subfield \mathbb{F}_{p^d} . The usual way of achieving such a representation is to use the fact that $\mathbb{F}_{p^k} \cong \mathbb{F}_{p^d}[X]/P(X)$ where P is an irreducible polynomial of degree k/d over \mathbb{F}_{p^d} . Elements of \mathbb{F}_{p^k} are represented by residue classes modulo P and these classes can in turn be represented by the polynomials over \mathbb{F}_{p^d} of degree less than k/d . The coefficients of these polynomials enable us to express the field elements as vectors over \mathbb{F}_{p^d} of length k/d ; therefore we generally require $k \log p$ bits to represent an element.

A well-known alternative method is to represent $\alpha \in \mathbb{F}_{p^k}$ by its minimal polynomial over a subfield \mathbb{F}_{p^d} . This is the unique monic irreducible polynomial F over \mathbb{F}_{p^d} such that $F(\alpha) = 0$. We always have $\deg(F) \leq k/d$.

Note that when $\deg(F) = k/d$ and $d < k$ then the k/d non-trivial coefficients of the minimal polynomial do not determine the element uniquely: if $\alpha_0 \in \mathbb{F}_{p^k}$ is a root of F then so are $\alpha_1 = \alpha_0^{p^d}$, $\alpha_2 = \alpha_0^{p^{2d}}$, ..., $\alpha_{k/d-1} = \alpha_0^{p^{k-d}}$. The α_i are the conjugates of α_0 over \mathbb{F}_{p^d} , and these elements are all represented by the same minimal polynomial over \mathbb{F}_{p^d} .

The minimal polynomial over \mathbb{F}_{p^d} of an element of \mathbb{F}_{p^k} will have degree k/d unless the element is contained in a subfield \mathbb{F}_{p^e} where $\mathbb{F}_{p^d} \subseteq \mathbb{F}_{p^e} \subset \mathbb{F}_{p^k}$. Thus at first sight it appears that for most elements of \mathbb{F}_{p^k} we would require $k \log p$ bits to specify the k/d non-trivial coefficients of the minimal polynomial over \mathbb{F}_{p^d} ,

and that therefore we need the same number of bits to represent elements as in the representation using residue classes discussed above. However in certain cases there exist relationships between the coefficients of the minimal polynomials that enable us to reduce the number of coefficients that are required and thereby make the representation more compact. This idea is used in both LUC and XTR; we describe these methods in the two examples at the end of this section.

We now introduce the subgroups of field extensions in which we work.

Definition 1. In a field \mathbb{F}_{p^k} we call a subgroup of prime order q with $q \mid \Phi_k(p)$ and $q \nmid k$ a *cyclotomic subgroup* and denote it by $G_{q,p,k}$. (Here $\Phi_k(p)$ denotes the k -th cyclotomic polynomial evaluated in p , see [8] and [15].)

We call the group of all elements of order dividing $\Phi_k(p)$ the (p, k) -*cyclotomic group* and denote it by $G_{p,k}$.

The original Diffie-Hellman protocol uses the $(p, 1)$ -cyclotomic group $G_{p,1}$, while Schnorr's variant is based in a cyclotomic subgroup $G_{q,p,1}$. LUC uses a cyclotomic subgroup $G_{q,p,2}$ and XTR uses $G_{q,p,6}$, as we next explain.

Example 1. The LUC system uses minimal polynomials to represent elements of a cyclotomic subgroup $G_{q,p,2}$ of $\mathbb{F}_{p^2}^*$. The minimal polynomial over \mathbb{F}_p of an element $h \in G_{q,p,2} \setminus \{1\}$ is

$$P_h = (X - h)(X - h^p) = X^2 - (h + h^p)X + h^{p+1} = X^2 - \text{Tr}_p(h)X + 1$$

where $\text{Tr}_p(h) \in \mathbb{F}_p$ denotes the trace of h over \mathbb{F}_p . Hence h can be represented by the polynomial $P_h \in \mathbb{F}_p[X]$ and this polynomial is completely determined by the value of $\text{Tr}_p(h)$. Thus only $\log p$ bits are required to represent elements of $G_{q,p,2}$ by their minimal polynomials, compared to the $2 \log p$ bits that would be required using a standard representation.

As already observed P_h does not determine h uniquely but determines both h and h^p , the conjugate of h over \mathbb{F}_p .

LUCDIF is a variant of Diffie-Hellman key exchange obtained by applying LUC to the conventional system described in the Introduction. In the LUCDIF variant Alice sends Bob $\text{Tr}_p(g^a)$ instead of g^a . Using the standard method for solving a quadratic equation Bob solves $X^2 - \text{Tr}_p(g^a)X + 1 = 0$ obtaining the solutions g^a and its conjugate g^{ap} . Bob can now use these solutions and his secret exponent b to calculate $(g^a)^b + (g^{ap})^b = g^{ab} + g^{abp} = \text{Tr}_p(g^{ab})$. Alice uses the same method to calculate the shared secret key $\text{Tr}_p(g^{ab})$ from the value $\text{Tr}_p(g^b)$ received from Bob.

The elements $\text{Tr}_p(g^a)$ and $\text{Tr}_p(g^b)$ that are communicated over the open channel are in \mathbb{F}_p and hence of length $\log p$ bits. This is half the size of the elements g^a and g^b that are exchanged in conventional Diffie-Hellman key exchange using the standard representation discussed at the beginning of this section.

Another benefit of LUCDIF is that the calculations that each party must perform are significantly quicker than in the conventional system. These calculations use so-called Lucas recurrent sequences. For full details the reader should consult [2], [22] (where the name 'LUCDIF' was proposed), [17], [16], [19] and [14].

Of course it is essential that the benefits achieved by applying LUC do not compromise the security of the system. In fact it is easily shown that breaking the LUCDIF variant is equivalent to breaking the conventional system.

Example 2. The XTR system represents elements of a cyclotomic subgroup $G_{q,p,6}$ by their minimal polynomials over \mathbb{F}_{p^2} . (This subgroup is called an XTR group in the XTR literature.) The minimal polynomial over \mathbb{F}_{p^2} of a non-identity element $h \in G_{q,p,6}$ is

$$\begin{aligned} P_h &= (X - h)(X - h^{p^2})(X - h^{p^4}) \\ &= X^3 - (h + h^{p^2} + h^{p^4})X^2 + (h^{p^2+1} + h^{p^4+1} + h^{p^4+p^2})X - h^{p^4+p^2+1} \\ &= X^3 - \text{Tr}_{p^2}(h)X^2 + (h^{p^2+1} + h^{p^4+1} + h^{p^4+p^2})X - h^{p^4+p^2+1} \end{aligned}$$

where Tr_{p^2} denotes the trace over \mathbb{F}_{p^2} .

Since $q \mid \Phi_6(p) = p^2 - p + 1$ and $p^2 - p + 1 \mid p^4 + p^2 + 1$ we know that the constant term of P_h is -1 . Furthermore the congruences $p^2 + 1 \equiv p$, $p^4 + 1 \equiv p^5$ and $p^4 + p^2 \equiv p^3$ modulo $p^2 - p + 1$ imply that the coefficient $h^{p^2+1} + h^{p^4+1} + h^{p^4+p^2}$ is equal to $\text{Tr}_{p^2}(h)^p$. Thus

$$P_h = X^3 - \text{Tr}_{p^2}(h)X^2 + \text{Tr}_{p^2}(h)^p X - 1,$$

which is completely determined by the value of $\text{Tr}_{p^2}(h) \in \mathbb{F}_{p^2}$. It follows that only $2 \log p$ bits are required to represent elements of $G_{q,p,6}$ by their minimal polynomials, which compares very favourably to the $6 \log p$ bits that would be required using a standard representation.

Clearly in order to apply XTR to a DL-based cryptosystem it is necessary to be able to perform certain computations using traces of elements of the XTR group. For example in Diffie-Hellman key exchange we must be able to compute $\text{Tr}_{p^2}(g^{xy})$ given $\text{Tr}_{p^2}(g^x)$ and y . Efficient methods for performing the calculations required for XTR variants of cryptosystems such as Diffie-Hellman key exchange and DSA have been developed by Lenstra and Verheul (see [10], [11], [12], [13]) and Lenstra and Stam (see [23]). As with LUC these methods are computationally more efficient than the corresponding calculations performed in $G_{q,p,6}$ without using traces.

We conclude this section by discussing some security issues.

The most effective known methods of (passive) attack against DL-systems are based on the Birthday Paradox or use of the Number Field Sieve. Birthday Paradox based algorithms (such as Pollard's rho algorithm [20]) have expected running times of order \sqrt{q} elementary operations in G , where q is the largest prime factor of the order of G . The Discrete Logarithm variant of the Number Field Sieve has a heuristic expected asymptotic running time of $L[p, 1/3, 1.923 + o(1)]$ (see [1] and [9]).

The security of the original Diffie-Hellman system, which uses $G_{p,1}$, depends not only on the size of p but also on that of the largest prime factor of $p - 1$; for adequate security this factor should have at least 160 bits. To resist Number Field Sieve attacks p should be at least 1024-bit.

For systems (like LUC, XTR) employing cyclotomic subgroups $G_{q,p,k}$ of \mathbb{F}_{p^k} the same requirements on the size of q apply: q should have at least 160 bits to be secure against Birthday Paradox attacks. The following lemma (see also [9, Lemma 2.4], as corrected by Minghua Qu) shows that the condition $q \nmid k$ ensures that every non-identity element of a cyclotomic subgroup $G_{q,p,k}$ lies outside every proper subfield of \mathbb{F}_{p^k} , and hence that $G_{q,p,k}$ is as secure against Number Field Sieve attacks as \mathbb{F}_{p^k} itself. This means that p should be chosen in such a way that $k \cdot \log p > 1024$. Hence for LUC p of at least 512 bits is recommended, and for XTR of at least 171 bits.

Lemma 1. *If $h \in G_{q,p,k} \setminus \{1\}$ then $h \notin \mathbb{F}_{p^d}$ for proper divisors d of k .*

Proof. Since $q \nmid k$ we have $\gcd(X^k - 1, kX^{k-1}) = 1$ in $\mathbb{F}_q[X]$ and thus $X^k - 1$ has no repeated roots in the algebraic closure of \mathbb{F}_q . As $X^k - 1 = \prod_{e|k} \Phi_e(X)$ and $\Phi_k(p) \equiv 0 \pmod q$ we see that $\Phi_e(p) \not\equiv 0 \pmod q$ for $e \mid k$, $e < k$. But for any proper divisor d of k we have

$$X^d - 1 = \prod_{e|d} \Phi_e(X) \mid \prod_{\substack{e|k \\ e < k}} \Phi_e(X),$$

so $p^d - 1 \not\equiv 0 \pmod q$. Thus the order of $\mathbb{F}_{p^d}^*$ is not a multiple of the order q of h .

3 Do more compact representations than XTR exist?

By representing elements of cyclotomic subgroups by their minimal polynomials over a subfield, LUC and XTR reduce the number of required bits per element by a factor 2 and 3 respectively. A natural question arises: can we do any better?

Both LUC and XTR provide evidence for the BPV conjecture mentioned in the Introduction, which can be informally stated as follows, using Euler's totient function ϕ :

Elements of $G_{q,p,k}$ can be represented with $\phi(k) \log p$ bits using minimal polynomials over some subfield of \mathbb{F}_{p^k} .

If the BPV conjecture were true the best size reduction (compared to a standard representation) is achieved when the ratio $k/\phi(k)$ is large. This happens when k is the product of distinct primes. LUC and XTR are the simplest such cases and the next value of k of interest would be $k = 2 \cdot 3 \cdot 5 = 30$. We shall investigate this case in Section 7.

We now present two more examples that provide further evidence in support of the BPV conjecture.

Example 3. Let k be a prime and let $h \in G_{q,p,k}$, with $h \neq 1$. The minimal polynomial P_h of h over \mathbb{F}_p has degree k and constant term equal to 1 if $k = 2$ and -1 otherwise (see Theorem 1). Therefore P_h is completely determined by the $k - 1$ coefficients of X, X^2, \dots, X^{k-1} . Since these coefficients are elements

of \mathbb{F}_p and $\phi(k) = k - 1$ it follows that elements of $G_{q,p,k}$ can be represented by $\phi(k) \log p$ bits, in support of the BPV conjecture.

Note that one can base generalisations of LUC on this example. In fact such a variant was published by G. Gong and L. Harn for $k = 3$ (see [7]). Here recurrent Lucas sequences similar to those used in LUC are employed. However this system has an ‘improvement factor’ $k/\phi(k)$ of just $3/2$.

Example 4. Let $k = 6$, so that the extension field has the same degree as in XTR (Example 2). In XTR we considered the minimal polynomial over \mathbb{F}_{p^2} of $h \in G_{q,p,6}$, $h \neq 1$. We now consider the minimal polynomial

$$P_h = \prod_{0 \leq i \leq 5} (X - h^{p^i}) = X^6 + a_5 X^5 + a_4 X^4 + a_3 X^3 + a_2 X^2 + a_1 X + a_0$$

of h over \mathbb{F}_p . The constant term $a_0 = 1$ since the order q of h divides $\Phi_6(p)$ which in turn divides $1 + p + p^2 + p^3 + p^4 + p^5$. Using $p^3 \equiv -1 \pmod{q}$ it is easily shown that $a_1 = a_5$ and $a_2 = a_4$ (cf. Corollary 1). We note that the value of the first elementary symmetric polynomial in the conjugates of h is $-a_5$ and the value of the second elementary symmetric polynomial in the conjugates of h is a_4 . Furthermore one can write a_3 (which is minus the value of the third elementary symmetric polynomial in the conjugates of h) as a symmetric polynomial of degree 2 in the conjugates of h . By the Fundamental Theorem of Symmetric Polynomials [18, Theorem 4.31] it follows that it is possible to write a_3 as a polynomial in a_5 and a_4 . In fact we have $a_3 = -a_5^2 + 2a_4 + 2a_5 - 2$, a relationship first noted in [4]. It follows that P_h is completely determined by a_5 and a_4 so that h can be represented by two elements of \mathbb{F}_p , that is by $\phi(6) \log p$ bits, in support of the conjecture.

The four examples that we have considered so far have demonstrated relationships that can hold between coefficients of minimal polynomials of elements of a cyclotomic subgroup $G_{q,p,k}$. In the next section we shall prove some general results concerning these relationships. We shall also formulate a weaker version of the BPV conjecture (Conjecture 3) that is more amenable to verification.

4 Coefficients of the minimal polynomials

We begin with the following theorem.

Theorem 1 *Let h be a generator of a cyclotomic subgroup $G_{q,p,k}$, where p is odd and $k \geq 2$. Let $X^{k/d} + a_{k/d-1}X^{k/d-1} + \dots + a_1X + a_0$ be the minimal polynomial of h over \mathbb{F}_{p^d} , for some d dividing k , with $d < k$. Then $a_0 = (-1)^{k/d}$, and if $k = 2\ell$ is even, $a_i = (-1)^{k/d} a_{k/d-i}^{p^\ell}$, for $i = 1, \dots, k/d - 1$.*

Proof. Write $h_j = h^{p^{dj}}$ for $j = 0, \dots, k/d - 1$. Then

$$X^{k/d} + a_{k/d-1}X^{k/d-1} + \dots + a_1X + a_0 = \prod_{j=0}^{k/d-1} (X - h_j),$$

and comparing coefficients we see that $a_i = (-1)^{k/d-i} \sigma_{k/d-i}(h_0, \dots, h_{k/d-1})$ for $i = 0, \dots, k/d - 1$, where $\sigma_n(h_0, \dots, h_{k/d-1})$ is the n -th elementary symmetric polynomial in the conjugates h_j of h . In particular

$$\begin{aligned} a_0 &= (-1)^{k/d} \sigma_{k/d}(h_0, \dots, h_{k/d-1}) \\ &= (-1)^{k/d} h_0 \cdots h_{k/d-1} = (-1)^{k/d} h^{1+p^d+p^{2d}+\cdots+p^{k-d}}. \end{aligned}$$

But $1 + p^d + p^{2d} + \cdots + p^{k-d} = (p^k - 1)/(p^d - 1)$ which is divisible by $\Phi_k(p)$ and hence by q , the order of h . Therefore $a_0 = (-1)^{k/d}$.

If $k = 2\ell$ is even then $p^k - 1 = (p^\ell - 1)(p^\ell + 1)$. Since the order q of h divides $p^k - 1$ but not $p^\ell - 1$ we have $p^\ell \equiv -1 \pmod{q}$ and therefore $h_j^{-1} = h_j^{p^\ell}$ for $j = 0, \dots, k/d - 1$. Furthermore, since $h_0 \cdot h_1 \cdots h_{k/d-1} = 1$ we have $\sigma_{k/d-i}(h_0, \dots, h_{k/d-1}) = \sigma_i(h_0^{-1}, \dots, h_{k/d-1}^{-1})$ for $i = 1, \dots, k/d - 1$. It follows that for $i = 1, \dots, k/d - 1$

$$\begin{aligned} \sigma_{k/d-i}(h_0, \dots, h_{k/d-1}) &= \sigma_i(h_0^{-1}, \dots, h_{k/d-1}^{-1}) = \sigma_i(h_0^{p^\ell}, \dots, h_{k/d-1}^{p^\ell}) \\ &= \sigma_i(h_0, \dots, h_{k/d-1})^{p^\ell} \quad (\text{characteristic } p) \\ &= ((-1)^i a_{k/d-i})^{p^\ell} = (-1)^i a_{k/d-i}^{p^\ell}. \end{aligned}$$

Therefore, as required, for $i = 1, \dots, k/d - 1$:

$$a_i = (-1)^{k/d-i} \sigma_{k/d-i}(h_0, \dots, h_{k/d-1}) = (-1)^{k/d} a_{k/d-i}^{p^\ell}.$$

Corollary 1 *If k is even and d divides $k/2$ then the minimal polynomial over \mathbb{F}_{p^d} of a generator of $G_{q,p,k}$ is palindromic: $a_i = a_{k/d-i}$ for $i = 0, \dots, k/d$.*

Proof. Write $\ell = k/2$. Elements of \mathbb{F}_{p^d} are invariant under p^ℓ -th powering since d divides ℓ . Hence, by the previous theorem, $a_i = (-1)^{k/d} a_{k/d-i}$ for $i = 0, \dots, k/d$. Since k/d is even the result follows.

Proposition 1 *Let $k = de$, with $e > 1$. Then for any element h of $G_{q,p,k}$ the minimal polynomial P_h over \mathbb{F}_{p^d} can be represented using the following number of elements of \mathbb{F}_{p^d} :*

- $e - 1$, if de is odd;
- $\frac{e-1}{2}$, if d is even and e is odd;
- $\frac{e}{2}$ if e is even.

Proof. We represent elements of $G_{q,p,k}$ by their minimal polynomials over the subfield of degree d . The constant coefficient is ± 1 , so $e - 1$ elements of \mathbb{F}_{p^d} suffice to represent elements of $G_{q,p,k}$. This covers the first case.

In the second and third cases k is even and by Theorem 1 only half of the remaining $e - 1$ coefficients are required. More precisely if e is odd we need $(e - 1)/2$ coefficients and if e is even we require $e/2$ coefficients.

Note that, unfortunately, this result cannot be used recursively since the coefficients a_i are not (in general) in a cyclotomic subgroup of \mathbb{F}_{p^d} .

Proposition 1 leaves a choice for d and e if k is composite, and some choices will offer better improvement factors than others. If k is even but not a power of two then Proposition 1 indicates that a good choice for e is the smallest odd divisor of k greater than 1. For example when k is divisible by 6 we can choose $e = 3$ and $d = k/3$ and thereby achieve an improvement ratio of 3. In the following example, which generalises Example 4, we show that the same improvement ratio can be achieved by taking $e = 6$.

Example 5. Let k be of the form $6d$, let $r = p^d$ and consider P_h , the minimal polynomial over \mathbb{F}_r of $h \in G_{q,p,k} \setminus \{1\}$:

$$P_h = \prod_{0 \leq i \leq 5} (X - h^{r^i}) = X^6 + a_5 X^5 + a_4 X^4 + a_3 X^3 + a_2 X^2 + a_1 X + a_0$$

where $a_i \in \mathbb{F}_r$.

The order of h is q , which divides $\Phi_{6d}(p)$. It is well-known (see [8]) that $\Phi_{6d}(X) \mid \Phi_6(X^d)$ so q divides $\Phi_6(r) = r^2 - r + 1$.

Arguing as in Example 2 we have

$$(X - h)(X - h^{r^2})(X - h^{r^4}) = X^3 - tX^2 + t^r X - 1,$$

where $t = \text{Tr}_{r,2}(h) \in \mathbb{F}_{r,2}$ is the trace of h over $\mathbb{F}_{r,2}$. From this we have

$$(X - h^r)(X - h^{r^3})(X - h^{r^5}) = X^3 - t^r X^2 + t^{r^2} X - 1.$$

Since $t^{r^2} = t$ it follows that

$$P_h = (X^3 - tX^2 + t^r X - 1)(X^3 - t^r X^2 + tX - 1),$$

and we see that not only is P_h palindromic, as Corollary 1 implies, with $a_0 = 1$, $a_1 = a_5 = -t - t^r$ and $a_2 = a_4 = t + t^r + t^{1+r}$, but also that

$$\begin{aligned} a_3 &= -2 - t^2 - t^{2r} \\ &= -(-t - t^r)^2 + 2(t + t^r + t^{1+r}) + 2(-t - t^r) - 2 \\ &= -a_5^2 + 2a_4 + 2a_5 - 2. \end{aligned}$$

This means that we only need a_5 and a_4 to specify P_h .

d	e	S	ratio: de/S
<i>odd</i>	<i>odd</i>	$d \cdot (e - 1)$	$e/(e - 1)$
<i>even</i>	<i>odd</i>	$d \cdot \frac{e-1}{2}$	$2e/(e - 1)$
<i>even</i>	<i>even</i>	$d \cdot \frac{e}{2}$	2
<i>odd</i>	<i>even</i>	$d \cdot \frac{e}{2}$	2
<i>any</i>	6	$d \cdot 2$	3

Table 1

The table summarises the results of Proposition 1 and Example 5 concerning the number of words S of size $\log p$ that suffices to represent the minimal polynomials of elements of $G_{q,p,de}$, and the improvement ratio de/S . Note that S is an upper bound; fewer words may do.

5 The Conjectures

We now work towards a more precise formulation of the BPV conjecture. Consider some $k = de$ with $e > 1$. Let h be an element of the (p, k) -cyclotomic group $G_{p,k}$ that is not contained in any proper subfield of \mathbb{F}_{p^k} and let $P_h^{(d)} = X^e + a_{e-1}X^{e-1} + \dots + a_1X + a_0$ be the minimal polynomial of h over \mathbb{F}_{p^d} . Note that a_{e-j} corresponds, up to sign, to the j -th elementary symmetric polynomial evaluated in the e conjugates of h over \mathbb{F}_{p^d} . By Theorem 1 we have $a_0 = (-1)^e$.

For $1 \leq i \leq e-1$ let $A_i = \{a_{e-1}, \dots, a_{e-i}\}$. We let u_d denote the smallest integer with the property that the set A_{e-1} of all non-trivial coefficients of $P_h^{(d)}$ can be recovered from A_{u_d} . Thus all coefficients of $P_h^{(d)}$ can be recovered from the first u_d elementary symmetric polynomials in the conjugates of h over \mathbb{F}_{p^d} but *not* from the first $u_d - 1$ polynomials.

We must address the question of what we mean by ‘recovering’ A_{e-1} from a subset A_i . Note that we should not simply state that all a_j can be expressed as polynomials in the elements of A_i , since the coefficients come from a finite field in which many relations will exist. It seems one requires the existence of such an expression *independent of p* , although perhaps dependent on d and e . However this is still not entirely satisfactory: the second part of Theorem 1 states that we can recover, for example, a_1 from $a_{k/d-1}$ using conjugates, i.e. in a manner which *does* depend on p . This means that our ‘recovery’ notion for $d > 1$ should imply the existence of polynomials with integer coefficients and degree independent of p that, when evaluated in the d conjugates over \mathbb{F}_p of the elements of A_i , will yield the other coefficients.

We shall introduce multivariate polynomials in indeterminates X_j , and evaluate the polynomials at the elements of some coefficient set A_i . It will be convenient to define the *weighted degree* of a monomial $X_1^{e_1} \dots X_n^{e_n}$ in $\mathbb{Z}[X_1, \dots, X_n]$ to be $\sum_{j=1}^n j \cdot e_j$ and the weighted degree of a polynomial P as the maximum of the weighted degrees of the monomials that appear in P (with non-zero coefficient). Note that X_j has weighted degree j in P . The motivation for this definition is that we shall evaluate X_j in a_{e-j} , which is symmetric of degree j in the conjugates of h over \mathbb{F}_{p^d} .

Observe that $G_{p,k}$ is asymptotically of size $p^{\phi(k)}$. Therefore in order to represent the whole of $G_{p,k}$ by the minimal polynomials of its elements over \mathbb{F}_{p^d} we must have $d \cdot u_d \geq \phi(k)$. Thus, for given values of k and d , we have an information-theoretic lower bound of $\lceil \phi(k)/d \rceil$ on the value of u_d . The conjecture states that in fact u_d is always *equal* to this lower bound.

We now come to our first formulation of the BPV conjecture.

Conjecture 1 ((d, e)-BPV) *Let $k = de$, with $e > 1$. Let u_d be the least value of u for which $Q_j \in \mathbb{Z}[X_1^{(0)}, \dots, X_1^{(d-1)}, X_2^{(0)}, \dots, X_2^{(d-1)}, \dots, X_u^{(0)}, \dots, X_u^{(d-1)}]$ exist, for $1 \leq j \leq e - u - 1$, such that for every prime p and every element $h \in G_{p,k}$ that is not contained in a proper subfield of \mathbb{F}_{p^k} , the coefficient a_j of $P_h^{(d)}$ is given by*

$$a_j = \bar{Q}_j(a_{e-1}, a_{e-1}^p, \dots, a_{e-1}^{p^{d-1}}, a_{e-2}, a_{e-2}^p, \dots, a_{e-2}^{p^{d-1}}, \dots, a_{e-u}, a_{e-u}^p, \dots, a_{e-u}^{p^{d-1}}),$$

for $1 \leq j \leq e - u - 1$, where \bar{Q}_j denotes Q_j with coefficients taken modulo p . Then $u_d = \lceil \phi(de)/d \rceil$.

Motivated by Example 4 we also formulate a strong form of the conjecture, including a bound on the (weighted) degree of the polynomials involved.

Conjecture 2 (strong (d, e) -BPV) *Let $k = de$, with $e > 1$. Let u_d^* be the smallest integer for which there exist polynomials Q_j as in Conjecture (d, e) -BPV with the additional requirement that the polynomials Q_j are of weighted degree at most u , where the weighted degree of $X_k^{(i)}$ is k (for $1 \leq k \leq u$ and $0 \leq i \leq d - 1$). Then $u_d^* = \lceil \phi(de)/d \rceil$.*

The main conjecture, which was stated informally in Section 3, can now be made precise. For $k > 1$ we define $\text{red}(k) = \min\{d \cdot u_d\}$, where the minimum is taken over all proper divisors d of k .

Conjecture 3 (k -BPV) *Let $k > 1$ be an integer. There exists a proper divisor d of k such that d divides $\phi(k)$ and for which $(d, k/d)$ -BPV holds. Therefore $\text{red}(k) = \phi(k)$.*

Conjecture 3 applies to all elements of $G_{p,k}$ that are not contained in a proper subfield of \mathbb{F}_{p^k} . Therefore if Conjecture 3 were true then the BPV conjecture, which we expressed earlier in terms of cyclotomic subgroups $G_{q,p,k}$, would certainly hold as well.

Finally, we formulate the obvious strengthened version of Conjecture 3, including a bound on the degree.

Conjecture 4 (strong k -BPV) *Let $k > 1$ be an integer. There exists a proper divisor d of k such that d divides $\phi(k)$ and for which **strong** $(d, k/d)$ -BPV holds.*

Our preparatory work on the coefficients implies the correctness of Conjecture 4 and hence of Conjecture 3 for a whole family of values of k .

Proposition 2 *Let $2^s p_1^{r_1} p_2^{r_2} \dots p_n^{r_n}$ be the prime factorisation of $k > 1$ with $2 < p_1 < \dots < p_n$. Then $\text{red}(k) \leq \phi(2^s p_1^{r_1}) p_2^{r_2} \dots p_n^{r_n}$. In particular, if k is of the form $2^s p_1^{r_1}$ then Conjecture 4 holds for k ; and $\text{red}(k) = \phi(k)$ in this case.*

Proof. If $k = 2^s$ and $s \geq 1$ then taking $d = 2^{s-1}$ and $e = 2$ in the *even-even* case of Proposition 1 gives the result since $\phi(2^s) = 2^{s-1}$. Similarly, if $k = p_1^{r_1}$ with $r_1 \geq 1$, the result follows from taking $d = p_1^{r_1-1}$ and $e = p_1$ in the *odd-odd* case as $\phi(p_1^{r_1}) = (p_1 - 1)p_1^{r_1-1}$. The general case of the first part of the result (where $r_1, s \geq 1$) follows by taking $d = 2^s p_1^{r_1-1} p_2^{r_2} \dots p_n^{r_n}$ and $e = p_1$ in the *even-odd* case. The final part follows directly, using the observation that $\text{red}(k) \geq \phi(k)$.

Proposition 2 implies that Conjecture 4 holds for all $k \leq 30$ with $k \neq 15, 21, 30$.

Example 6. The first case of real interest of Conjecture k -BPV is $k = 30$, as there an improvement ratio of $30/8 > 3$ might be obtained. By virtue of Proposition 1 it follows that, by choosing $e = 3$, we can represent elements of the cyclotomic group $G_{p,30}$ as minimal polynomials using a single coefficient from the subfield of p^{10} elements, so using $10 \log p$ bits. Therefore $\text{red}(30) \leq 10$.

Conjecture 30-BPV states that in fact only $\phi(30) \log p = 8 \log p$ bits are necessary. More specifically the conjecture says that for some divisor d of $\phi(30) = 8$, the minimal polynomial over \mathbb{F}_{p^d} of any element of $G_{p,30}$ can be generated by eight elements of \mathbb{F}_p . For $d = 2$, for example, it could be that only the four highest of the non-trivial coefficients of minimal polynomials over \mathbb{F}_{p^2} are independent, and that the others can be expressed as polynomial expressions in these. This would represent a significant improvement on the upper bound provided by Proposition 1, which states that the seven highest non-trivial coefficients are sufficient to generate the others.

The 25 pairs (d, e) with $de \leq 30$ for which Proposition 1 and Example 5 do *not* provide a proof of Conjecture 2 are listed in the table at the end of Section 7.

6 The Magma programs

In order to test the conjectures formulated in the previous section we performed some experiments using the computer algebra system MAGMA [3].

Algorithm 0 (Find relations) *Input: integers p, k, d, u, v, j .*

Output: a set \mathcal{Q} of polynomials in $\mathbb{Z}[X_1, \dots, X_u, Y]$.

Description:

Determine a prime divisor q of $\Phi_k(p)$ not dividing k (Lemma 1), and a generator h of $G_{q,p,k}$ (e.g. taking the $(p^k - 1)/q$ -th power of a primitive element g of \mathbb{F}_{p^k}).

Next, generate the finite set S of all sequences $[s_1, \dots, s_u]$ with $\sum_{i=1}^u i \cdot s_i \leq v$.

Now generate $s = \#S$ random elements h_1, \dots, h_s of $G_{q,p,k}$, for example by taking random powers of h , determine the minimal polynomials

$$P_{h_i}^{(d)} = X^{k/d} + a_{k/d-1}X^{k/d-1} + \dots + a_1X + a_0,$$

of these elements over \mathbb{F}_{p^d} and evaluate $m(h_i, \mathbf{s}) = a_{e-1}^{s_1} \cdot a_{e-2}^{s_2} \cdot \dots \cdot a_{e-u}^{s_u}$ for all $\mathbf{s} \in S$. Let M be the square $s \times s$ matrix with entries in \mathbb{F}_{p^d} , the i -th row of which consists of the monomials $m(h_i, \mathbf{s})$, for \mathbf{s} ranging over S . Let $\mathbf{w} \in \mathbb{F}_{p^d}^s$ consist of the coefficients a_j (with j given by the input) of the polynomials $P_{h_i}^{(d)}$, for $i = 1, 2, \dots, s$.

Solve the linear system of equations $M\mathbf{c} = \mathbf{w}$ for $\mathbf{c} \in \mathbb{F}_p^s$. If the solution space is non-empty, translate each element \mathbf{c} from the solution space back to a polynomial relation $C \in \mathbb{Z}[X_1, \dots, X_u]$ via

$$\mathbf{c} \mapsto C = \sum_{\mathbf{s} \in S} c_{\mathbf{s}} X_1^{s_1} \cdots X_u^{s_u} - Y,$$

where on the right we interpret the component $c_s \in \mathbb{F}_p$ of the vector \mathbf{c} as an integer by taking the least integer representative for its residue class modulo p .

Finally, determine the Gröbner basis \mathcal{Q} of the ideal generated by these relations in $\mathbb{Q}[X_1, \dots, X_u, Y]$.

This ends the description of the algorithm.

The output of the algorithm consists of polynomials in u variables that form a basis for all polynomial relations $Q(a_{e-1}, \dots, a_{e-u}) - a_j = 0$ between the coefficients of the minimal polynomial for generators of the cyclotomic subgroup $G_{q,p,k}$, satisfying the condition that the weighted degree of Q is at most v .

To verify Conjecture 2 for a pair (d, e) we apply the following algorithm, with input d, e and with $w = \lceil \phi(de)/d \rceil$.

Algorithm 2 *Input: integers d, e, w .*

Output: either 'false' or sets \mathcal{Q}_j of candidate polynomials for Conjecture 2.

Description:

Let $u = \lceil \phi(de)/d \rceil$.

Repeat the following step for $j = e - u - 1, e - u - 2, \dots, 1$ in succession, terminating with output 'false' when an empty set \mathcal{Q}_j is encountered, and with sets $\mathcal{Q}_j, j = e - u - 1, \dots, 1$, as output otherwise:

Choose a prime number p , and apply Algorithm 0 with input $p, k = de, d, u, v = w, j$ to determine a set \mathcal{Q}_j .

If Algorithm 2 returns 'false', Conjecture 2 is refuted for the given values of d, e as no polynomial relation exists (for at least one j) of weighted degree at most u_d that works modulo p . Otherwise it returns candidate polynomials \mathcal{Q}_j expressing a_j in a_{e-1}, \dots, a_{e-u} . These candidates have only been proven to work for a single prime number p ; to increase confidence one would test the candidates for different values of p .

A (less effective) alternative to Algorithm 2 consists of a *single* application of Algorithm 0 rather than $e - u - 1$ successive ones, by replacing in the input for Algorithm 0 the values of u and v by $e - 1$, and putting $j = 0$. The result will be that Algorithm 0 will attempt to find all algebraic relations between all a_i 's (of weighted degree bounded by v) in one go; if the Conjectured relations exist, the Gröbner basis will exhibit them all. This approach is only feasible for very small values of de (see Example 7).

Algorithm 2 rarely succeeds; it is designed to refute Conjecture 2 for pairs d, e . Likewise, the following algorithm is designed to refute Conjecture 1.

Algorithm 1 *Input: integers d, e .*

Output: either 'false' or sets \mathcal{Q}_j of candidate polynomials for Conjecture 1.

Description:

Repeatedly apply Algorithm 2 with input triples d, e, w , until sets \mathcal{Q}_j are returned, starting with $w = \lceil \phi(de)/d \rceil$, and incrementing w by 1 when Algorithm 2 returns 'false'.

It should now also be clear how to attempt to refute (or prove) Conjectures 3, 4: apply Algorithms 1, 2 for all pairs (d, e) of divisors of k with d dividing $\phi(k)$.

As stated, the algorithms do not look for dependencies involving the \mathbb{F}_p -conjugates of a_{e-1}, \dots, a_{e-u} . The reason for this is that initially we attempt to find dependencies that do not involve the proper conjugates; the algorithm can easily be modified to include them, but doing this blows up the number of variables in the monomials by a factor d^u . We have omitted this from the description of the algorithms for the sake of clarity.

The Gröbner basis of the ideal is determined to detect dependencies between relations that are found. In our experiments usually one of three things happened: either no relation was found, or many relations were found due to the fact that the prime was chosen too small, or a few relations were found that generated an ideal with a Gröbner basis consisting of a single polynomial relation expressing the dependency of a_{e-u-1} on a_{e-1}, \dots, a_{e-u} . See [5], for example, for a discussion of Gröbner bases.

The main feature of Algorithm 0 is that of converting the problem of finding a polynomial relation to finding the kernel of a matrix over a finite field: the columns of the matrix correspond to the monomials, and the existence of an algebraic relation between the coefficients implies a dependency between the evaluations of the monomials at the coefficients, that is between the columns of the matrix. Thus the problem is reduced to linear algebra over \mathbb{F}_{p^d} .

The experiments we carried out deviated slightly from the description in this section: the indeterminate Y in Algorithm 0 was given weight $u + 1$ and was allowed to appear with exponent larger than 1 in the polynomial relations (see Example 10). For that reason our searches started at weight (a multiple of) $\lceil \phi(de)/d \rceil + 1$, exceeding the minimal value predicted by Conjecture 2.

Example 7. As a first example we present the output of our algorithm for $k = 4$.

Conjecture 1 holds for $(d, e) = (2, 2)$ since by Theorem 1 the minimal polynomials over \mathbb{F}_{p^2} of elements of $G_{p,4} \setminus \mathbb{F}_{p^2}$ are of the form $X^2 + a_1X + 1$, with $a_1 \in \mathbb{F}_{p^2}$.

The other case for $k = 4$ is $d = 1, e = 4$. The minimal polynomials over \mathbb{F}_p of elements of $G_{p,4} \setminus \mathbb{F}_{p^2}$ are of the form $X^4 + a_3X^3 + a_2X^2 + a_1X + a_0, a_0 = 1$.

In a run of Algorithm 2 we used $p = 5$; since $\Phi_4(5) = 5^2 + 1 = 26$, we look at $G_{13,5,4}$ in $G_{5,4}$. To find possible algebraic relations between a_3, a_2 and a_1 we take $u = 2$; we choose $w = 3$. The only monomials besides Y we obtain are $X_1^2, X_2, X_1, 1$. One run of our algorithm produced two dependencies in the matrix M , corresponding to an ideal with Gröbner basis $X_1 - Y, X_2 - Y^2 + Y + 1$. The first of these expresses that $a_3 = a_1$, as we expect by Theorem 1, but the second is an ‘accident’ caused by the fact that we have chosen p and thereby q to be very small. Indeed, in this case several minimal polynomials coincided (since the 12 non-trivial elements have just 3 different minimal polynomials). This illustrates why small primes p should be avoided.

If we invoke the algorithm with $p = 101$ instead, we immediately find a single relation $a_3 - a_1 = 0$ and no others. As a matter of fact, if we increase the parameters u and w to 3 and 4, the result will be a Gröbner basis $a_3 - a_1, a_0 - 1$: the minimal polynomials are always palindromic, and we have rediscovered Corollary 1 for this case.

Refuting Conjecture 1 (without the degree bounds) would involve looking at evaluations of *all* possible monomials in u coefficients. But since there is only a finite number of (different powers of) elements in \mathbb{F}_{p^k} anyway, this is still a finite task! However, the necessary computation can only be done if p is very small (say 2 or 3), in which case we run into problems similar to those in the previous example for small p , namely that the order of $G_{p,k}$ would be smaller than the number of monomials and we would obtain many unwanted identities.

7 Experimental Results

In this section we describe the experiments we have performed to test the conjectures. We looked at all cases with $k = de \leq 30$ still left open, as summarised in the table at the end of this section. We comment on some interesting cases, in order of ascending k .

Example 8. $k = 6$

This provided a test case for our programs (see the earlier examples). There are three pairs (d, e) to consider.

The case $d = 3, e = 2$ is trivial, since in the quadratic extension \mathbb{F}_{p^6} over \mathbb{F}_{p^3} elements are given by a single non-trivial element of the palindromic minimal polynomial; elements are thus represented by one element of \mathbb{F}_{p^3} , which is in accordance with Conjecture 1.

With $d = 1, e = 6$ Algorithm 2 finds (within a few seconds) the relations $-a_5^2 + 2a_4 + 2a_5 - 2 = a_3, a_4 = a_2$ and $a_5 = a_1$, for example with $p = 211$. Thus elements of a degree six field can be represented by the elements a_5 and a_4 from the prime field. This proves, as noted before, both Conjecture 2, and hence Conjecture 1, for $(d, e) = (1, 6)$, as well as Conjecture 4 and hence Conjecture 3, for $k = 6$.

With $d = 2, e = 3$ we have a cubic extension with k even. The standard Algorithm 2 does not find small relations; however, if we include the conjugates a_2^p and a_1^p as well as a_2 and a_1 , then with $p = 29$ the algorithm produces the relation $a_2 + a_1^p = 0$ (and also, in fact, $a_1^{p^2} - a_1 = 0$ in the Gröbner basis). Note the conflicting constraints on p once we include the conjugates: we want p to be large to avoid spurious relation in a small field, whereas we want it to be small since we get monomials including $a_i^{p^j}$ in the relations. The relation $a_2 + a_1^p = 0$ means we can represent the degree 6 field by a single element a_1 from the quadratic subfield; the element a_2 can then be recovered. This proves Conjecture 1 for $(d, e) = (2, 3)$ and Conjecture 3 for $k = 6$ (again).

Example 9. $k = 9$

Conjecture 2 and Conjecture 4 for $d = e = 3$ are covered by Proposition 1.

To achieve the same efficiency in the full extension ($d = 1, e = 9$) one would have to express a_2 and a_1 in terms of a_8, a_7, \dots, a_3 . Our experiments with Algorithm 2 show that no such relations exist with $u = 6$ and $w = 7$. That is, Conjecture 2 is false for this case.

In order to investigate whether a relaxation as in Conjecture 1 with regard to the degree of the polynomials involved would hold, we increased the search bound in Algorithm 0 to $w = 28$. In this and the cases to follow, we took w as a multiple of $u+1$ and took w also as an upper bound on the weighted degree of the polynomials when all variables are taken into consideration; that is, we searched for polynomials in $\mathbb{Z}[X_1, \dots, X_u, Y]$ for which $\sum_{i=1}^u i \cdot s_i + (u+1) \cdot s \leq w$, where s is the exponent of Y . In the current case that simply means that we allowed polynomials involving Y up to the 4-th power. Thus we even consider relations for a_2 involving a_2^4 .

No relations were found with $w = 28$. The computation involved computing the kernel of a 8561×8561 matrix over \mathbb{F}_p (using $p = 2003$).

Example 10. $k = 10$

The case $d = 2, e = 5$ can be done using 2 elements from \mathbb{F}_{p^2} by Proposition 1, proving Conjecture 1 for $(d, e) = (2, 5)$ but also Conjecture 3 for $k = 10$.

For $d = 1, e = 10$ Proposition 1 shows that 5 elements of \mathbb{F}_p suffice; Conjecture 1 predicts that 4 should be enough. We therefore invoke the Algorithm with $u = 4$; we found the following relation, but only after raising the search limit to $w = 15$ (using $p = 1009$):

$$\begin{aligned}
& a_9^8 + 2 \cdot a_7^7 - 8 \cdot a_9^6 \cdot a_8 - 2 \cdot a_9^6 \cdot a_7 + a_9^6 \cdot a_5 - 12 \cdot a_9^5 \cdot a_8 + 4 \cdot a_9^5 \cdot a_7 + 4 \cdot a_9^5 \cdot a_5 \\
& - 4 \cdot a_9^5 + 21 \cdot a_9^4 \cdot a_8^2 + 12 \cdot a_9^4 \cdot a_8 \cdot a_7 - 2 \cdot a_9^4 \cdot a_8 \cdot a_6 - 6 \cdot a_9^4 \cdot a_8 \cdot a_5 + 2 \cdot a_9^4 \cdot a_8 \\
& + 2 \cdot a_9^4 \cdot a_7^2 - 2 \cdot a_9^4 \cdot a_7 \cdot a_5 + 12 \cdot a_9^4 \cdot a_7 + a_9^4 \cdot a_6^2 - 10 \cdot a_9^4 \cdot a_6 + 2 \cdot a_9^4 \cdot a_5 - 3 \cdot a_9^4 \\
& + 20 \cdot a_9^3 \cdot a_8^2 - 16 \cdot a_9^3 \cdot a_8 \cdot a_7 - 4 \cdot a_9^3 \cdot a_8 \cdot a_6 - 16 \cdot a_9^3 \cdot a_8 \cdot a_5 + 16 \cdot a_9^3 \cdot a_8 \\
& - 6 \cdot a_9^3 \cdot a_7^2 + 8 \cdot a_9^3 \cdot a_7 - 12 \cdot a_9^3 \cdot a_6 + 2 \cdot a_9^3 \cdot a_5^2 + 4 \cdot a_9^3 - 20 \cdot a_9^2 \cdot a_8^3 \\
& - 20 \cdot a_9^2 \cdot a_8^2 \cdot a_7 + 8 \cdot a_9^2 \cdot a_8^2 \cdot a_6 + 10 \cdot a_9^2 \cdot a_8^2 \cdot a_5 - 8 \cdot a_9^2 \cdot a_8^2 - 8 \cdot a_9^2 \cdot a_8 \cdot a_7^2 \\
& + 4 \cdot a_9^2 \cdot a_8 \cdot a_7 \cdot a_6 + 8 \cdot a_9^2 \cdot a_8 \cdot a_7 \cdot a_5 - 32 \cdot a_9^2 \cdot a_8 \cdot a_7 - 4 \cdot a_9^2 \cdot a_8 \cdot a_6^2 \\
& - 2 \cdot a_9^2 \cdot a_8 \cdot a_6 \cdot a_5 + 32 \cdot a_9^2 \cdot a_8 \cdot a_6 + 4 \cdot a_9^2 \cdot a_8 - 2 \cdot a_9^2 \cdot a_7^3 + a_9^2 \cdot a_7^2 \cdot a_5 \\
& + 12 \cdot a_9^2 \cdot a_7 \cdot a_6 + 16 \cdot a_9^2 \cdot a_7 \cdot a_5 - 4 \cdot a_9^2 \cdot a_7 - 4 \cdot a_9^2 \cdot a_6^2 - 6 \cdot a_9^2 \cdot a_6 \cdot a_5 \\
& + 4 \cdot a_9^2 \cdot a_5^2 + 2 \cdot a_9^2 \cdot a_5 + 8 \cdot a_9^2 - 8 \cdot a_9 \cdot a_8^3 + 16 \cdot a_9 \cdot a_8^2 \cdot a_7 + 8 \cdot a_9 \cdot a_8^2 \cdot a_6 \\
& + 12 \cdot a_9 \cdot a_8^2 \cdot a_5 - 16 \cdot a_9 \cdot a_8^2 + 12 \cdot a_9 \cdot a_8 \cdot a_7^2 - 16 \cdot a_9 \cdot a_8 \cdot a_7 - 8 \cdot a_9 \cdot a_8 \cdot a_6 \cdot a_5 \\
& + 16 \cdot a_9 \cdot a_8 \cdot a_6 - 4 \cdot a_9 \cdot a_8 \cdot a_5^2 - 8 \cdot a_9 \cdot a_8 - 4 \cdot a_9 \cdot a_7^2 \cdot a_5 + 4 \cdot a_9 \cdot a_7^2 \\
& + 8 \cdot a_9 \cdot a_7 \cdot a_6^2 - 16 \cdot a_9 \cdot a_7 \cdot a_6 - 2 \cdot a_9 \cdot a_7 \cdot a_5^2 + 8 \cdot a_9 \cdot a_7 \cdot a_5 - 8 \cdot a_9 \cdot a_6^2 \\
& - 16 \cdot a_9 \cdot a_6 \cdot a_5 + 8 \cdot a_9 \cdot a_6 + 2 \cdot a_9 \cdot a_5^2 + 4 \cdot a_9 \cdot a_5 + 4 \cdot a_8^4 + 8 \cdot a_8^3 \cdot a_7 - 8 \cdot a_8^3 \cdot a_6 \\
& - 4 \cdot a_8^3 \cdot a_5 + 8 \cdot a_8^3 + 8 \cdot a_8^2 \cdot a_7^2 - 8 \cdot a_8^2 \cdot a_7 \cdot a_6 - 4 \cdot a_8^2 \cdot a_7 \cdot a_5 + 16 \cdot a_8^2 \cdot a_7 \\
& + 4 \cdot a_8^2 \cdot a_6^2 + 4 \cdot a_8^2 \cdot a_6 \cdot a_5 - 16 \cdot a_8^2 \cdot a_6 + a_8^2 \cdot a_5^2 - 8 \cdot a_8^2 \cdot a_5 + 4 \cdot a_8 \cdot a_7^3 \\
& - 4 \cdot a_8 \cdot a_7^2 \cdot a_6 - 2 \cdot a_8 \cdot a_7^2 \cdot a_5 - 16 \cdot a_8 \cdot a_7 \cdot a_6 - 16 \cdot a_8 \cdot a_7 \cdot a_5 + 8 \cdot a_8 \cdot a_7 \\
& + 16 \cdot a_8 \cdot a_6^2 + 8 \cdot a_8 \cdot a_6 \cdot a_5 - 8 \cdot a_8 \cdot a_6 - 4 \cdot a_8 \cdot a_5 - 8 \cdot a_8 + a_7^4 - 4 \cdot a_7^3 \\
& - 4 \cdot a_7^2 \cdot a_6 - 6 \cdot a_7^2 \cdot a_5 + 8 \cdot a_7^2 + 8 \cdot a_7 \cdot a_6^2 + 8 \cdot a_7 \cdot a_6 \cdot a_5 - 8 \cdot a_7 \cdot a_6 + 2 \cdot a_7 \cdot a_5^2 \\
& + 4 \cdot a_7 \cdot a_5 - 8 \cdot a_6^3 - 4 \cdot a_6^2 \cdot a_5 + 12 \cdot a_6^2 + 2 \cdot a_6 \cdot a_5^2 + 4 \cdot a_6 \cdot a_5 + a_5^3 + 3 \cdot a_5^2 - 4
\end{aligned}$$

A single run with these parameters took around 10 seconds. The size of the matrix, determined by the number of monomials involved, is 408×408 .

This relation poses some interesting questions; since the equation is of degree at least 3 in each of the variables, in general there will not be a *unique* solution for the variable a_5 , given values for the a_9, \dots, a_6 . Moreover, the polynomial is irreducible when we consider it as a polynomial in $F[a_i]$ for all $i \in \{9, 8, 7, 6, 5\}$, with F the field of rational functions in the other four variables.

Using this — impractical — relation, an improvement factor of $10/4$ is achieved; less than in XTR but more than in LUC.

Example 11. $k = 12, 24$

For $k = 12$ and $d = 1$ we found polynomials Q_7 and Q_6 expressing a_7 and a_6 in a_{11}, \dots, a_8 ; these polynomials are of weighted degree 15 and 18 respectively. As in the previous case they contain powers of a_7 and a_6 greater than 1.

For $k = 24$, $d = 2$ we found the same relations as for $k = 12$, $d = 1$.

Example 12. $k = 30$

Finally, the most interesting case.

k	d	e	$\lceil \phi(k)/d \rceil \cdot d$	Prop. 1
30	1	30	8	15
30	2	15	8	14
30	3	10	9	15
30	5	6	10	10
30	6	5	12	12
30	10	3	10	10
30	15	2	15	15

As before we compare the conjectured and proven bounds on the number of elements of \mathbb{F}_p that suffices. This shows that three cases of Conjecture 1 are still open. A quick run of Algorithm 2 showed that Conjecture 2 is false in each of the three cases $(3, 10)$, $(2, 15)$ and $(1, 30)$.

The table also shows that to prove Conjecture 3 for $k = 30$ we either need to prove that 8 elements of \mathbb{F}_p or 4 elements of \mathbb{F}_{p^2} will suffice to generate all coefficients. Our further search for relations in these cases had no success either; the search bounds are given below.

p	k	d	u	w	$\#S$
1009	30	1	8	27	10269
1009	30	1	11	24	6720
1009	30	1	14	25	9012
71	30	2	4	10	3616
71	30	2	5	6	1920
101	30	2	6	7	5760

The last column lists the number of monomials taken into consideration and hence the number of minimal polynomials generated. Note that these results (for $d = 2$) refer to a modification of Algorithm 2, discussed in Section 6, to include conjugates of the coefficients.

For the remaining open cases (see the table below) we searched for relations in vain. For each line in the table we ran Algorithm 0 for every u in the range from the conjectured value (inclusive) up to the proven bound (exclusive). For values of k exceeding 20 we ran Algorithm 0 only with $w = u + 1$ (thus only testing Conjecture 2), while for $k \leq 20$ we went further (in an attempt to prove Conjecture 1), by taking $w = 2(u + 1)$ or even $w = 3(u + 1)$.

When $d > 1$ we also ran the modified algorithm, taking the conjugates into account; only in the cases $k = 21, d = 3, u = 5$, and $k = 27, d = 3, u = 6, 7$ the resulting computation involved square matrices that were too large for us to deal with. The conjectured improvement factors in both cases are less than 2.

The table lists all 25 cases with $de \leq 30$ for which Proposition 1 and Example 5 do *not* provide a proof of Conjecture 2. It lists the value $\lceil \phi(k)/d \rceil$ for u_d^* predicted by Conjecture 2, the correct value as obtained by our experiments, and the upper bound S/d implied by Proposition 1 (cf. Table 1).

k	(d, e)	$\lceil \phi(k)/d \rceil$	u_d^*	S/d	k	(d, e)	$\lceil \phi(k)/d \rceil$	u_d^*	S/d
9	(1, 9)	6	8	8	24	(1, 24)	8	12	12
10	(1, 10)	4	5	5	24	(2, 12)	4	6	6
12	(1, 12)	4	6	6	24	(3, 8)	3	4	4
14	(1, 14)	6	7	7	25	(1, 25)	20	24	24
15	(1, 15)	8	14	14	26	(1, 26)	12	13	13
15	(3, 5)	3	4	4	27	(1, 27)	18	26	26
18	(1, 18)	6	9	9	27	(3, 9)	6	6, 7, 8	8
18	(2, 18)	3	4	4	28	(1, 28)	12	14	14
20	(1, 20)	8	10	10	28	(2, 14)	6	7	7
20	(2, 10)	4	5	5	30	(1, 30)	8	15	15
21	(1, 21)	12	20	20	30	(2, 15)	4	7	7
21	(3, 7)	4	5, 6	6	30	(3, 10)	3	5	5
22	(1, 22)	10	11	11					

Theorem 2 *Conjecture 2 is false for all pairs (d, e) covered by the table, with the possible exception of the case $(3, 9)$. For all (d, e) with $de \leq 30$, with $(3, 7)$ and $(3, 9)$ possibly excepted, the true value of u_d^* equals the upper bound implied by Proposition 1. Moreover, Conjecture 4 is false for $k = 30, 21, 15$, i.e. the cases ≤ 30 not covered by Proposition 1.*

8 Conclusion

Based on generalisations of the LUC and XTR methods we have formulated precise and verifiable versions of the Brouwer-Pellikaan-Verheul conjecture posed in [4]. By experiment we have shown that it is unlikely that a compact representation of elements exists in extension fields of degree thirty, providing some evidence that XTR cannot be improved with respect to compactness of representation.

Our experiments leave open the possibility that the conjectures hold with polynomials of large degree, which most likely would be of no practical value.

References

1. M. Adleman, J. DeMarrais *A subexponential algorithm over all finite fields*, CRYPTO '93 Proc., Springer-Verlag, 147-158.
2. D. Bleichenbacher, W. Bosma, A.K. Lenstra, *Some remarks on Lucas-Based Cryptosystems*, CRYPTO '95 Proceedings, Springer-Verlag, pp. 386-396.
3. W. Bosma, J.J. Cannon, C. Playoust, *The Magma Algebra System I: The User Language*, Journal of Symbolic Computation **24** (1997), 235-265.
4. A.E. Brouwer, R. Pellikaan, E.R. Verheul, *Doing more with fewer bits*, Proceedings Asiacypt99, LNCS 1716, Springer-Verlag 1999, 321-332.
5. D. Cox, J. Little, D. O'Shea, *Ideals, Varieties, and Algorithms*, Springer, 1992.
6. W. Diffie, M.E. Hellman, *New directions in cryptography*, IEEE Trans. on IT **22**, 1976, 644-654.
7. G. Gong, L. Harn, *Public key cryptosystems based on cubic finite field extensions*, IEEE Trans. on I.T., November 1999.
8. S. Lang, *Algebra*, Addison-Welsey, 1993.
9. A.K. Lenstra, *Using Cyclotomic Polynomials to Construct Efficient Discrete Logarithm Cryptosystems over Finite Fields*, Information Security and Privacy - ACISP97 Proceedings (Sydney 1997), Lect. Notes in Comp. Sci. 1270, Springer-Verlag, pp. 127-138.
10. A.K. Lenstra, E.R. Verheul, *The XTR public key system*, Proceedings of Crypto 2000, LNCS 1880, Springer-Verlag, 2000, 1-19; available from www.ecstr.com.
11. A.K. Lenstra, E.R. Verheul, *Key improvements to XTR*, Proceedings of Asiacypt 2000, LNCS 1976, Springer-Verlag, 2000, 220-223; available from www.ecstr.com.
12. A.K. Lenstra, E.R. Verheul, *Fast irreducibility and subgroup membership testing in XTR*, Proceedings of the 2001 Public Key Cryptography conference, LNCS 1992, Springer-Verlag, 2001, 73-86; available from www.ecstr.com.
13. A.K. Lenstra, E.R. Verheul, *An overview of the XTR public key system*, In: Public-Key Cryptography and Computational Number Theory, Walter de Gruyter, 2001, 151-180.
14. R. Lidl, W.B. Müller, *Permutation Polynomials in RSA-cryptosystems*, Crypto '83 Proceedings, Plenum Press, pp. 293-301.
15. R. Lidl, H. Niederreiter, *Finite Fields*, Addison-Wesley, 1983.
16. W.B. Müller, *Polynomial functions in modern cryptology*, Contributions to general Algebra 3, Proceedings of the Vienna Conference (1985), pp. 7-32. Proceedings, Springer-Verlag, pp. 50-61.
17. W.B. Müller, W. Nöbauer, *Cryptanalysis of the Dickson-Scheme*, Eurocrypt '85 Proceedings, Springer-Verlag, pp. 50-61.
18. W.K. Nicholson, *Introduction to abstract algebra*, PWS-Kent Publishing Company, Boston, 1993.
19. W. Nöbauer, *Cryptanalysis of the Rédei Scheme*, Contributions to general Algebra 3, Proceedings of the Vienna Conference (1985), pp. 255-264.
20. J.M. Pollard, *Monte Carlo methods for index computation (mod p)*, Math. Comp., **32** (1978), 918-924.
21. C.P. Schnorr, *Efficient signature generation by smart cards*, Journal of Cryptology, **4** (1991), 161-174.
22. P. Smith, C. Skinner, *A public-key cryptosystem and a digital signature system based on the Lucas function analogue to discrete logarithms*, Asiacypt '94 proceedings, Springer-Verlag, pp. 357-364.
23. M. Stam, A.K. Lenstra, *Speeding Up XTR*, Proceedings of Asiacypt 2001, LNCS 2248, Springer-Verlag, 2001, 125-143; available from www.ecstr.com.