Polly Cracker, Revisited^{*}

Martin R. Albrecht¹, Pooya Farshim², Jean-Charles Faugère¹,
and Ludovic Perret¹

¹ INRIA, Paris-Rocquencourt Center, SALSA Project
UPMC Univ Paris 06, UMR 7606, LIP6, F-75005, Paris, France
CNRS, UMR 7606, LIP6, F-75005, Paris, France

² Department of Computer Science, Darmstadt University of Technology, Germany
malb@lip6.fr, farshim@cased.de, jean-charles.faugere@inria.fr,
ludovic.perret@lip6.fr

Abstract. We initiate the formal treatment of cryptographic constructions (“Polly Cracker”) based on the hardness of computing remainders modulo an ideal over multivariate polynomial rings. We start by formalising the relation between the ideal remainder problem and the problem of computing a Gröbner basis. We show both positive and negative results. On the negative side, we define a symmetric Polly Cracker encryption scheme and prove that this scheme only achieves bounded CPA security. Furthermore, we show that a large class of algebraic transformations cannot convert this scheme to a fully secure Polly-Cracker-style scheme. On the positive side, we formalise noisy variants of the ideal membership, ideal remainder, and Gröbner basis problems. These problems can be seen as natural generalisations of the **LWE** problem and the approximate GCD problem over polynomial rings. We then show that noisy encoding of messages results in a fully **IND-CPA**-secure somewhat homomorphic encryption scheme. Our results provide a new family of somewhat homomorphic encryption schemes based on new, but natural, hard problems. Our results also imply that Regev’s **LWE**-based public-key encryption scheme is (somewhat) *multiplicatively* homomorphic for appropriate choices of parameters.

Keywords: Polly Cracker, Gröbner bases, **LWE**, Noisy encoding, Homomorphic encryption, Public-key encryption, Provable security.

1 Introduction

BACKGROUND. Homomorphic encryption [19] is a cryptographic primitive which allows performing arbitrary computations over encrypted data. From an

^{*} The work described in this paper has been supported by the Royal Society grant JP090728 and by the Commission of the European Communities through the ICT program under contract ICT-2007-216676 (ECRYPT-II). M. Albrecht, J.-C. Faugère, and L. Perret were also supported by the french ANR under the Computer Algebra and Cryptography (CAC) project (ANR-09-JCJCJ-0064-01) and the EXACTA project (ANR-09-BLAN-0371-01). P. Farshim was funded in part by the US Army Research laboratory, the UK Ministry of Defense and was accomplished under Agreement Number W911NF-06-3-0001. Due to space limitations this work is only an extended abstract of the full work available in [1].

algebraic perspective, this homomorphic feature can be seen as the ability to evaluate multivariate polynomials over ciphertexts. Hence, an instantiation of homomorphic encryption over multivariate polynomials is perhaps the most natural strategy.

Indeed, let $\mathcal{I} \subset P = \mathbb{F}[x_0, \dots, x_{n-1}]$ be some ideal. We can encrypt a message $m \in P/\mathcal{I}$ as $c = f + m$ for f randomly chosen in $\mathcal{I}$. Decryption is performed by computing remainders modulo $\mathcal{I}$. From the definition of an ideal the homomorphic features of this scheme follow. The problem of computing remainders modulo an ideal was solved by Buchberger in [8], where he introduced the notion of Gröbner bases, and gave an algorithm for computing such bases.

In fact, all known doubly homomorphic schemes are based on variants of the ideal remainder problem over various rings. For example in [13] the ring $\langle p \rangle \in \mathbb{Z}$ for p an odd integer is considered. In [19] ideals in a number field play the same role (cf. [29]). One can even view Regev’s LWE-based public-key encryption scheme [25] in this framework. Finally, we note that the construction displayed above is essentially Polly Cracker (PC) [17]. However, despite their simplicity, our confidence in PC-style schemes has been shaken as almost all such proposals have been broken [15]. In fact, it is a long standing open research challenge to propose a secure PC-style encryption scheme [5].

CONTRIBUTIONS & ORGANISATION. Our contributions can be summarised as follows: 1) we initiate the formal treatment of PC-style schemes and characterise their security; 2) we show the impossibility of converting such schemes to fully IND-CPA-secure schemes through a large class of transformations; 3) we introduce natural noisy variants of classical problems related to Gröbner bases which also generalise previously considered noisy problems; and 4) we present a new somewhat (and doubly) homomorphic encryption scheme based on a new class of computationally hard problems.

In more detail, after settling notation in Section 2, we formalise various problems from commutative algebra in the language of game-based security definitions in Section 3. In particular, we show that computing remainders modulo an ideal with overwhelming probability is equivalent to computing a Gröbner basis for zero-dimensional ideals. We then show that deciding ideal membership and computing ideal remainders are equivalent for certain choices of parameters. We then introduce a symmetric variant of Polly Cracker and characterise its security guarantees. We show that this scheme achieves *bounded* IND-CPA security, and that this level of security is the best that one can hope for: we give an attacker which breaks the cryptosystem once enough ciphertexts are obtained.

In Section 5, we show the security limitations of the constructed scheme are in some sense *intrinsic*. More precisely, we show that a large class of algebraic transformation cannot turn this scheme into a (fully) IND-CPA secure and additively homomorphic PC-style scheme.

To go beyond this limitation, we consider a constructions where the encoding of messages is randomised. To prove security for such schemes, we consider noisy variants of the ideal membership and related problems. These can be seen as natural generalisations of the (decisional) LWE and the approximate GCD

problems over polynomial rings (Section 6). After formalising and justifying the hardness of the noisy assumptions in Section 7, we show that noisy encoding of messages can indeed be used to construct a fully IND-CPA-secure somewhat homomorphic scheme. This result also implies that Regev’s LWE-based public-key scheme is *multiplicatively* homomorphic under appropriate choices of parameters. Our result, together with a standard symmetric-to-asymmetric conversion for homomorphic schemes, provides a positive answer to the long standing open problem proposed by Barkee et al. [5]. In addition, we provide a new family of somewhat homomorphic schemes which are based on new natural variants of well-studied hard problems. Due to space limitations, we discuss concrete parameter choices and include a reference implementation in the full version of the paper [1]. There, we also show how our scheme allows proxy re-encryption of ciphertexts. This re-encryption procedure can be seen as trading noise for degree in ciphertexts. That is, we can control the growth of the ciphertext size due to multiplication by tolerating more noise. We note that this technique was recently and independently developed in [7]. In [1], we also show that our scheme achieves a limited form of key-dependent message (KDM) security in the standard model, where the least significant bit of the constant term of the key is encrypted. We leave it as an open problem to adapt the techniques of [2] to achieve full KDM security for the Polly Cracker with noise scheme.

1.1 Related Work

Polly Cracker. In 1993, Barkee et al. wrote a paper [5] whose aim was to dispel the urban legend that “Gröbner bases are hard to compute”. Another goal of this paper was to direct research towards *sparse* systems of multivariate equations. To do so, the authors proposed the most obvious dense Gröbner-based cryptosystem, namely an instantiation of the construction mentioned at the beginning of the introduction. In their scheme, the public key is a set of polynomials $\{f_0, \dots, f_{m-1}\} \subset \mathcal{I}$ which is used to construct an element $f \in \mathcal{I}$. Encryption of messages $m \in P/\mathcal{I}$ are computed as $c = \sum h_i f_i + m = f + m$ for $f \in \mathcal{I}$. The private key is a Gröbner basis G which allows to compute $m = c \bmod \mathcal{I} = c \bmod G$. As highlighted in [5] this scheme can be broken using results from [12] (cf. Theorem 2). At about the same time, and independently from the work of Barkee et al., Fellows and Koblitz [17] proposed a framework for the design of public-key cryptosystems. The ideas in [17] were similar to Barkee et al.’s, but differed in some details. However, the main instantiation of such a system was the Polly Cracker cryptosystem. Subsequently, a variety of sparse PC-style schemes were proposed. The focus on sparse polynomials aimed to prevent the attack based on Theorem 2, yet almost all of these schemes were broken. We point the reader to [15] for a good survey of various constructions and attacks. Currently, the only PC-style scheme which is not broken is the scheme in [9]. This scheme is based on binomial ideals (which in turn are closely related to lattices). Not only can our constructions be seen as instantiations of Polly Cracker (with and without noisy encoding of messages), they also allow security proofs based on

the hardness of computational problems related to (multivariate) polynomial ideals with respect to random systems.

Homomorphic Encryption. With respect to doubly (i.e., additively and multiplicatively) homomorphic schemes, a number of different hardness assumptions and constructions appeared in the literature. These include the Ideal Coset Problem of Gentry [19], the approximate GCD problem over the Integers of van Dijk et al. [13], the Polynomial Coset Problem as proposed by Smart and Vercauteren in [29], the Approximate Unique Shortest Vector Problem, the Subgroup Decision Problem, and the Differential Knapsack Vector Problem which appear in [23]. The main difference between our work and previous work is that we base the security of our somewhat homomorphic scheme on *new* computational problems related to ideals over multivariate polynomial rings. Furthermore, due to the versatility of Gröbner basis theory, our work can be seen as a generalisation of a number of known schemes and their underlying hardness assumptions. However, while our construction is doubly homomorphic and reasonably efficient for low multiplicative circuit depths, it is currently an open problem how to make it bootstrappable and hence turn it into a fully homomorphic scheme.

$\mathcal{MQ}$ Cryptography. Our work bears some connection with public-key cryptosystems based on the hardness of solving multivariate quadratic equations ($\mathcal{MQ}$). The difference is that our cryptographic constructions enjoy strong reductions to the known and hard problem of solving a *random* system of equations, whereas the bulk of work in $\mathcal{MQ}$ cryptography relies on heuristic security arguments [14]. In contrast, our work is more in the direction of research initiated by Berbain et al. [6] who proposed a stream cipher whose security was reduced to the difficulty of solving a system of random multivariate quadratic equations over $\mathbb{F}_2$. Note also that the concept of adding noise to a system of multivariate equations has been also proposed by Gouget and Patarin in [21] for the design of an authentication scheme. Our work, however, presents a more general and complete treatment of problems related to ideals over multivariate polynomials – both with and without noise – and aims to provide a formal basis to assess the security of cryptosystems based on such problems.

2 Preliminaries

NOTATION. We write $x \leftarrow y$ for assigning value y to a variable x , and $x \leftarrow_{\$} X$ for sampling x from a set X uniformly at random. If A is a probabilistic algorithm we write $y \leftarrow_{\$} A(x_1, \dots, x_n)$ for the action of running A on inputs $x_1, \dots, x_n$ with uniformly chosen random coins, and assigning the result to y . For a random variable X we denote by $[X]$ the support of X , i.e., the set of all values that X takes with non-zero probability. We use ppt for probabilistic polynomial-time. We call $\eta(\lambda)$ negligible if $|\eta(\lambda)| \in \lambda^{-\omega(1)}$.

COMMUTATIVE ALGEBRA NOTATION. In [1] we recall some basic definitions related to Gröbner bases. For a more detailed treatment we refer to, for instance, [10]. We consider a polynomial ring $P = \mathbb{F}[x_0, \dots, x_{n-1}]$ over some finite

field (typically $\mathbb{F}_q$), some monomial ordering on elements of P , and a set of polynomials $f_0, \dots, f_{m-1}$. We denote by $M(f)$ the set of all monomials appearing in $f \in P$. By $LM(f)$ we denote the leading monomial appearing in $f \in P$ according to the chosen term ordering. We denote by $LC(f)$ the coefficient $\in \mathbb{F}$ corresponding to $LM(f)$ in f and set $LT(f) = LC(f) \cdot LM(f)$. We denote by $P_{<d}$ the set of polynomials of degree $< d$ (and analogously for the $>, \leq, \geq$, and $=$ relations). We define $P_{=0}$ as the underling field including $0 \in \mathbb{F}$. We define $P_{<0}$ as zero. Finally, we denote by $M_{<m}$ the set of all monomials $< m$ for some monomial m (and analogously for the $>, \leq, \geq$, and $=$ relations). We assume the usual power product representation for elements of P .

3 Gröbner Basis and Ideal Membership Problems

Following [11], we define a *computational polynomial ring scheme*. This is a general framework allowing to discuss in a concrete way the different families of rings that may be used in cryptographic applications. More formally, a computational polynomial ring scheme $\mathcal{P}$ is a sequence of probability distribution of *polynomial ring descriptions* $(\mathbf{P}_\lambda)_{\lambda \in \mathbb{N}}$. A polynomial ring description P specifies various algorithms associated with P such as computing ring operations, sampling elements, testing membership, encoding of elements, ordering of monomials, etc. We assume each polynomial ring distribution is over $n = n(\lambda)$ variables, for some polynomial $n(\lambda)$, and is over a finite prime field of size $q(\lambda)$.

In this work we denote by $\text{GBGen}(1^\lambda, P, d)$ an arbitrary ppt algorithm which outputs a reduced Gröbner basis G for some zero-dimensional ideal $\mathcal{I} \subset P$ such that every element of G is of degree at most d . Of particular interest to this paper is the Gröbner basis generation algorithm shown in Algorithm 1 called $\text{GBGen}_{\text{dense}}(\cdot)$. (Algorithm $\text{ReduceGB}(\cdot)$ is given in [1].) We show in [1] that $\text{GBGen}_{\text{dense}}(\cdot)$ returns a Gröbner basis. Throughout the paper we assume an implicit dependency of various parameters associated with P on the security parameter. Thus, we drop λ to ease notation.

Algorithm 1: Algorithm $\text{GBGen}_{\text{dense}}(1^\lambda, P, d)$

```

1 begin
2   if  $d = 0$  then return  $\{0\}$ ;
3   for  $0 \leq i < n$  do
4     for  $m_j \in M_{<x_i^d}$  do
5        $c_{ij} \leftarrow \$_{\mathbb{F}_q}$ ;  $g_i \leftarrow g_i + c_{ij}m_j$ ;
6   return  $\text{ReduceGB}(\{x_0^d + g_0, \dots, x_{n-1}^d + g_{n-1}\})$ ;
7 end
```

We can now formally define the problem of computing a Gröbner basis.

Definition 1. *The Gröbner basis problem is defined through the game denoted $\text{GB}_{\mathcal{P}, \text{GBGen}(\cdot), d, b, m}$ as shown in Figure 1. The advantage of a ppt algorithm $\mathcal{A}$ in*

solving the GB problem is defined as the probability of winning the game (i.e., the game returning **T**). An adversary is legitimate if it calls the **Sample** procedure at most $m = m(\lambda)$ times.

Initialize ($1^\lambda, \mathcal{P}, d$): begin $P \leftarrow_{\$} \mathcal{P}_\lambda$; $G \leftarrow_{\$} \text{GBGen}(1^\lambda, P, d)$; return $(1^\lambda, P)$; end	Sample (): begin $f \leftarrow_{\$} P_{\leq b}$; $f \leftarrow f - (f \bmod G)$; return f ; end	Finalize (): begin return $(G = G')$; end
--	--	---

Fig. 1. Game $\text{GB}_{\mathcal{P}, \text{GBGen}(\cdot), d, b, m}$

We show in [1] that **Sample** returns elements of degree b which are uniformly distributed in $\langle G \rangle$. We recall that given a Gröbner basis G of an ideal $\mathcal{I}$, $r = f \bmod \mathcal{I} = f \bmod G$ is the normal form of f with respect to the ideal $\mathcal{I}$. We sometimes drop the explicit reference to $\mathcal{I}$ when it is clear from the context which ideal we are referring to, and simply refer to r as the normal form of f . Computing normal forms is the ideal remainder problem which we formalise below.

Definition 2. *The ideal remainder problem is defined through the game shown in Figure 2: $\text{IR}_{\mathcal{P}, \text{GBGen}(\cdot), d, b, m}$. The advantage of a ppt algorithm $\mathcal{A}$ in solving the IR problem is defined as the probability of winning the game minus $1/q^{\dim_{\mathbb{F}_q}(P/\langle G \rangle)}$. An adversary is legitimate if it calls the **Sample** procedure at most $m = m(\lambda)$ times.*

Initialize ($1^\lambda, \mathcal{P}, d$): begin $P \leftarrow_{\$} \mathcal{P}_\lambda$; $G \leftarrow_{\$} \text{GBGen}(1^\lambda, P, d)$; return $(1^\lambda, P)$; end	Sample (): begin $f \leftarrow_{\$} P_{\leq b}$; $f' \leftarrow (f \bmod G)$; return $f - f'$; end	Challenge (): begin $f \leftarrow_{\$} P_{\leq b}$; return f ; end	Finalize (r'): begin $r \leftarrow f \bmod G$; return $r = r'$; end
--	--	---	---

Fig. 2. Game $\text{IR}_{\mathcal{P}, \text{GBGen}(\cdot), d, b, m}$

In Lemma 1 below we prove a weak form of equivalence between the above problems. That is, we require that the IR adversary returns the correct answer with an *overwhelming* probability. This is due to the restriction that **Sample** can only be called a bounded number of times, and thus one cannot amplify the success probability of the IR adversary through repetition. The weak statement is sufficient in our context.

Lemma 1. *If the GB problem is hard, then the IR problem is weakly hard (i.e., cannot be solved with overwhelming probability). Furthermore, if the IR problem is hard then so is the GB problem.*

The precise theorem statement and a proof is given in [1]. Informally, the reduction of the GB problem to the IR problem works as follows. Consider an arbitrary element g_i in the Gröbner basis G . We can write g_i as $m_i + \tilde{g}_i$ for some $\tilde{g}_i < g_i$ and $m_i = \text{LM}(g_i)$. Now, assume the normal form of m_i is r_i and suppose that $r_i < m_i$. This implies that $m_i = \sum_{j=0}^{n-1} h_j g_j + r_i$ for some $h_i \in P$. Hence, we have $m_i - r_i \in \langle G \rangle$, an element $\in \langle G \rangle$ with leading monomial m_i . Repeat this process for all monomials up to and including degree d and accumulate the results $m_i - r_i$ in a list $\tilde{G}$. The list $\tilde{G}$ is a list of elements $\in \langle G \rangle$ with $\text{LM}(\tilde{G}) \supseteq \text{LM}(G)$ which implies $\tilde{G}$ is a Gröbner basis. We note that this is the core idea behind the FGLM algorithm [16].

The decisional variant of the IR problem is to decide whether the normal form of some element modulo an ideal is zero or not, i.e., whether this element is in the ideal or not. This is the ideal membership problem formalised below.

Definition 3. *The ideal membership problem is defined through the the game denoted $\text{IM}_{\mathcal{P}, \text{GBGen}(\cdot), d, b, m}$ as shown in Figure 3. The advantage of a ppt algorithm $\mathcal{A}$ in solving IM is defined as twice the probability of winning the game minus 1. An adversary is legitimate if it calls the **Sample** procedure at most $m = m(\lambda)$ times.*

Initialize ($1^\lambda, \mathcal{P}, d$): begin $P \leftarrow_{\$} \mathcal{P}_\lambda$; $G \leftarrow_{\$} \text{GBGen}(1^\lambda, P, d)$; $c \leftarrow_{\$} \{0, 1\}$; return $(1^\lambda, P)$; end	Sample (): begin $f \leftarrow_{\$} P_{\leq b}$; $f' \leftarrow f \bmod G$; return $f - f'$; end	Challenge (): begin $f \leftarrow_{\$} P_{\leq b}$; if $c = 1$ then $f \leftarrow f - (f \bmod G)$; return f ; end	proc. Finalize (c'): begin return $(c = c')$; end
--	--	--	---

Fig. 3. Game $\text{IM}_{\mathcal{P}, \text{GBGen}(\cdot), d, b, m}$

Clearly any adversary which can solve the IR problem can also solve the IM problem. However, if the search space of reminders modulo $\langle G \rangle$ is sufficiently small, i.e., when $q^{\dim_{\mathbb{F}_q}(P/\langle G \rangle)} = \text{poly}(\lambda)$, and under similar assumptions as for Lemma 1, one can also perform the converse reduction. That is, one can solve the IR problem using an oracle for the IM problem. Lemma 2 below proves this equivalence for the special case of $\text{GBGen}_{\text{dense}}(\cdot)$. Once again, this is sufficient in our context. As before, for Lemma 2 to be meaningful we require that the IM adversary returns the correct answer with *overwhelming* probability.

Lemma 2. *If the IR problem is hard, then the IM problem is weakly hard for poly-sized $q^{\dim_{\mathbb{F}_q}(P/\langle G \rangle)}$. Furthermore, if the IM problem is hard, then the IR problem is also hard.*

Informally, the construction of an IR adversary from an IM adversary proceeds as follows. Let $\tilde{f}$ be the challenge polynomial. The attacker simply exhaustively searches all elements of the $\mathbb{F}_q$ vector space $P/\langle G \rangle$ until the right remainder r is found. This occurs if $f - r \in \langle G \rangle$ and can be then detected using an IM

adversary. However, there is a technical difficulty here. In general, the attacker does not necessarily know the support of $P/\langle G \rangle$ and hence cannot know how to construct r . However, in our case we assume that $\text{GBGen}(\cdot) = \text{GBGen}_{\text{dense}}(\cdot)$ and this difficulty does not arise. In a more general setting, we would have to discover $P/\langle G \rangle$ as well (cf. proof of Lemma 4). See [1] for the proof.

Complexity estimation about Gröbner basis computations [1], together with the above results, lead to the following hardness assumptions.

Definition 4. *Let $\mathcal{P}$ be such that $n(\lambda) = \Omega(\lambda)$. Assume $b - d > 0$, $b > 1$, and that $m(\lambda) = c \cdot n(\lambda)$ for a constant $c \geq 1$. Then the advantage of any ppt algorithm in solving the GB/IR/IM problem is negligible as function of λ .*

4 Symmetric Polly Cracker: Noise-Free Version

4.1 Homomorphic Symmetric Encryption

SYNTAX. A *homomorphic symmetric-key encryption scheme* (HSKE) is specified by four ppt algorithms: 1) $\text{Gen}(1^\lambda)$ is the key generation algorithm and returns a key pair (SK, PK) , a message space $\text{MsgSp}(\text{PK})$ and a function space $\text{FunSp}(\text{PK})$. 2) $\text{Enc}(\text{m}, \text{SK})$ is the encryption algorithm and returns a ciphertext c . 3) $\text{Eval}(\text{c}_0, \dots, \text{c}_{t-1}, C, \text{PK})$ is the evaluation algorithm and outputs a ciphertext c_{evl} . 4) $\text{Dec}(\text{c}_{\text{evl}}, \text{SK})$ is the deterministic decryption algorithm and returns either a message m or a special failure symbol $\perp$.

CORRECTNESS. An HSKE scheme is correct if for any $\lambda \in \mathbb{N}$, any $(\text{SK}, \text{PK}) \in [\text{Gen}(1^\lambda)]$, any t messages $\text{m}_i \in \text{MsgSp}(\text{PK})$, any $\text{c} \in [\text{Enc}(\text{m}, \text{SK})]$, any circuit $C \in \text{FunSp}(\text{PK})$, any t ciphertexts $\text{c}_i \in [\text{Enc}(\text{m}_i, \text{PK})]$, and any evaluated ciphertext $\text{c}_{\text{evl}} \in [\text{Eval}(\text{c}_0, \dots, \text{c}_{t-1}, C, \text{PK})]$, we have that $\text{Dec}(\text{c}_{\text{evl}}, \text{SK}) = C(\text{m}_0, \dots, \text{m}_{t-1})$. We do not necessarily require correctness over freshly created ciphertexts.

COMPACTNESS. An HSKE scheme is compact if there exists a fixed polynomial bound $B(\cdot)$ so that for any key pair $(\text{SK}, \text{PK}) \in [\text{Gen}(1^\lambda)]$, any circuit $C \in \text{FunSp}(\text{PK})$, any set of t messages $\text{m}_i \in \text{MsgSp}(\text{PK})$, any ciphertext $\text{c}_i \in [\text{Enc}(\text{m}_i, \text{SK})]$, and any evaluated ciphertext $\text{c}_{\text{evl}} \in [\text{Eval}(\text{c}_0, \dots, \text{c}_{t-1}, C, \text{PK})]$, the size of c_{evl} is at most $B(\lambda + |C(\text{m}_0, \dots, \text{m}_{t-1})|)$ (independently of the size of C).

The syntax, correctness, and compactness of a homomorphic public-key encryption scheme is defined similarly.

4.2 The Scheme

In this section we formally define the (noise-free) symmetric Polly Cracker encryption scheme. We present a family of schemes parameterised not only by the underlying computational polynomial ring scheme $\mathcal{P}$, but also by a Gröbner basis generation algorithm, which itself depends on a degree bound d , and a second degree bound b . Our parameterised scheme, which we write as $\mathcal{SPC}_{\mathcal{P}, \text{GBGen}(\cdot), d, b}$, is presented in Figure 4. The message space is $P/\mathcal{I}$.

Correctness of evaluation can be verified by a straight-forward calculation. This scheme is not compact since multiplications square the size of the ciphertext.

$\text{Gen}_{\mathcal{P}, \text{GBGen}(\cdot), d, b}(1^\lambda):$	$\text{Enc}(m, \text{SK}):$	$\text{Dec}(c, \text{SK}):$	$\text{Eval}(c_0, \dots, c_{t-1}, C, \text{PK}):$
begin $P \leftarrow_{\$} \mathcal{P}_\lambda;$ $G \leftarrow_{\$} \text{GBGen}(1^\lambda, P, d);$ $\text{SK} \leftarrow (G, P, b);$ $\text{PK} \leftarrow (P, b);$ return $(\text{SK}, \text{PK});$ end	begin $f \leftarrow_{\$} P_{\leq b};$ $f' \leftarrow f \bmod G;$ $f \leftarrow f - f';$ $c \leftarrow m + f;$ return $c;$ end	begin $m \leftarrow c \bmod G;$ return $m;$ end	begin apply the Add and Mult gates of C over $P;$ return the result; end

Fig. 4. The (noise-free) Symmetric Polly Cracker scheme $\mathcal{SPC}_{\mathcal{P}, \text{GBGen}(\cdot), d, b}$

4.3 Security

We will show that the above scheme only achieves a weak version of chosen-plaintext security, which allows access to a limited number of ciphertexts.

Definition 5. The m -time IND-BCPA security of a (homomorphic) symmetric-key encryption scheme SKE is defined through a game $\text{IND-BCPA}_{m, \text{SKE}}$, which is similar to IND-CPA except that the adversary can query its encryption and left-or-right oracles a total of at most $m = m(\lambda)$ times. We say SKE is m -IND-BCPA secure if the advantage of any ppt adversary $\mathcal{A}$, defined as twice the probability of winning the game minus 1 is negligible.

Theorem 1. The scheme in Figure 4 is m -IND-BCPA secure iff the IM problem is hard.

See [1] for the proof. As a corollary, observe that when $m(\lambda) = \mathcal{O}(\lambda^b)$ one can construct an adversary which breaks the $\text{IND-BCPA}_{m, \text{SKE}}$ security of $\mathcal{SPC}$ in polynomial time. Thus we can only hope to achieve security in the bounded model for this scheme.

5 Symmetric-to-Asymmetric Conversion

Our goal for the rest of the paper is to convert the above scheme to one which is both fully IND-CPA secure and somewhat homomorphic. Once we achieve this, it is possible to construct a public-key scheme using the homomorphic features of the symmetric scheme by applying various generic conversions. In the literature there are two prominent such conversions:

- (A) Publish a set of encryptions of zero F_0 as part of the public key. To encrypt $m \in \{0, 1\}$ compute $c = \sum_{f_i \in S} f_i + m$ where S is a sparse subset of F_0 [13].
- (B) Publish two sets F_0 and F_1 of encryptions of zero and one as part of the public key. To encrypt $m \in \{0, 1\}$ compute $c = \sum_{f_i \in S_0} f_i + \sum_{f_j \in S_1} f_j$, with S_0 and S_1 being sparse subsets of F_0 and F_1 respectively such that the parity of $|S_1|$ is m . Decryption checks whether $\text{Dec}(c, \text{SK})$ is even or odd [27].

The security of the above transformations rests upon the (computational) indistinguishability of asymmetric ciphertexts from those produced directly using the symmetric encryption algorithm. As noted above, since $\mathcal{SPC}$ is not IND-CPA

secure the above transformations cannot be used. However, one could envisage a larger class of transformations which might lead to a fully secure additively homomorphic SKE (or equivalently an additively homomorphic PKE) scheme. In this section we rule out a large class of such transformations. To this end, we consider PKE schemes which lie within the following design methodology.

1. The secret key is the Gröbner basis G of a zero-dimensional ideal $\mathcal{I} \subset P$. The decryption algorithm computes $c \bmod \mathcal{I} = c \bmod G$ (perhaps together with some post-processing such as a mod 2 operation). Thus, the message space is (essentially) $P/\mathcal{I}$. We assume that $P/\mathcal{I}$ is known.
2. The public key consists of elements $f_i \in P$. We assume that the remainder of these elements modulo the ideal $\mathcal{I}$, i.e., $r_i := f_i \bmod \mathcal{I}$, are known.
3. A ciphertext is computed using ring operations. In other words, it can be expressed as $f = \sum_{i=0}^{N-1} h_i f_i + r$. Here f_i are as in the public key, h_i are some polynomials (possibility depending on f_i), and r is an encoding in $P/\mathcal{I}$ of the message.
4. The construction of the ciphertext does not encode knowledge of $\mathcal{I}$ beyond f_i . That is, we have $\left(\sum_{i=0}^{N-1} h_i f_i + r\right) \bmod \mathcal{I} = \sum_{i=0}^{N-1} h_i r_i + r$. Hence we have that $\left(\sum_{i=0}^{N-1} h_i r_i + r\right) \in P/\mathcal{I}$ as an element of P .
5. The security of the scheme relies on the fact that elements f produced at step (3) are computationally indistinguishable from random elements in $P_{\leq b}$.

Condition 4 imposes some real restrictions on the set of allowed transformation, but strikes a reasonable balance between allowing a general statement without ruling out too large a class of conversions. It requires that the r_i and r do not encode any information about the secret key. We currently require this restriction on the “expressive power” of r_i and r so as to make a general impossibility statement. If r_i and r produce a non-zero element in $\mathcal{I}$ using some arbitrary algorithm $\mathcal{A}$, we are unable to prove anything about the transformation. Furthermore, it is plausible that for any given $\mathcal{A}$ a similar impossibility result can be obtained if the remaining conditions hold.

Note that the two transformations above are special linear cases of this methodology. For transformation (A) we have that $f_i \in \mathcal{I}$ (hence $r_i = 0$), $h_i \in \{0, 1\}$ and $r = m$. For transformation (B) we have $r_i = 0$ if $f_i \in F_0$, $r_i = 1$ if $f_i \in F_1$, $h_i \in \{0, 1\}$, and $r = 0$.

To show that any conversion of the above form cannot lead to an IND-CPA-secure public-key scheme, we will use the following theorem which was also used in [5] to discourage the use of Gröbner bases for public-key schemes.

Theorem 2 ([12]). *Let $\mathcal{I} = \langle f_0, \dots, f_{m-1} \rangle$ be an ideal in the polynomial ring $P = \mathbb{F}[x_0, \dots, x_{n-1}]$, h be such that $\deg(h) \leq D$, and*

$$h - (h \bmod \mathcal{I}) = \sum_{i=0}^{m-1} h_i f_i, \text{ where } h_i \in P \text{ and } \deg(h_i f_i) \leq D.$$

Let G be the output of some Gröbner basis computation algorithm up to degree D . Then $h \bmod \mathcal{I}$ can be computed by polynomial reduction of h via G .

The main result of this section is a consequence of the above theorem. It essentially states that uniformly sampling elements of the ideal up to some degree is equivalent to compute a Gröbner basis for the ideal. Note that in itself Theorem 2 does not provide this result, since there is no assumption about the “quality” of h . Hence, to prove this result we first show that the above methodology implies sampling as in Theorem 2 but with uniformly random output. Theorem 2 then allows us to compute normal forms which (because of the randomness of h) allows the computation of a Gröbner basis by Lemma 1. The proof of Theorem 3 is given in [1].

Theorem 3. *Let $G = \{g_0, \dots, g_{s-1}\}$ be the reduced Gröbner basis of the zero-dimensional ideal $\mathcal{I}$ in the polynomial ring $P = \mathbb{F}[x_0, \dots, x_{n-1}]$ where each $\deg(g_i) \leq d$. Assume that $P/\mathcal{I}$ is known. Furthermore, let $F = \{f_0, \dots, f_{N-1}\}$ be a set of polynomials with known $r_i := f_i \bmod \mathcal{I}$. Let $\mathcal{A}$ be a ppt algorithm which given F produces elements $f = \sum h_i f_i + r$ with $\deg(f) \leq b$, $h_i \in P$, $b \leq B$, $\deg(h_i f_i) \leq B$, and $(f \bmod \mathcal{I}) = \sum h_i r_i + r$. Suppose further that the outputs of $\mathcal{A}$ are computationally indistinguishable from random elements in $P_{\leq b}$. Then there exists an algorithm which computes a Gröbner basis for $\mathcal{I}$ from F in $\mathcal{O}(n^{3B})$ field operations.*

Therefore, if for some degree $b \geq d$ computationally uniform elements of $P_{\leq b}$ can be produced using the public key $f_0, \dots, f_{N-1}$, there is an attacker which recovers the secret key $g_0, \dots, g_{s-1}$ in essentially the same complexity. Hence, while conceptually simple and provably secure up to some bound, our symmetric Polly Cracker scheme $\mathcal{SPC}_{\mathcal{P}, \text{GBGen}(\cdot), d, b}$ does not provide a valid building block for constructing a fully homomorphic public-key encryption scheme.

REMARK. Although the above impossibility result is presented for public-key encryption schemes, due to the equivalence result of [27], it also rules out the existence of additively homomorphic symmetric PC-style schemes with full IND-CPA security.

6 Gröbner Bases with Noise

In this section, we introduce noisy variants of the problems presented in Section 3. The goal is to lift the restriction on the number of samples that the adversary can obtain, and following a similar design methodology to Polly Cracker, construct an IND-CPA-secure scheme. That is, we consider problems which naturally arise if we consider noisy encoding of messages in $\mathcal{SPC}$. Similarly to [13,26] we expect a problem which is efficiently solvable in the noise-free setting to be hard in the noisy setting. We will justify this assumption in Section 6.1 by arguing that our construction can be seen as a generalisation of [13,26]. The games below will be parameterised by a noise distribution. The discrete Gaussian distribution – denoted for $\chi_{\alpha, q}$ for standard deviation αq and modulus q – is of particular interest to us (cf. [25]).

We now define a noisy variant of the Gröbner basis problem. The task here is still to compute a Gröbner basis for some ideal $\mathcal{I}$. However, we are now only

given access to a noisy sample oracle which provides polynomials which are not necessarily in $\mathcal{I}$ but rather are “close” approximations to elements of $\mathcal{I}$. Here the term “close” is made precise using a noise distribution χ on $P/\mathcal{I}$.

Definition 6. *The Gröbner basis with noise problem is defined through the game $\text{GBN}_{\mathcal{P}, \text{GBGen}(\cdot), d, b, \chi}$ as shown in Figure 5. The advantage of a ppt algorithm $\mathcal{A}$ in solving the GBN problem is the probability of winning the game.*

Initialize ($1^\lambda, \mathcal{P}, d$):	Sample (·):	Finalize (G'):
begin	begin	begin
$P \leftarrow_{\$} \mathcal{P}_\lambda;$	$f \leftarrow_{\$} P_{\leq b};$	$\text{return } (G = G');$
$G \leftarrow_{\$} \text{GBGen}(1^\lambda, P, d);$	$e \leftarrow_{\$} \chi;$	end
return $(1^\lambda, P);$	$f \leftarrow f - (f \bmod G) + e;$	
end	return $f;$	
	end	

Fig. 5. Game $\text{GBN}_{\mathcal{P}, \text{GBGen}(\cdot), d, b, \chi}$

The essential difference between the noisy and noise-free versions of the GB problem is that by adding noise we have eliminated the restriction on the adversary to call the **Sample** oracle a bounded number of times. The choice of χ greatly influences the hardness of the GBN problem.

As in the noise-free setting, we can ask various questions about the ideal $\mathcal{I}$ spanned by G . One such example is solving the ideal remainder problem with access to noisy samples from $\mathcal{I}$.

Definition 7. *The ideal remainder with noise problem is defined through the game $\text{IRN}_{\mathcal{P}, \text{GBGen}(\cdot), d, b, \chi}$ as shown in Figure 6. The advantage of a ppt algorithm $\mathcal{A}$ is defined as the probability of winning the game minus $1/q(\lambda)^{\dim_{\mathbb{F}}(P/\langle G \rangle)}$.*

Initialize ($1^\lambda, \mathcal{P}, d$):	Sample (·):	Challenge (·):	Finalize (r'):
begin	begin	begin	begin
$P \leftarrow_{\$} \mathcal{P}_\lambda;$	$f \leftarrow_{\$} P_{\leq b};$	$f \leftarrow_{\$} P_{\leq b};$	$r'' = f \bmod G;$
$G \leftarrow_{\$} \text{GBGen}(1^\lambda, P, d);$	$e \leftarrow_{\$} \chi;$	return $f;$	return $r' = r'';$
return $(1^\lambda, P);$	$f \leftarrow f - (f \bmod G) + e;$	end	end
end	return $f;$		
	end		

Fig. 6. Game $\text{IRN}_{\mathcal{P}, \text{GBGen}(\cdot), d, b, \chi}$

In fact, the above two problems are equivalent as shown in the lemma below. Compared to the noise-free version, we no longer need the IM adversary to be overwhelmingly successful, as there are no restrictions on the number of calls that can be made to the **Sample** procedure. The proof is given in [1].

Lemma 3. *The IRN problem is hard iff the GBN problem is hard.*

Similarly to the noise-free setting, the ideal membership with noise (IMN) problem is the decisional variant of the IRN (and hence the GBN) problem. However,

in the noisy setting we have the choice between a noisy and noise-free challenge polynomial. In the definition below noisy challenges are provided and the adversary wins the game if he can distinguish whether an element was sampled uniformly from $P_{\leq b}$ or from $\mathcal{I} + \chi$.

Definition 8. *The ideal membership with noise problem is defined through the game $\text{IMN}_{\mathcal{P}, \text{GBGen}(\cdot), d, b, \chi}$ as shown in Figure 7. The advantage of a ppt algorithm A in solving the IMN problem is as twice the probability of winning the game minus 1.*

<u>Initialize($1^\lambda, \mathcal{P}, d$):</u>	<u>Sample():</u>	<u>Challenge():</u>	<u>Finalize(c'):</u>
begin $P \leftarrow_{\$} \mathcal{P}_\lambda;$ $G \leftarrow_{\$} \text{GBGen}(1^\lambda, \mathcal{P}, d);$ $c \leftarrow_{\$} \{0, 1\};$ return $(1^\lambda, P);$ end	begin $f \leftarrow_{\$} P_{\leq b};$ $e \leftarrow_{\$} \chi;$ $f' \leftarrow f \bmod G;$ $f \leftarrow f - f' + e;$ return $f;$ end	begin $f \leftarrow_{\$} P_{\leq b};$ if $c = 1$ then $e \leftarrow_{\$} \chi;$ $f \leftarrow f - (f \bmod G) + e;$ return $f;$ end	begin return $(c' = c);$ end

Fig. 7. Game $\text{IMN}_{\mathcal{P}, \text{GBGen}(\cdot), d, b, \chi}$

Our definition of the IMN problem can be seen as an instantiation of Gentry’s ideal coset problem [18] since both problems require distinguishing uniformly chosen elements in $P_{\leq b}$ from those in $\mathcal{I} + \chi$. Our problem, however, assumes noisy samples since it is clear from Section 3 that otherwise the problem is easy.

Again, we would like to have a decision-to-search reduction; that is, we would like to have an equivalence between the IRN and IMN problems. This equivalence holds when the search space of remainders is polynomial in λ , namely when $q(\lambda)^{\dim_{\mathbb{F}_q}(\mathcal{P}(\lambda)/\text{GBGen}(\cdot))} = \text{poly}(\lambda)$. The intuition behind this reduction is that the adversary can exhaustively search the quotient ring and use the IMN oracle to verify his guess. Once again, a technical difficulty arises as the adversary does not know the search space $P/\mathcal{I}$ and thus has to discover it during the attack. Again, the IMN adversary provides an oracle to accomplish this. This is formalised in the lemma below whose proof is in [1].

Lemma 4. *The IMN problem is hard iff the IRN problem is hard for poly-sized $q^{\dim_{\mathbb{F}_q}(P/\langle G \rangle)}$.*

Hence GBN is equivalent to IRN and IRN is equivalent to IMN under some additional assumptions about the size $P/\mathcal{I}$. Finally, for $d = 1$ (but arbitrarily b) we show that if we can solve the GBN problem on average, then we can also solve it for worst-case instances. This in turn increases our confidence in hardness of the GBN problem. The proof of the follow lemma is given in [1].

Lemma 5. *If the GBN problem is worst-case hard, then it is also average-case hard.*

6.1 Hardness Assumptions and Justifications

Let us now investigate the hardness of the GBN, IRN, and IMN problems.

RELATION TO LWE. It is easy to see that GBN can be considered as a non-linear generalisation of LWE if $q = \text{poly}(n)$ is a prime. In other words, we have equivalence between these problems when $b = d = 1$ in GBN. This is formalised below (proof is in [1]).

Lemma 6. *If the LWE problem is hard then the GBN problem is also hard for $b = d = 1$.*

In the noise-free setting we assume that solving systems of equations of degree greater than 1 is harder than solving those of degree 1. More generally, we assume that equations of degree $b > b'$ are harder to solve than those of degree b' . Intuitively, equations of degree b' can be seen as those of degree b where the coefficients of higher degree monomials are set to zero. However, formalising this intuition for an adversary which expects uniformly distributed equations of degree b seems futile since producing such equations is equivalent to solving the system by Theorem 3.

In the noisy setting this equivalence (i.e., Theorem 3) between sampling and solving no longer holds. However, we still need to deal with the distribution of noise. One strategy to show that difficulty increases with the degree parameter b is to allow for an increase of the noise level in the samples. We formalise this below (a proof is given in [1]).

Lemma 7. *If the GBN problem is hard for degree $2b$ with noise $\chi_{\sqrt{N}\alpha^2q,q}$, $N = \binom{n+b}{b}$, then it is also hard for degree b with noise $\chi_{\alpha,q}$.*

RELATION TO THE APPROXIMATE GCD PROBLEM. The GBN problem for $n = 1$ is the approximate GCD problem over $\mathbb{F}_q[x]$. Contrary to the approximate GCD problem over the integers (cf. [13]), this problem has not yet received much attention, and hence it is unclear under which parameters it is hard. However, as we discuss in [1], the notion of a Gröbner basis can be extended to $\mathbb{Z}[x_0, \dots, x_{n-1}]$, which in turn implies a version of the GBN problem over $\mathbb{Z}$. This can be seen as a direct generalisation of the approximate GCD problem in $\mathbb{Z}$.

THE CASE $q = 2$. Recall that if $b = d = 1$ we have an equivalence with the LWE problem (or the well-known problem of learning parity with noise (LPN) if $q = 2$). More generally, for $d = 1$ we can reduce Max-3SAT instances to GBN instances by translating each clause individually to a Boolean polynomial. However, in Max-3SAT the number of samples is bounded and hence this reduction only shows the hardness of GBN with a bounded number of samples. Still, the Gröbner basis returned by an arbitrary algorithm $\mathcal{A}$ solving GBN using a bounded number of samples will provide a solution to the Max-3SAT problem. Vice versa, we may convert a GBN instance for $d = 1$ to a Max-SAT instance (more precisely Partial Max-Sat) by running an ANF to CNF conversion algorithm [4].

KNOWN ATTACKS. Finally, we consider known attacks to understand the difficulty of the GBN problem. Recall that if $b = 1$ Lemma 6 states that we can solve the LWE problem if we can solve the GBN problem. The converse also applies. Indeed, for any $b \geq d$ and $d = 1$ the best known attack against the

GBN problem for $d = 1$ is to reduce it to the LWE problem, similarly to the linearisation technique used for solving non-linear systems of equations in the noise-free setting. Let $N = \binom{n+b}{b}$ be the number of monomials up to degree b . Let $\mathcal{M} : P \rightarrow \mathbb{F}_q^N$ be a function which maps polynomials in P to vectors in $\mathbb{F}_q^N$ by assigning the i -th component of the image vector the coefficient of the i -th monomial $\in M_{\leq b}$. Then, in order to reduce GBN with n variables and degree b to LWE with N variables, reply to each LWE **Sample** query by calling the GBN **Sample** oracle to retrieve f , compute $v = \mathcal{M}(f)$ and return (a, b) with $a = (v_{N-1}, \dots, v_1)$ and $b = -v_0$. When the LWE adversary queries **Finalize** on s , query the GBN **Finalize** on $[x_0 - s_0, \dots, x_{n-1} - s_{n-1}]$. Correctness follows from the correctness of linearisation in the noise-free setting [3]. Furthermore, the LWE problem in N variables and with respect to the discrete Gaussian noise distribution $\chi_{\alpha, q}$ is considered to be hard if $\alpha \geq 3/2 \cdot \max(\frac{1}{q}, 2^{-2\sqrt{N \log q \log d}})$ for an appropriate choice of δ which is the quality of the approximation for the shortest vector problem. With current lattice algorithms $\delta = 1.01$ is hard and 1.005 infeasible [24].

Perhaps the most interesting attack on LWE from the perspective of this work is that due to Arora and Ge [3] which reduces the problem of solving linear systems with noise to the problem of solving (structured) non-linear noise-free systems. We may apply this technique directly to GBN, i.e., without going to LWE first, and reduce it to GB with large b . However, it seems this approach does not improve the asymptotic complexity of the attack. Finally, certain conditions to rule out exhaustive search must be imposed.

Definition 9. *Let $b, d \in \mathbb{N}$ with $b \geq d \geq 1$. Let $\mathcal{P}$ be a polynomial ring distribution and $\chi_{\alpha, q}$ be the discrete Gaussian distribution. Suppose the parameters n , α , and q (all being a function of λ) satisfy the following set of conditions: 1) $n \geq \sqrt[b]{\lambda}$; 2) $(\alpha q)^{nd^n} \approx 2^\lambda$ so exhaustive search over the noise or the secret key space is ruled out; 3) $\alpha q \geq 8$ as suggested in [22]; and 4) for $N := \binom{n+b}{b}$, and $\delta := 1.005$ we have $\alpha \geq 3/2 \cdot \max\{\frac{1}{q}, 2^{-2\sqrt{N \log q \log \delta}}\}$, and hence the best known attacks against the LWE problem are ruled out [24, 28]. Then the advantage of any ppt algorithm in solving the GBN, IRN, and IMN problems is negligible.*

7 Polly Cracker with Noise

In this section we present a fully IND-CPA-secure PC-style symmetric encryption scheme. Our parameterised scheme, $\mathcal{SPCN}_{\mathcal{P}, \text{GBGen}(\cdot, d, b, \chi)}$, is shown in Figure 8. Here we represent elements in $\mathbb{F}_q$ as integers in the interval $(-\lfloor \frac{q}{2} \rfloor, \lfloor \frac{q}{2} \rfloor]$. This representation is also used in the definition of noise. All the computations are performed in the ring P as generated by **Gen**. Furthermore we assume that $\gcd(2, q) = 1$. This condition is needed for the correctness and the security of our scheme. The message space is $\mathbb{F}_2$ (although we remark that this can be generalised to other small fields). Correctness of evaluation up to overflows can be established by a straight-forward calculation.

PERMITTED CIRCUITS. Circuits composed of **Add** and **Mul** gates can be seen as multivariate Boolean polynomials in t variables over $\mathbb{F}_2$. We can consider the

$\text{Gen}_{\mathcal{P}, \text{GBGen}(\cdot), d, b, \chi}(1^\lambda):$	$\text{Enc}(m, \text{SK}):$	$\text{Dec}(c, \text{SK}):$	$\text{Eval}(c_0, \dots, c_{t-1}, C, \text{PK}):$
begin $P \leftarrow_{\$} \mathbf{P}_\lambda;$ $G \leftarrow_{\$} \text{GBGen}(1^\lambda, P, d);$ $\text{SK} \leftarrow (G, P, b, \chi);$ $\text{PK} \leftarrow (P, b, \chi);$ return (SK, PK); end	begin $f \leftarrow_{\$} P \equiv b;$ $f' \leftarrow f \bmod G;$ $f \leftarrow f - f';$ $e \leftarrow_{\$} \chi;$ $c \leftarrow f + 2e + m;$ return $c;$ end	begin $m' \leftarrow c \bmod G;$ $m \leftarrow m' \bmod 2;$ end	begin apply Add and Mul gates of C over $P;$ return the result; end

Fig. 8. The Symmetric Polly Cracker with Noise scheme $\mathcal{SPCN}_{\mathcal{P}, \text{GBGen}(\cdot), d, b, \chi}$

generalisation of this set of polynomials to $\mathbb{F}_q$ (i.e., the coefficients are in $\mathbb{F}_q$). In order to define the set of permitted circuits (which will be parameterised by $\alpha > 0$) we first embed the Boolean polynomials into the ring of polynomials over $\mathbb{Z}$. For $\chi_{\alpha, q}$ we have that the probability of the noise being larger than $k\alpha q$ is $< \exp(-k^2/2)$. We now say that a circuit is valid if for any $(s_0, \dots, s_{t-1})$ with $s_i \leq t\alpha q$ we have that the outputs are less than q for some parameter t . This restriction ensures that no overflows occur when polynomials are evaluated over $\mathbb{F}_q$. In [1] we discuss how to set α and q in order to allow for evaluation of polynomials of some fixed degree μ and provide a Sage implementation [30].

COMPACTNESS. Additions do not increase the size of the ciphertext, but they do increase the size of the error by at most one bit. Multiplications square the size of the ciphertext and the bit-size of the the noise by approximately $\log(5e_0e_1)$ bits. In [1] we also provide a discussion on how to trade ciphertext size with noise, an avenue which is investigated independently in [7]. The theorem below, which is proven in [1], states the security properties of the above scheme.

Theorem 4. *If the IMN problem is hard, then the scheme in Figure 8 is secure.*

The above theorem together with the recent results in [27] which establish the equivalence of symmetric and asymmetric homomorphic encryption schemes leads to the first provably secure public-key encryption scheme from assumptions related to Gröbner bases for random systems. This provides a positive answer to the challenges raised by Barkee et al. [5] (and later also by Gentry [18]). We note here that the transformation – as briefly described in Section 5 – only use the additive features of the scheme and does not require full homomorphicity.

Acknowledgments. We would like to thank Carlos Cid for valuable feedback and discussions on this work. We would also like to thank Frederik Armknecht for helpful discussions on an earlier draft of this work.

References

1. Albrecht, M.R., Farshim, P., Faugère, J.-C., Perret, L.: Polly Cracker, revisited. Cryptology ePrint Archive, Report 2011/289 (2011)
2. Applebaum, B., Cash, D., Peikert, C., Sahai, A.: Fast Cryptographic Primitives and Circular-Secure Encryption Based on Hard Learning Problems. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 595–618. Springer, Heidelberg (2009)

3. Arora, S., Ge, R.: New Algorithms for Learning in Presence of Errors. In: Aceto, L., Henzinger, M., Sgall, J. (eds.) ICALP 2011, Part I. LNCS, vol. 6755, pp. 403–415. Springer, Heidelberg (2011)
4. Bard, G.V., Courtois, N.T., Jefferson, C.: Efficient methods for conversion and solution of sparse systems of low-degree multivariate polynomials over $\text{GF}(2)$ via SAT-solvers. Cryptology ePrint Archive, Report 2007/024 (2007)
5. Barkee, B., Can, D.C., Ecks, J., Moriarty, T., Ree, R.F.: Why you cannot even hope to use Gröbner bases in public key cryptography: An open letter to a scientist who failed and a challenge to those who have not yet failed. *J. of Symbolic Computations* 18(6), 497–501 (1994)
6. Berbain, C., Gilbert, H., Patarin, J.: QUAD: A multivariate stream cipher with provable security. *J. Symb. Comput.* 44(12), 1703–1723 (2009)
7. Brakerski, Z., Vaikuntanathan, V.: Efficient fully homomorphic encryption from (standard) LWE. To appear in FOCS 2011 (2011)
8. Buchberger, B.: Ein Algorithmus zum Auffinden der Basiselemente des Restklassenrings nach einem nulldimensionalen Polynomideal. PhD thesis, Universität Innsbruck (1965)
9. Caboara, M., Caruso, F., Traverso, C.: Lattice Polly Cracker cryptosystems. *Journal of Symbolic Computation* 46, 534–549 (2011)
10. Cox, D., Little, J., O’Shea, D.: *Ideals, Varieties, and Algorithms*, 3rd edn. Springer, Heidelberg (2005)
11. Cramer, R., Shoup, V.: Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack. *SIAM J. of Computing*, 167–226 (2003)
12. Dickenstein, A., Fitchas, N., Giusti, M., Sessa, C.: The membership problem for unmixed polynomial ideals is solvable in single exponential time. *Discrete Appl. Math.* 33(1-3), 73–94 (1991)
13. van Dijk, M., Gentry, C., Halevi, S., Vaikuntanathan, V.: Fully Homomorphic Encryption Over the Integers. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 24–43. Springer, Heidelberg (2010)
14. Ding, J., Yang, B.-Y.: Multivariate public key cryptography. In: *Post-Quantum Cryptography*, pp. 193–234. Springer, Heidelberg (2009)
15. dit Vehel, F.L., Marinari, M.G., Perret, L., Traverso, C.: A survey on Polly Cracker systems. In: *Gröbner Bases. Coding and Cryptography*, pp. 285–305. Springer, Heidelberg (2009)
16. Faugère, J.-C., Gianni, P.M., Lazard, D., Mora, T.: Efficient computation of zero-dimensional Gröbner bases by change of ordering. *Journal of Symbolic Computation* 16, 329–344 (1993)
17. Fellows, M., Koblitz, N.: Combinatorial cryptosystems galore! In: *Finite Fields: Theory, Applications, and Algorithms. Contemporary Mathematics*, vol. 168, pp. 51–61. AMS (1994)
18. Gentry, C.: A fully homomorphic encryption scheme. PhD thesis, Stanford University (2009)
19. Gentry, C.: Fully homomorphic encryption using ideal lattices. In: *ACM Symposium on Theory of Computing*, pp. 169–178 (2009)
20. Gentry, C., Halevi, S.: Implementing Gentry’s Fully-Homomorphic Encryption Scheme. In: Paterson, K.G. (ed.) EUROCRYPT 2011. LNCS, vol. 6632, pp. 129–148. Springer, Heidelberg (2011)
21. Gouget, A., Patarin, J.: Probabilistic Multivariate Cryptography. In: Nguyễn, P.Q. (ed.) VIETCRYPT 2006. LNCS, vol. 4341, pp. 1–18. Springer, Heidelberg (2006)

22. Lindner, R., Peikert, C.: Better Key Sizes (and Attacks) for LWE-Based Encryption. In: Kiayias, A. (ed.) CT-RSA 2011. LNCS, vol. 6558, pp. 319–339. Springer, Heidelberg (2011)
23. Melchor, C.A., Gaborit, P., Herranz, J.: Additively Homomorphic Encryption with d -Operand Multiplications. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 138–154. Springer, Heidelberg (2010)
24. Micciancio, D., Regev, O.: Lattice-based cryptography. In: Post-Quantum Cryptography, pp. 147–191. Springer, Heidelberg (2009)
25. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM* 56, 34:1–34:40 (2009)
26. Regev, O.: The learning with errors problem. In: IEEE Conference on Computational Complexity 2010, pp. 191–204 (2010)
27. Rothblum, R.: Homomorphic Encryption: From Private-Key to Public-Key. In: Ishai, Y. (ed.) TCC 2011. LNCS, vol. 6597, pp. 219–234. Springer, Heidelberg (2011)
28. Rückert, M., Schneider, M.: Estimating the security of lattice-based cryptosystems. *Cryptology ePrint Archive*, Report 2010/137 (2010)
29. Smart, N.P., Vercauteren, F.: Fully Homomorphic Encryption with Relatively Small Key and Ciphertext Sizes. In: Nguyen, P.Q., Pointcheval, D. (eds.) PKC 2010. LNCS, vol. 6056, pp. 420–443. Springer, Heidelberg (2010)
30. Stein, W.A., et al.: Sage Mathematics Software. The Sage Development Team, Version 4.7.0 (2011), <http://www.sagemath.org>