Polly Cracker, Revisited, Revisited

Gottfried Herold*

Ruhr-Universität Bochum
Horst Görtz Insitute for IT-Security
Bochum, Germany
gottfried.herold@rub.de

Abstract. In this paper, we consider the Polly Cracker with Noise (PCN) cryptosystem by Albrecht, Farshim, Faugère, and Perret (Asiacrypt 2011), which is a public-key cryptosystem based on the hardness of computing Gröbner bases for noisy random systems of multivariate equations. We examine four settings, covering all possible parameter ranges of PCN with zero-degree noise. In the first setting, the PCN cryptosystem is known to be equivalent to Regev's LWE-based scheme. In the second, it is known to be at most as secure as Regev's scheme. We show that for one other settings it is equivalent to a variants of Regev's with less efficiency and in the last setting it is completely insecure and we give an efficient key-recovery attack. Unrelated to the attack, we also fix some flaws in the security proofs of PCN.

Keywords: Polly Cracker with Noise, Learning with Errors, Gröbner bases, Cryptanalysis.

1 Introduction

Background. By the term Polly Cracker-type cryptosystem, we mean a family of cryptosystems starting from the early 1990s that propose to base their security on the difficulty of computing Gröbner bases ([8,2]). In its public key version and the most simple form, the public key is an ideal I in a polynomial ring (given by sufficiently many polynomials of degree b from I) and the secret key is a Gröbner basis for I consisting of polynomials of degree $d \leq b$. These systems mostly lack a formal treatment of security and almost all of them have been broken due fundamental limitations in the construction([2,1]). See [7] for a good survey on various instantiations and attacks.

Recently, at Asiacrypt 2011, Albrecht, Farshim, Faugère, and Perret [1] proposed a new cryptosystem called Polly Cracker with Noise (PCN) that tries to overcome these limitations. Their cryptosystem can be seen both as a high-dimensional generalization of Regev's LWE-based scheme [12] and a noisy generalization of the Polly Cracker-style cryptosystems. They also give a formal proof of security, based on the hardness of computational problems related to Gröbner

* Due to space limitations, this version does not contain the proofs of Thm. 3. These are contained in the full version, available on eprint.

M. Fischlin, J. Buchmann, and M. Manulis (Eds.): PKC 2012, LNCS 7293, pp. 17–33, 2012.

bases and ideals in multivariate polynomial rings. Note that this paper refers mainly to the full version of [1] on eprint, which contains more material than the proceedings version.

One of the appealing features of the PCN cryptosystem comes from its ideal-theoretic framework. In this framework it is prominently visible that the PCN cryptosystem, which contains LWE as a special case, is both multiplicatively and additively homomorphic for a limited number of operations. For the special case of LWE, the recent fully homomorphic scheme by Brakerski and Vaikuntanathan from FOCS 2011 [4] can be represented in this framework.

Our Contributions. Our first result is that the Polly Cracker with Noise cryptosystem with zero-degree noise is either insecure or does not offer any security benefit (although still a conceptual one) compared to Regev's scheme. For $b > d > 1$, we present an efficient attack that recovers the secret key from the public key. For $d = 1$, the security of the PCN cryptosystem is at most that of Regev's scheme by [1]. For $d = b > 1$, PCN has the same security as Regev's schem, but with less efficiency. The only remaining case $b = d = 1$ is exactly Regev's scheme by [1].

Note that zero-degree noise is used for the homomorphic properties claimed in [1], cf. Sect. 2.3.

As a second result, we point out flaws in the security proofs of [1], giving counterexamples to the statements claimed therein. We then give corrected proofs for $d = 1$, thereby showing their security proofs only work for $d = 1$. Note that the attack against $b > d > 1$ is unrelated to these flaws. Due to space limitations, the proofs are only contained in the full version, available on eprint.

Organization of this Work. This work is organized as follows: In Section 2, we start by introducing some notation and recalling the Polly Cracker with Noise cryptosystem and its security assumptions. In Section 3, we relate the PCN cryptosystem to Regev's scheme for $b = d$ and for $d = 1$.

In Section 4, we give counterexamples to the security proofs of [1] and give corrected statements for $d = 1$.

In Section 5, we present our key-recovery attack for $b > d > 1$.

2 The Polly Cracker with Noise Cryptosystem

2.1 Gröbner Bases

In this section, we introduce some notation and recall some facts regarding Gröbner bases [5]. For a more detailed exposition, see e.g. [6].

Let $P = \mathbb{F}_q[X_1, X_2, \ldots, X_n]$ be a polynomial ring and $<$ be a fixed monomial ordering for its monomials. For a subspace $Q \subset P$, we denote by $Q_{<k}, Q_{=k}, Q_{\leq k}$ the restriction of Q to polynomials of total degree $< k, = k, \leq k$, respectively. We shall always assume that q is odd, for simplicity prime, and that the monomial ordering is compatible with the total degree of monomials (e.g. `deglex` or `degrevlex`), i.e. $\deg f < \deg g$ implies $f < g$ for all monomials $f, g \in P$,

where deg denotes total degree. W.l.o.g. we may assume $X_1 < X_2 < \ldots < X_n$. For a polynomial $f \in P$, let $\mathrm{LC}(f), \mathrm{LM}(f), \mathrm{LT}(f)$ denote the leading coefficient, monomial and term, respectively. We always represent polynomials $f \in P, f = \sum_{m \leq \mathrm{LM}(f)} f_m \cdot m$ by their dense coefficient representation, i.e. the list of the f_m. Note that for degree-compatible $<$, the length of this list is at most $\dim P_{\leq \deg f} = \binom{n+\deg f}{\deg f}$, which is polynomial in n for fixed $\deg f$.

Definition 1. Gröbner basis
Let $I \subset P$ be an ideal. A finite set $G = \{g_1, \ldots, g_l\}$ is called a Gröbner basis for I if G generates I as an ideal and if for every $f \in I$, there is a $g_i \in G$ such that $\mathrm{LM}(g_i) \mid \mathrm{LM}(f)$.

If additionally, $\mathrm{LC}(g_i) = 1$ *for all i and no term of g_i is divisible by* $\mathrm{LC}(g_j)$ *for $i \neq j$, we call G a* reduced *Gröbner basis.*

Every ideal $I \subset P$ has a Gröbner basis G. If one additionally insists on G being reduced, G is unique. For any $f \in P$, we can use the multivariate polynomial division algorithm to compute the remainder, denoted $f \bmod G$. The central property of a Gröbner basis G is that $f \bmod G$ is unique. We use this property to identify P/I with the set of remainders, thus viewing $P/I \subset P$. As a vector space, P/I is generated by those monomials not divisible by any $\mathrm{LM}(g_i)$ and we always have $P = (P/I) \oplus I$.

2.2 Polly Cracker with Noise

In this section, we briefly recall the (symmetric key variant of the) Polly Cracker with Noise(PCN) cryptosystem.

The secret key of this cryptosystem is a Gröbner basis G for some ideal $I \subset P$. Ciphertexts are noisy samples from I, where the message is appropriately embedded in the noise. More precisely, we encrypt a message bit $M \in \{0, 1\}$ as $f + 2e + M$, where $f \leftarrow_\$ I$ and $e \leftarrow_\$ \mathcal{X}$ from some noise distribution $\mathcal{X}$ on P/I. We can decrypt c by computing $M = (c \bmod G) \bmod 2$, provided the noise e is small enough.

In more detail, let us consider $P = \mathbb{F}_q[X_1, \ldots, X_n]$ and $<$ as above. We will also need to fix some integers $0 < d \leq b$, which will denote the degree of the Gröbner basis polynomials and the message polynomials, respectively. The parameters $q = q(\lambda), n = n(\lambda)$ will be implicitly functions of the security parameter λ, with $\log q = \mathrm{poly}(\lambda)$ (sometimes even $q = \mathrm{poly}(\lambda)$), $n = \mathrm{poly}(\lambda)$ and $n^d = \Omega(\lambda)$ (so $\mathrm{poly}(n) = \mathrm{poly}(\lambda)$). Note that we assume b, d *not* to depend on the security parameter λ.

The secret key of our cryptosystem will be a (reduced) Gröbner basis $G = \{g_1, \ldots, g_n\}$ for some ideal I, so we need an algorithm to generate Gröbner bases. In general, we require $\mathbf{Gen}(1^\lambda)$ to be a ppt algorithm outputting a reduced Gröbner basis $G = \{g_1, \ldots, g_k\}$ for an ideal $I \subsetneq P$ with $\deg g_i \leq d$.

For definiteness, we will restrict our attention in Sect. 5 to the key generation algorithm suggested in [1] called $\mathrm{GBGen}_{\mathrm{dense}}$.

Algorithm 1. $\mathrm{GBGen}_{\mathrm{dense}}$

function $\mathrm{GBGen}_{\mathrm{dense}}(1^\lambda)$:
  for $i = 1$ to n **do**
    $g_i \leftarrow X_i^d$
    for all monomials $m \in P_{\leq d}$ with $m < X_i^d$ and $m \neq X_j^d$ for any j **do**
      $g_{i,m} \leftarrow_\$ F_q$ uniformly
      $g_i \leftarrow g_i + g_{i,m} \cdot m$
  return $G = \{g_1, \ldots, g_n\}$

Writing each g_i as $g_i = \sum_m g_{i,m} \cdot m$ where m runs over the possible monomials of P, $\mathrm{GBGen}_{\mathrm{dense}}$ sets the leading term of g_i to be X_i^d. The coefficients of smaller monomials are chosen uniformly and independently at random.

Buchberger's first criterion (cf. [3, Lemma 5.66, p. 222] or [6, section 2.9, pp. 99–108]) guarantees that this is indeed a Gröbner basis for its generated ideal $I = (g_1, \ldots, g_n)$. Setting all coefficients of X_j^d in g_i to be 0 for $i \neq j$ guarantees that G is a *reduced* Gröbner basis. Note that sampling these coefficients at random as well and then reducing the Gröbner basis afterward, as originally done in [1], gives the same output distribution.

We denote by $\mathcal{Q} = P/I$ the quotient ring and identify it with a subspace $\mathcal{Q} \subset P$ as above, such that $P = I \oplus \mathcal{Q}$.

With G generated by $\mathrm{GBGen}_{\mathrm{dense}}$, $\mathcal{Q}$ is always finite-dimensional and a basis is given by $\left\{X_1^{t_1} \cdots X_n^{t_n} \mid t_i < d\right\}$. Note that this does not depend on the randomness of $\mathrm{GBGen}_{\mathrm{dense}}$ and for simplicity we shall always assume that $\mathcal{Q} \subset P$ is finite-dimensional and a basis for $\mathcal{Q}_{\leq b}$ is publicly known, even for general **Gen**. It follows that for $d = 1$, $\mathcal{Q} = \mathbb{F}_q$ is just the field of constants in P. In the case $d > 1$, the full quotient $\mathcal{Q}$ has exponential dimension $\dim \mathcal{Q} = d^n$, essentially due to the lack of a fixed bound on *total* degree. In this case, our cryptosystem will only make use of the polynomially-dimensional subspace $\mathcal{Q}_{\leq b} \subset \mathcal{Q}$.

Let $\mathcal{X}$ be an efficiently sampleable noise distribution on $\mathcal{Q}_{\leq b}$. The distributions we will later be concerned with will be either uniform or discrete Gaussian distributions on vector sub-spaces. In the case of Gaussians, this will mean we independently sample each coefficient of $e \leftarrow \mathcal{X}$ in a particular basis from a discrete Gaussian distribution.

By the *support* S of a probability distribution Φ on a finite set Ω, we mean those elements of Ω that are assigned a non-zero probability by Φ.

Using the Gröbner basis G for I, we can obtain noisy samples from $I_{\leq b} + \mathcal{X}$ by applying algorithm 2.[1]

By SampleI without subscript, we denote the special case of noiseless sampling from $I_{\leq b}$ (i.e. with $e = 0$ above).

Following [1], we note that SampleI actually samples uniformly from $I_{\leq b}$ and also give an alternative sampling algorithm, whose equivalence we will need later on:

[1] Identifying a set with the uniform distribution on it, $I_{\leq b} + \mathcal{X}$ actually is the output distribution of the algorithm.

Algorithm 2. $\text{SampleI}_{\mathcal{X}}$

1: **function** $\text{SampleI}_{\mathcal{X}}$(G,b):
2: $f \leftarrow_{\$} P_{\leq b}$ uniformly
3: $e \leftarrow_{\$} \mathcal{X}$
4: $f := f - (f \bmod G) + e$
5: **return** f.

Lemma 1. *For any Gröbner basis $G = (g_1, \ldots, g_m)$ for I,* $\text{SampleI}(G, b)$ *yields uniform samples from $I_{\leq b}$.*

Furthermore, if $\deg g_i = d_i \leq b$ *for all g_i and the underlying monomial ordering is compatible with* deg*, we have the following alternative sampling algorithm, which gives the same distribution:*

Let $t_i \leftarrow_{\$} P_{\leq b-d_i}$ uniformly for $i \in \{1, \ldots, m\}$ and sample $f \in I$ as $f = \sum_{i=1}^{m} t_i \cdot g_i$.

Proof. Clearly, both ways of sampling give us polynomials from $I_{\leq b}$. We observe that both $f \mapsto f \bmod G$ and $(t_1, \ldots t_m) \mapsto f = \sum_{i=1}^{m} t_i \cdot g_i$ are $\mathbb{F}_q$-linear maps. Since surjective linear maps preserve uniform distributions, both resulting distributions are uniform on their respective supports.

For SampleI, the support is clearly all of $I_{\leq b}$, since we may choose any element from $I_{\leq b}$ in step 2 of the algorithm.

For the alternative sampling, we note that for $f \in I_{\leq b}$, the multivariate polynomial division algorithm for $f \bmod G$ gives us a (typically non-unique) representation $f = \sum_i t_i g_i$. Since $<$ is compatible with deg, the intermediate results in that computation have degree $\leq b$, which ensures that $\deg t_i \leq b - d_i$. This already proves the claim.

To encrypt a message bit $M \in \{0, 1\}$, we proceed as follows:

Algorithm 3. Enc_G

1: **function** Enc_G(M):
2: $f \leftarrow \text{SampleI}(G)$
3: $e \leftarrow_{\$} \mathcal{X}$
4: $c := f + 2e + M$
5: **return** c

Accordingly, decryption of a ciphertext $c \in P_{\leq b}$ is performed by the following algorithm, where for $f \in P$, $f_{=0}$ denotes the constant coefficient of f:

Algorithm 4. Dec_G

1: **function** Dec_G(c):
2: $M := (c \bmod G)_{=0} \in \{-\lfloor \frac{q}{2} \rfloor, \ldots, \lfloor \frac{q}{2} \rfloor\}$
3: **return** $M \bmod 2$

Decryption is correct, provided that for the noise $2e \leftarrow_\$ 2\mathcal{X}$ we have $|2e_{=0}| < \lfloor \frac{q}{2} \rfloor$. If $\mathcal{X}$ is a sufficiently narrow discrete Gaussian distribution, this will be the case with overwhelming probability.

Remark 1. Embedding the message in the noise
In algorithm Enc_G above, the message M is merely one bit and is embedded only in the degree 0 term of the noise. Hence, in algorithm Dec_G, we also take only the degree 0-coefficient $(c \bmod G)_{=0}$. In particular, this means that fake ciphertexts c not generated by Enc_G still decrypt to a bit, even if $c \bmod G$ is not in the support of $2\mathcal{X} + \{0, 1\}$. Alternatively, we could output an error in the latter case.[2]

In fact, in [1] it is implicitly assumed (and also implemented that way in the reference implementation) that the noise is completely contained in degree 0. Unfortunately, these issues are not addressed in [1] and we will show in Sect. 5 that for $d > 1$ this choice renders the system insecure for $b > d$. For $b = d$ or $d = 1$, compare the following Sect. 3, where we show that these choice offer no benefit compared to $b = d = 1$. For $b = d = 1$, the PCN cryptosystem is a reformulation of Regev's scheme.

Actually, if the message is contained only in degree 0, the coefficients belonging to the monomials of $\mathcal{Q}_{\leq d}$ other than the constant term of a ciphertext polynomial c are completely irrelevant for decryption (cf. Prop. 1, which is a special case of that).

So unless one wants to detect fake ciphertexts as mentioned above or make use of the multiplicative homomorphic properties (cf. Sect. 2.3), one should really use uniform noise for those coefficients (or just leave those coefficients out of the ciphertext altogether).

In this work we will consider the more general setting, where the message is contained in degree 0, but the noise distribution $\mathcal{X}$ on $\mathcal{Q}_{\leq b}$ is arbitrary. When we assume that the noise is concentrated in degree 0, we will explicitly state that.

2.3 Homomorphic Properties and Public Key Version

One of the appealing aspects of the PCN cryptosystem is that it is somewhat homomorphic:

$P \to \mathcal{Q}, f \mapsto f \bmod G$ is actually a ring map. This means that for ciphertexts $c_1 = f_1 + 2e_1 + M_1, c_2 = f_2 + 2e_2 + M_2$, with $f_i \in I_{\leq b}, e_i \in \mathcal{Q}_{\leq b}, M_i \in \{0, 1\}$, we have

$$c_1 + c_2 = (f_1 + f_2) + 2(e_1 + e_2) + M_1 + M_2$$

and

$$c_1 \cdot c_2 = g + 2\,(2e_1e_2 + e_1M_2 + e_2M_1 \bmod G) + (M_1M_2 \bmod G),$$

where $g \in I_{\leq 2b}$

[2] Note that if the support of $2\mathcal{X} + \{0, 1\}$ is a vector space, a CPA-attacker can check for this error himself, so this does not affect security.

From this, we get $\mathrm{Dec}_G(c_1) \dotplus \mathrm{Dec}_G(c_2) = \mathrm{Dec}(c_1 \dotplus c_2)$, provided that the noise of the sum/product does not grow too large.

For sums, this implies that for a sufficiently narrow Gaussian $\mathcal{X}$, the cryptosystem supports a limited number of homomorphic additions at the cost of increased noise, and still decrypts correctly with overwhelming probability.

Note that this also holds in the case that we embed several bits into one ciphertext, provided the noise is narrow coefficient-wise. Via the usual generic construction [14], these additive somewhat homomorphic properties allow to convert the secret key cryptosystem into a public key cryptosystem by publishing a sufficient amount of encryptions of 0 as the public key. Note that the same applies to Regev's scheme [12] described below in Section 3. For simplicity, in this work we deal with the secret key versions of both schemes, but it is easy to see that everything carries over directly to the public-key setting.

For multiplications, if the noise is concentrated in degree 0, we get that the noise is approximately multiplied for each multiplication of ciphertexts[3], so we can also perform a limited number of homomorphic multiplications.

If $d > 1$ and $\mathcal{X}$ is not supported in degree 0, this will actually fail if done naïvely. The reason is that even if all coefficients of e_1, e_2 are small, $e_1 \cdot e_2 \bmod G$ might have large coefficients due to reduction mod G.

This is the case even if the coefficients of the Gröbner basis polynomials are small; take for example the reduced Gröbner basis $G = (g_1, \ldots, g_{2n}) \subset \mathbb{F}_q[X_1, \ldots X_{2n}]$ with

$g_1 = X_1^2 - a_1, \quad g_2 = X_2^2 - a_2,$
$g_{2i} = X_{2i}^2 - a_{2i} X_{2i-2} X_{2i-3},$
$g_{2i+1} = X_{2i+1}^2 - a_{2i+1} X_{2i-2} X_{2i-1}$ for $i \geq 1$ and $a_i \in \mathbb{F}_q$ small.

Then for $e_1 = e_2 = X_{2n} X_{2n-1} \in \mathcal{Q}_{\leq 2}$, we have $e_1 \cdot e_2 \bmod G = \prod_{i=1}^{2n} a_i$, which is exponentially large.

This observation makes it highly desirable to concentrate the noise and message in degree 0. Unfortunately, this renders the system insecure (cf. Sect. 5) unless $d = 1$ or $b = d$. By the results of Section 3, in the latter cases, we should rather use $b = d = 1$.

2.4 Security Assumptions

[1] introduced the following three security problems related to the PCN cryptosystem:

Definition 2. *The* Gröbner basis with noise (GBN) problem $\mathrm{GBN}_{n,\mathbf{Gen},d,b,\mathcal{X}}$ *for parameters as above is defined as follows:*

Let $G \leftarrow \mathbf{Gen}(1^\lambda)$ *be a reduced Gröbner basis. Given access to a sampling oracle for* $\mathrm{Sample}\mathrm{I}_{\mathcal{X}}$*, the task is to find* G*. The advantage for a (ppt) algorithm* A *in solving the* $\mathrm{GBN}_{n,\mathbf{Gen},d,b,\mathcal{X}}$ *problem ist given as*

$$\mathbf{Adv}^{gbn}_{n,\mathbf{Gen},d,b,\mathcal{X},A}(\lambda) = \mathbf{Pr}[A \text{ solves the } \mathrm{GBN}_{n,\mathbf{Gen},d,b,\mathcal{X}}\text{-}problem] - \frac{1}{|\mathcal{G}|},$$

[3] Note that this also increases the total degree, which can be adressed by reencryption techniques[4], but this will not be important for us here.

where $\mathcal{G}$ is the set of possible secret keys and the probability is over the coins of **Gen**, SampleI *and A.*

Note that we always assume that $|\mathcal{G}|$ is exponential.

Definition 3. *The* Ideal remainder with noise problem $\mathrm{IRN}_{n,\mathbf{Gen},d,b,\mathcal{X}}$ *for parameters as above is defined as follows:*

Let $G \leftarrow \mathbf{Gen}(1^\lambda)$ and a uniformly random challenge $x \leftarrow_\$ P_{\leq b}$. Given x and access to a sampling oracle for $\mathrm{SampleI}_{\mathcal{X}}$, the task is to find $x \bmod G \in \mathcal{Q}_{\leq b}$. The advantage for a ppt algorithm B for this problem is given as

$$\mathbf{Adv}^{irn}_{n,\mathbf{Gen},d,b,\mathcal{X},B}(\lambda) = \mathbf{Pr}[B \text{ solves the } \mathrm{IRN}_{n,\mathbf{Gen},d,b,\mathcal{X}}\text{-problem}] - \frac{1}{|\mathcal{Q}_{\leq b}|},$$

where the probability is over the coins of **Gen**, SampleI, B *and the uniform choice of the challenge x.*

Note that this definition of advantage implicitly assumes that $\mathcal{Q}_{\leq b}$ is known to the attacker.

Definition 4. *The* Ideal membership with noise (IMN) problem $\mathrm{IMN}_{n,\mathbf{Gen},d,b,\mathcal{X}}$ *for parameters as above is defined as follows:*

Let $G \leftarrow \mathbf{Gen}(1^\lambda)$. Given access to a sampling oracle for $\mathrm{SampleI}_{\mathcal{X}}$, the task is to distinguish a challenge polynomial x drawn either as $x \leftarrow_\$ \mathrm{SampleI}_{\mathcal{X}}$ or as a uniform $x \in_R P_{\leq b}$. The advantage for a ppt algorithm C for this is given as

$$\mathbf{Adv}^{imn}_{n,\mathbf{Gen},d,b,\mathcal{X},C}(\lambda) = \mathbf{Pr}[C^{\mathrm{SampleI}_{\mathcal{X}}()}(x) = 1] - \mathbf{Pr}[C^{\mathrm{SampleI}_{\mathcal{X}}()}(u) = 1]$$

where $x \leftarrow_\$ \mathrm{SampleI}_{\mathcal{X}}, u \in_R \mathcal{Q}_{\leq b}$ and the probability is over the coins of **Gen**, SampleI, C *and choices of x or u. Note that we differ by a factor 2 from [1].*

The security assumption made in [1] is that for appropriate choice of parameters, namely $b \leq d \leq 1$ arbitrary, $\mathbf{Gen} = \mathrm{GBGen}_{\mathrm{dense}}$ and $\mathcal{X}$ a sufficiently broad discrete Gaussian distribution on $\mathbb{F}_q$, the advantage for any ppt algorithm is negligible for GBN / IRN / IMN.

Also, it was claimed in [1] that all of these assumptions and the IND-CPA-security of PCN are essentially equivalent:

1. The GBN problem is hard iff the IRN problem is hard.
2. For polynomially-sized $\mathcal{Q}_{\leq b}$, IRN is hard iff IMN is hard.
3. If IMN is hard, the PCN cryptosystem is IND-CPA-secure.

As their proofs of 1 and 2 contain errors (amongst other things, the reduction presents the wrong distributions to the algorithms), we will redo the proofs for 1 and 2 in Sect. 4.

Unfortunately, we will have to make additional assumptions compared to [1], most importantly we have to assume $d = 1$ for the $\Rightarrow$ direction in the first proof and for the $\Leftarrow$ direction of the second. We will also give a counterexample indicating that these additional assumptions are necessary.

3 Relations to LWE and Regev's Scheme

We will now relate the PCN cryptosystem to LWE and show that the cases $b = d$ and $d = 1$ both reduce to Regev's LWE-based scheme. Let us briefly recall the LWE distribution, the LWE assumption and Regev's scheme from [12], which has a reduction to the LWE assumption:

Definition 5. *Learning with Errors (LWE)*
Let Φ be some noise distribution on a finite field $\mathbb{F}_q$ and $n \in \mathbb{N}$ and $s \in \mathbb{F}_q^n$. The LWE distribution $\mathcal{L}_{s,\Phi}$ on $\mathbb{F}_q^n \times \mathbb{F}_q$ is obtained by sampling $a_1, \ldots, a_n \in \mathbb{F}_q$ uniformly random, $e \leftarrow_\$ \Phi$ and outputting $(a_1, \ldots, a_n, \sum a_i s_i + e)$.

The computational LWE problem $\mathrm{LWE}_{n,q,\Phi}$ is the following problem: For uniformly random $s \in \mathbb{F}_q^n$, compute s when given oracle access to $\mathcal{L}_{s,\Phi}$.

The decisional LWE problem $\mathrm{DLWE}_{n,q,\Phi}$ is the following problem: For uniformly random $s \in \mathbb{F}_q^n$, distinguish $x \in_R \mathbb{F}_q^{n+1}$ from $x \leftarrow_\$ \mathcal{L}_{s,\Phi}$ when given oracle access to $\mathcal{L}_{s,\Phi}$.

The LWE assumption (for $q, \mathcal{X}$ given functions of n) states that any ppt algorithm can only solve these problems with negligible advantage.

Definition 6. *Regev's scheme*
Let Φ be some noise distribution on a finite field $\mathbb{F}_q$. In its secret key version[4], Regev's scheme generates a secret key $s = (s_1, \ldots, s_n) \in \mathbb{F}_q^n$ uniformly. We encrypt a message $M \in \{0, 1\}$ by sampling $a = (a_1, \ldots, a_n) \in \mathbb{F}_q^n$ randomly, $e \leftarrow_\$ \Phi$ and defining the ciphertext as $(a, \langle a, s\rangle + 2e + M)$, where $\langle a, s\rangle = \sum a_i s_i$ is the scalar product.

Decryption recovers $2e + M$ and, from that, M itself, provided e is small enough.

As already noted in [1], Regev's scheme is equivalent to the PCN cryptosystem for $d = b = 1$. To see that, we can identify Regev's secret s with the Gröbner basis $G = (X_1 + s_1, \ldots, X_n + s_n)$. We identify ciphertexts (a, b) with linear polynomials $\sum a_i X_i + b$ and Φ with $\mathcal{X}$.

In fact, such a relationship also holds for $b = d > 1$ and for $b > d = 1$, where the cases with $d = 1$ were already discussed in [1]:

Theorem 1. *Relationship of PCN with LWE for $b = d$ or $d = 1$*

- *For $b = d$, the IND-CPA-security of PCN (with parameters $q, \mathcal{X}, b, d, n$) is equivalent to the IND-CPA-security of Regev's scheme (with parameters $q, \mathcal{X}, n$).*
- *For $d = 1$, there exists a tight security reduction from the IND-CPA-security of PCN (with parameters $q, \mathcal{X}, b, d, n$) to the IND-CPA-security of Regev's scheme (with parameters $q, \mathcal{X}, \binom{n+b}{b}$).*
- *For $b = d = 1$, the PCN cryptosystem is a reformulation of Regev's scheme.*

[4] The public-key version is obtained by using Rothblum's construction [14] just as with PCN and all observations carry over directly to the public-key versions of both schemes.

Proof. For $b = d$, this follows from proposition 1 below, showing that in this case the PCN cryptosystem is a redundant version of Regev's scheme. For $d = 1$, this follows from proposition 2 below, showing that in this case the PCN cryptosystem is a structured version of Regev's scheme. The case $b = d = 1$ was already discussed above.

Regarding ciphertext length, recall that the PCN-ciphertexts are $\binom{n+b}{b}$ elements from $\mathbb{F}_q$. As a consequence, for $b = d > 1$ we have a loss in efficiency, but no gain in security. For $b > 1, d = 1$, we have no gain in efficiency (apart from a shorter secret key compared to Regev's) and potentially a loss in security. Therefore, there is little point in using the PCN cryptosystem for $b = d$ or $d = 1$ unless $b = d = 1$.

Proposition 1. *Relation of PCN with LWE for $b = d$*
Consider the case $b = d$ and assume that $\mathcal{X}$ outputs $e \leftarrow_\$ \mathcal{X}, e = \sum_m e_m \cdot m$, where the sum runs over the monomials and the e_m are chosen independently, their distribution possibly depending on m (This is the case if the noise is contained in degree 0). Then the PCN cryptosystem is essentially[5] a reformulation of (the secret key version) of the amortized[6] variant [11] of Regev's scheme, where each monomial m of $\mathcal{Q}_{\leq b}$ corresponds to one parallel instance of Regev's original scheme.

To see this, consider a PCN-ciphertext c. By lemma 1, c is of the form $c = \sum t_i \cdot g_i + 2e + M$ for $e \leftarrow_\$ \mathcal{X}$ with $t_i \in \mathbb{F}_q$. Let us write $c = \sum_m c_m \cdot m$ for the monomials m of c. Then for $1 \neq m \in \mathcal{Q}_{\leq b}$, the coefficients of the ciphertext are $c_m = \sum t_i \cdot g_{i,m} + 2e_m$ and $c_1 = \sum t_i \cdot g_{i,1} + 2e_1 + M$. These are noisy random linear combinations of the secret $g_{i,m}$ as in Regev's scheme. The other $m \notin \mathcal{Q}_{\leq b}$ are $m = X_i^d$ and there we have $c_{X_i^d} = t_i$. It follows that the ciphertexts are exactly as in the amortized variant of Regev's.

When taking that point of view for general $b = d > 1$, beware that by construction, for some $m \in \mathcal{Q}_{=b}$ and some $j \in \{1, \ldots, n\}$ we can have $m \not< X_j^d$, so $g_{m,X_j^d} = 0$. In that case, the corresponding LWE-instance has a secret key from $\mathbb{F}_q^{n'}$ for some $n' < n$. In particular, for $\mathrm{GBGen}_{\mathrm{dense}}$ and $m = X_n^{b-1} X_{n-1}$ we have $n' = 1$. Of course, since the message is contained in degree 0, only the Regev-instance for the constant monomial $m = 1$ is relevant and the above is not an issue. The other coefficients (apart from the X_i^d) are superfluous, not only for the ciphertexts but also for the secret key, since these coefficients are independent of the $g_{i,1}$ and the message. It follows that for $b = d$, the security of the PCN cryptosystem does not depend on d at all, but the efficiency degrades with d. Note that if the e_m are not independent, this might only help the attacker.

[5] The only difference is that for some of the parallel instances, the secret key has fewer coordinates.

[6] This amortized variant just runs parallel instances of Regev's, where the random coefficients a of the noisy linear combinations $\langle a, s \rangle + e$ are shared between instances.

Proposition 2. *Reduction from PCN to LWE for $d = 1$, b arbitrary*
Consider the case $d = 1$, b arbitrary. Then the PCN cryptosystem can be viewed as a structured version of Regev's scheme. There is a reduction from the (IND-CPA-)security of PCN to the (IND-CPA-)security of Regev's original scheme, as already noted in [1].

To see this, first observe that for $d = 1$, the secret Gröbner basis of the PCN cryptosystem is necessarily of the form $G = (X_1 - s_1, \ldots, X_n - s_n)$ for $s = (s_1, \ldots, s_n)$. We then have $f \bmod G = f(s)$, so SampleI(G, b) just gives us polynomials $f - f(s) \cdot 1$ for $f \in P_{\leq b}$ uniformly. For a monomial $m \neq 1$ of $P_{\leq b}$, let $\widetilde{s_m} := m(s) \in \mathbb{F}_q$. It follows that PCN-ciphertexts are of the form $\sum_{m \neq 1} a_m \cdot m - (\sum_{m \neq 1} a_m \widetilde{m_s}) \cdot 1$, where the $a_m \in \mathbb{F}_q$ are uniform and the sums run over the monomials m of $P_{\leq b}$ (except the constant one). This implies that PCN-instances are nothing but Regev-instances with a structured secret key $\widetilde{m_s}$.

Our reduction just has to remove that structure from the key. This can be done as in [13] by rerandomizing the secret:

Our reduction chooses $t_m \in \mathbb{F}_q$ uniformly for $m \neq 1$ monomial of $\mathbb{F}_q$. Then we bijectively transform any PCN-ciphertext $c = \sum_{m \neq 1} a_m m + b$ into a Regev-ciphertext $T_t(c) = (\boldsymbol{a}, b + \sum_{m \neq 1} t_m a_m)$. These ciphertexts are distributed as Regev-ciphertexts with uniform secret $\widetilde{\boldsymbol{s}} + \boldsymbol{t}$ with the same a_i and the same noise $e \leftarrow_\$ \mathcal{X}$.

4 Security Proofs

In this section, we clarify the relationships between the different security assumptions we recalled in Sect. 2.4 and the security of the PCN cryptosystem. We will first give counterexamples, showing that, under the LWE assumption, the GBN, IRN and IMN problems are not equivalent for general $d > 1$, refuting the claims from [1]. We will then give corrected proofs for $d = 1$.

In order to make the proofs for $d = 1$ work, we need to impose the following technical restriction on $\mathcal{X}$:

Definition 7. *We call a noise distribution $\mathcal{X}$ on $\mathcal{Q}_{\leq b}$ recognizable with noise, if for every $p' = \mathrm{poly}(\lambda)$ there exists a ppt algorithm D that, given oracle access to $\mathcal{X}_{a,p}$ with $p \leq p'$, outputs a with overwhelming probability for uniform $a \leftarrow_\$ \mathcal{Q}_{\leq b}$. Hereby, $\mathcal{X}_{a,p}$ is defined as a distribution that, with probability $\left(1 - \frac{1}{p(\lambda)}\right)$, outputs a uniform $x \leftarrow_\$ \mathcal{Q}_{\leq b}$, and otherwise (with probability $\frac{1}{p(\lambda)}$) outputs $x = e + a$ for $e \leftarrow_\$ \mathcal{X}$.*

We remark that a discrete Gaussian distribution with polynomial standard deviation is recognizable with noise (using as D the majority vote).

Theorem 2. IRN hard $\Leftrightarrow$ GBN hard, IRN hard $\Leftrightarrow$ IMN hard
Assume that the LWE assumption holds for some $q = \mathrm{poly}(n)$ and some noise distribution Φ on $\mathbb{F}_q$ that is recognizable with noise. Then there exists an instantiation for **Gen** *with $\mathcal{X}$ recognizable with noise (and, in particular, distinguishable*

from uniform), such that the IRN problem is easy, but both the GBN problem and the IMN problem are hard, contradicting the proofs from [1].

Proof. Consider the case $b = d = 2$ and let $q = \mathrm{poly}(n)$ and Φ be such that the LWE assumption holds for q and Φ. We consider **Gen** that outputs reduced Gröbner bases of the form $G = (X_1^2, X_2 + s_2 X_1, X_3 + s_3 X_1, \ldots, X_n + s_n X_1)$. Then $\mathcal{Q} = \mathcal{Q}_{\leq b}$ is generated by X_1 and 1 as a vector space. For the noise distribution $e_1 X_1 + e_2 \leftarrow_\$ \mathcal{X}$, we take $e_1 \leftarrow_\$ \Phi$ and e_2 uniform from $\mathbb{F}_q$.

By construction, the constant coefficient of all Gröbner base polynomials is 0, so for any $f \in I_{\leq b}$ we have $f \bmod G = f_{=0} + r(f) X_1$ for some $r(f) \in \mathbb{F}_q$. This already implies that we can guess the remainder by guessing $r(f)$ with noticeable probability $\frac{1}{q}$, compared to $|\mathcal{Q}_{\leq b}| = q^2$, giving a non-negligible advantage for the IRN problem.

Now let $f \leftarrow_\$ \mathrm{SampleI}_{\mathcal{X}}$ with $f = f^{(2)} + f^{(1)} + f^{(0)}$ be the homogenous parts of degree 2, 1 and 0. Since $P_{\geq 2} \subset I$ and the noise in degree 0 is uniform, we get that $f^{(2)}$ and $f^{(0)}$ are independently uniform and independent of $f^{(1)}$. Let us write $f = e + \sum_i t_i g_i$ with $g_i \in G, t_i \in P, e \in \mathcal{Q}$. Since $\deg g_i \geq 1$, $f^{(1)}$ only depends on e and the degree-0 part of the t_i.

It follows that $f^{(1)} = bX_1 + a_2 X_2 + \ldots a_n X_n$ with a_i uniform and $b = \sum a_i s_i + e_1$ with $e_1 \leftarrow_\$ \Phi$, i.e. $f^{(1)}$ is distributed as $\mathcal{L}_{s,\Phi}$. It follows that the IMN problem is equivalent to the $\mathrm{DLWE}_{n-1,q,\Phi}$-problem and the GBN problem is equivalent to the $\mathrm{LWE}_{n-1,q,\Phi}$-problem, both of which we assumed to be hard.

Remark 2. Separation of IRN and GBN.
There is also a separation between IRN and GBN, if we assume that the LWE-assumption holds for some q and some Gaussian noise. Namely, take $b = d = 2$ and let **Gen** output Gröbner bases of the form $X_1^2 - s_1, \ldots X_n^2 - s_n$ with $s_i \in \mathbb{F}_q$ independent and uniformly. Note that there are no linear terms here. As noise distribution choose Gaussian noise, concentrated in degree 0. Then the GBN problem is is hard if the LWE assumption holds (cf. Prop. 1). However, IMN is easy, because noisy samples from the ideal have no linear terms.

Note that we assumed Φ to be recognizable with noise to satisfy the requirements from [1] and all requirements from Thm 3, apart from $d = 1$, below. Without that restriction on the noise, we may take Φ to be uniform and get an information-theoretical variant of Thm. 2 without the need for an LWE assumption.

For $d = 1$, the statements from [1] actually hold. The reason why we can make the proof work only in that case is that we need an amplification step, for which our rerandomization strategy only works for $d = 1$.

Theorem 3. IRN hard $\Leftrightarrow$ GBN hard $\Leftrightarrow$ IMN hard for $d = 1, q = \mathrm{poly}(n)$ and $\mathcal{X}$ recognizable with noise
For any **Gen**$, \mathcal{X}, b \leq d$*, we have:*

1. *If the IMN problem is hard, the PCN cryptosystem is IND-CPA-secure.*
2. *If $d = 1, \mathcal{X}$ recognizable with noise, then the IRN problem is hard iff the GBN problem is hard.*

3. *If $q = \mathrm{poly}(n), d = 1, \mathcal{X}$ distinguishable from uniform, then then IRN problem is hard iff the IMN problem is hard.*

Proof. Statement 1 is proven in [1]. Statement 2 and 3 are proven in the appendix of the full version.

5 Attack on Low-Dimensional Noise

In this section, we present our main contribution. We will present a polynomial time CPA-attack against the PCN cryptosystem that recovers the secret key, if $b > d > 1$, using that the noise is contained in degree 0. Note that all concrete parameter choices of [1] use $d = 1$, but this attack still violates the explicit security assumption, which is stated for general d.

Throughout this section we assume that $d > 1$ and that $\mathcal{X}$ is supported in degree $\leq k$. Furthermore, we assume for simplicity that we are using $\mathbf{Gen} = \mathrm{GBGen}_{\mathrm{dense}}$. Using the notation from Alg. 1 above, let us write the secret key as $G = (g_1, \ldots, g_n)$ with $g_i = X_i^d + \sum_m g_{i,m} \cdot m$. Our attack will derive linear equations for the $g_{i,m}$.

The intuition behind the attack is the following:

Since the support of $\mathcal{X}$ is contained in $\mathcal{Q}_{\leq k}$, all ciphertexts are contained in a vector sub-space $N := I_b \oplus \mathcal{Q}_{\leq k} \subsetneqq P_{\leq b}$. We can recover this vector space N via a CPA-oracle (In the public-key variant, N is directly given by the public key). Note that the dimension of N is known, namely, it is $\dim N = \dim P_{\leq b} - \dim \mathcal{Q}_{\leq b} + \dim \mathcal{Q}_{\leq k} = O(n^b)$.

Of course, since $g_i \in I_{\leq b} \subset N$, the secret Gröbner base polynomials must also lie in this subspace. If the inclusion is proper, this directly translates into linear equations for the $g_{i,m}$. Unfortunately (for the attacker), these equations do not yet determine the g_i: We may add any error term $h \in \mathcal{Q}_{\leq k}$ to g_i and we still have $g_i + h \in N$.

To overcome this, we make use of the fact that I is an ideal, so $t \cdot g_i \in I_{\leq b} \subset N$ for any polynomial t with $\deg t \leq b - d$. Roughly speaking, this effectively also multiplies the error term by t and we will use this to move the error out of $\mathcal{Q}_{\leq k}$. In order to move the error completely out of $\mathcal{Q}_{\leq k}$, we need to multiply by polynomials t of degree $> k$, so we expect our attack to work whenever $b - d > k$. In particular, for $k = 0$, this strategy will recover the secret key for $b > d$.

Note that we will also cover the case $d > k$:

Remember that for $d > 1$, $\mathcal{Q}$ has exponential (in n) dimension and contains polynomials of degree $> d$ for n large. In the case $d > k$, we will not get any useful information from $g_i \in N$. But for $k < b$ we still have $\mathcal{Q}_{\leq k} \subsetneqq \mathcal{Q}_{\leq b} \subsetneqq \mathcal{Q}$ for n large. This means we will get some useful equations from $t \cdot g_i \in N$ for polynomials t with $b - d \geq \deg t \geq d - k$

We now present the actual algorithm and then we will give a rigorous analysis.

Algorithm 5. ppt attack against PCN cryptosystem with low-degree noise

Input: $1^n, k, b, d$, access to a CPA oracle
Output: Secret key $g_{i,m}$

1: $N :=$ a vector space basis of $\mathcal{Q}_{\leq k}$
2: **repeat**
3: Create $f \in I_{\leq b} \oplus \mathcal{Q}_{\leq k}$ an encryption of 0
4: $N := N \cup \{f\}$
5: **until** $\dim(\operatorname{Span} N) = \dim(I_{\leq b} \oplus \mathcal{Q}_{\leq k})$
▷ We now have $\operatorname{Span}(N) = I_{\leq b} \oplus \mathcal{Q}_{\leq k}$
6: Write N as a matrix and perform Gaussian elimination to obtain $\operatorname{Span}(N) = \ker A$ for linear $A : P_{\leq b} \to \mathbb{F}_q^{\dim \mathcal{Q}_{\leq b} - \dim \mathcal{Q}_{\leq k}}$.
7: **for** $i = 1$ to n **do**
8: Let $E_i := \emptyset$ be the set of equations for the $g_{i,m}$.
9: **for all** monomials $t \in P_{\leq b-d}$ **do**
10: Add the inhomogenous linear equations $A(t \cdot g_i) = 0$ in the variables $g_{i,m}$ to E_i.
11: Solve the system of equations E_i
12: **return** A solution $\overline{g_{i,m}}$ for each of the E_i

Theorem 4. Algorithm 5 is Correct and Runs in Polynomial Time with Overwhelming Probability
With overwhelming probability, algorithm 5 runs in polynomial time $O(n^{2b+d+1})$.

If $n > k, d > 1$ and $b - d > k$, the algorithm outputs the secret key.

In particular, for $k = 0$, that is, for noise concentrated in degree 0, the algorithm gives an efficient key-recovery attack for $b > d > 1$.

More precisely, we claim that, if $n > k$, we have $\overline{g_{i,m}} = g_{i,m}$ whenever $\deg m > k - (b - d)$.

For any other m with $\deg m \leq k - (b - d)$, $\overline{g_{i,m}}$ may be chosen arbitrarily by the algorithm (the solution of the E_i is not unique if such monomials exist).

Proof. Let us start with the running time:
In line 2 to 5, we use the CPA-oracle to obtain $f \in I_{\leq b} \oplus \mathcal{Q}_{\leq k}$ (Note here that if the message is embedded only within $\mathcal{Q}_{\leq k}$ as well, *any* ciphertext will do).

Since the $I_{\leq b}$-component of f is uniform, after $O(\dim I_{\leq b}) = O(n^b)$ steps, we will eventually obtain all of $I_{\leq b} \oplus \mathcal{Q}_{\leq k}$ with overwhelming probability.

After that, the running time of the algorithm is dominated by solving the E_i. Each E_i consists of $\dim P_{\leq b-d} \cdot (\dim \mathcal{Q}_{\leq b} - \dim \mathcal{Q}_{\leq k}) = O(n^{2b-d})$ equations in at most $\dim P_{\leq d} = O(n^d)$ unknowns. Since $2b - d > d$, this gives a running time of $O(n^{2b-d}) \cdot O(n^d) \cdot O(n^d) = O(n^{2b+d})$ for solving each E_i, hence a total running time of $O(n^{2b+d+1})$ to solve all the E_i (cf. Rmk 3 below).

We now turn to the correctness statement:

By construction, $\operatorname{Span}(N) = I_{\leq b} \oplus \mathcal{Q}_{\leq k}$, starting from line 6.

Since $g_i \in I_{\leq b} \subset I_{\leq b} \oplus \mathcal{Q}_{\leq k}$, we have $A(g_i) = 0$. Making use of the fact that I is an ideal, we also have $t \cdot g_i \in I_{\leq b}$ and hence $A(t \cdot g_i) = 0$ for $\deg t \leq b - d$.

It follows that the equations we derive for the $g_{i,m}$ are correct, that is, the $g_{i,m}$ satisfy the equations E_i.

Note that in line 10, we rewrite the linear equation $A(tg_i)$ as an equation in the $g_{i,m}$. Implicitly, we add the equations $g_{i,m} = 0$ for $m > X_i^d$, $g_{i,X_j^d} = 0$ for $i \neq j$ and $g_{i,X_i^d} = 1$ at this point. Since we set the coefficient of X_i^d to be 1 in that last equation, the resulting system of equations E_i is a system of inhomogeneous linear equations.

Now, the E_i might have more than one solution, apart from the secret key $g_{i,m}$.

To show that the coefficients for monomials m with $\deg m \leq k - (b - d)$ are undetermined, we first observe that $P_{\leq k} \bmod G = \mathcal{Q}_{\leq k}$, so $I_{\leq b} \oplus \mathcal{Q}_{\leq k} = I_{\leq b} + P_{\leq k}$. Consequently, by Lemma 1 $I_{\leq b} \oplus \mathcal{Q}_{\leq k}$ are exactly all elements of the form $f = e + \sum t_i g_i$ with arbitrary $e \in P_{\leq k}$, $t_i \in P_{\leq (b-d)}$. The coefficients of the g_i of degree $\leq k - (b - d)$ then only affect the coefficients of f of degree $\leq k$, which are uniform due to e. So $I_{\leq b} + P_{\leq k}$ does not depend on the coefficients of g_i of degree $\leq k - (b - d)$, which implies that these coefficients span a subspace of the kernel of the E_i.

To show that for $n \geq k + 1, \deg m > k - (b - d)$, we have $\overline{g_{i,m}} = g_{i,m}$, let $\overline{g_i}, \overline{g_i}'$ be 2 solutions for E_i and $h = \overline{g_i} - \overline{g_i}'$. We need to show $\deg h \leq k - (b - d)$ (which means $h = 0$ if the right-hand side is negative).

By construction of the E_i, we know that $A(t \cdot h) = 0$, or equivalently $t \cdot h \in I_{\leq b} \oplus \mathcal{Q}_{\leq k}$, for all $t \in P_{\leq b-d}$. The other equations coming from the restrictions on the set of monomials that can appear in the $\overline{g_i}, \overline{g_i}'$ imply that h can only contain coefficients for the set of monomials $\{m \mid m < X_i^d, m \neq X_j^d \text{ for any } j\}$. This implies that $h \in \mathcal{Q}_{\leq d}$, in particular, $\deg h \leq d$.

We will show that $\deg h \leq k - \alpha$ for $0 \leq \alpha \leq b - d$, using induction on α. For $\alpha = b - d$, the claim then follows.

For the base case $\alpha = 0$, we already observed that $h \in \mathcal{Q}_{\leq d}$. Setting $t = 1$ in $A(th) = 0$ yields $h \in I_{\leq b} \oplus \mathcal{Q}_{\leq k}$. Together, these give $h \in \mathcal{Q}_{\leq k} \cap \mathcal{Q}_{\leq d}$, so $\deg h \leq k$ as desired.

For the inductive step, assume $\deg h \leq k - \alpha$ for $\alpha < (b - d)$. Assume w.l.o.g. that $h \neq 0$, since otherwise we are done. Let $H = \mathrm{LT}(h)$ be the leading term. Since the monomial order is degree-compatible, $\deg H = \deg h$. We need to show that $\deg H < k - \alpha$.

For this, choose a monomial t of degree $\deg t = \alpha + 1 \leq b - d$ such that $t \cdot H \in \mathcal{Q} = \mathrm{Span}\{X_1^{v_1} \cdots X_n^{v_n} \mid v_i < d \text{ for all } i\}$. This can be accomplished for $d > 1$ and $n \geq k + 1$ by choosing $t = X_{i_1} \cdots X_{i_{\alpha+1}}$ a product of $\alpha + 1$ pairwise different variables, disjoint from those of H.[7] By the properties of a monomial order, $\mathrm{LT}(t \cdot h) = t \cdot H$. Since $t \cdot H \in \mathcal{Q}$, this is not reduced modulo G, so we have $\mathrm{LT}((t \cdot h) \bmod G) = t \cdot H$. Since $A(t \cdot h) = 0$, we have $t \cdot h \in I_{\leq b} \oplus \mathcal{Q}_{\leq k}$. This implies $(t \cdot h) \bmod G \in \mathcal{Q}_{\leq k}$, in particular $(t \cdot h) \bmod G$ has degree at most k. It follows that H has degree at most $k - \deg t = k - \alpha - 1$. This finally proves the theorem.

Remark 3. Algorithm 5 was optimized for simplicity of analysis. We can get a better running time by using the highly structured nature of the equations on

[7] Note that if $k < d$, any t of degree $\alpha + 1$ will do without the restriction on n.

the E_i. In particular, we don't need all $t \in P_{\leq b-d}$, as the proof above shows and we also don't need to solve the E_i separately for $1 \leq i \leq n$.

Also, we would like to remark that Algorithm 5 also gives an attack to the underlying GBN, IRN and IMN problems; in particular the existence of this attack is not related to the flaws in security proof of [1] we pointed out in section 4.

6 Conclusion and Open Problems

We have seen that for $d > 1$, the security reductions from [1] will no longer work and there arise problems in choosing a noise distribution. Concentrating the noise in low degree makes the scheme insecure unless $b = d$, so the obvious way to go is to spread the noise over the full quotient. We remark that it might be possible to retain the homomorphic properties by using a different strategy to generate the Gröbner basis, allowing multiplicative homomorphic properties in Ring-LWE [10] style. We leave this as an open problem.

Acknowledgements. We would like to thank Martin Albrecht for valuable discussions and helpful comments.

References

1. Albrecht, M.R., Farshim, P., Faugère, J.-C., Perret, L.: Polly Cracker, Revisited. In: Lee, D.H. (ed.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 179–196. Springer, Heidelberg (2011); Cryptology ePrint Archive, Report 2011/289, `http://eprint.iacr.org/`
2. 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)
3. Becker, T., Weispfenning, V.: Gröbner bases: a computational approach to commutative algebra. Graduate Texts in Mathematics. Springer (1993)
4. Brakerski, Z., Vaikuntanathan, V.: Efficient fully homomorphic encryption from (standard) LWE. To appear in FOCS 2011 (2011); Cryptology ePrint Archive, Report 2011/344 (2011), `http://eprint.iacr.org/`
5. Buchberger, B.: Ein Algorithmus zum Auffinden der Basiselemente des Restklassenrings nach einem nulldimensionalen Polynomideal. PhD thesis, Universität Innsbruck (1965)
6. Cox, D., Little, J., O'Shea, D.: Ideals, Varieties, and Algorithms, 3rd edn. Springer (2005)
7. 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 (2009)
8. Fellows, M., Koblitz, N.: Combinatorial cryptosystems galore! In: Finite Fields: Theory, Applications, and Algorithms. Contemporary Mathematics, vol. 168, pp. 51–61. AMS (1994)

9. Gentry, C.: A fully homomorphic encryption scheme. PhD thesis, Stanford University (2009)
10. Lyubashevsky, V., Peikert, C., Regev, O.: On Ideal Lattices and Learning with Errors over Rings. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 1–23. Springer, Heidelberg (2010)
11. Peikert, C., Vaikuntanathan, V., Waters, B.: A Framework for Efficient and Composable Oblivious Transfer. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 554–571. Springer, Heidelberg (2008)
12. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing (STOC 2005), pp. 84–93. ACM (2005)
13. Regev, O.: The learning with errors problem (invited survey). In: IEEE Conference on Computational Complexity, pp. 191–204. IEEE Computer Society Press (2010)
14. 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)