

Cryptographic Schemes Based on the ASASA Structure: Black-box, White-box, and Public-key (Extended Abstract)

Alex Biryukov, Charles Bouillaguet, and Dmitry Khovratovich

CSC& SnT, University of Luxembourg, Luxembourg, and University of Lille-1, France
{alex.biryukov,dmitry.khovratovich}@uni.lu,
charles.bouillaguet@univ-lille1.fr

Abstract. In this paper we pick up an old challenge to design public key or white-box constructions from symmetric cipher components. We design several encryption schemes based on the ASASA structure ranging from fast and generic symmetric ciphers to compact public key and white-box constructions based on generic affine transformations combined with specially designed low degree non-linear layers. While explaining our design process we show several instructive attacks on the weaker variants of our schemes¹.

Keywords: ASASA, multivariate cryptography, white-box cryptography, cryptanalysis, algebraic, symmetric.

1 Introduction

Since the development of public key cryptography in the late 1970's it has been an open challenge to diversify the set of problems on which such primitives were built as well as to find faster alternatives, since most public key schemes were several orders of magnitude slower than symmetric ones. One of the directions was to design public key schemes from symmetric components. As public key encryption requires trapdoors, they have been hidden in secret affine layers [39], field representations [43], biased S-boxes and round functions [45]; however most of these schemes were broken [35,42]. We recall that a typical symmetric cipher is built from layers of affine transformations (A) and S-boxes (S), a design principle dating back to Shannon. It is thus natural to see what designs can be made from such components. Whereas the classical cipher AES-128 consists of 10 rounds with 19 layers in total, it is striking that a lot of effort has been put into designing public-key schemes with only 3 layers, using the ASA (affine-substitution-affine) structure. This has indeed been the mainstream of what is known as *multivariate cryptography*. However, in this case, the non-linear layer is usually an ad-hoc monolithic function over the full state, as opposed to an array of independent S-boxes.

It has been known that the scheme SASAS with two affine and three nonlinear layers is vulnerable to a structural attack if the nonlinear layer consists of

¹ The full version of our paper is available at Eprint [8].

several independent S-boxes [10]. The scheme ASA, though secure for a random monolithic S-box, has been shown weak in concrete multivariate proposals. In the seemingly unrelated area of *white-box cryptography* the ASA approach to build obfuscated lookup tables failed multiple times. This suggests exploring the shortest scheme unbroken so far — the ASASA construction with injective S-boxes — in the application to symmetric (black-box), public-key, and white-box cryptography. Let us overview the related areas.

Retrospective of multivariate cryptography

The idea of multivariate cryptography dates back to the Shannon’s idea that recovering the secrets in any cryptographic scheme could be reduced to solving particular systems of (boolean) equations. Since nearly all forms of cryptology *implicitly* rely on the hardness of solving some kind of equation systems, then it must be possible to design cryptographic schemes that *explicitly* rely on the hardness of this problem. In multivariate public-key schemes, the public-key itself is a system of polynomial equations in several variables. It is well-known that solving such systems is NP-hard, even when the polynomials are quadratic (hence the name of the MQ problem, which stands for Multivariate Quadratic polynomial systems). An additional advantage of the MQ cryptosystems is that they seem invulnerable to quantum algorithms and hence are candidates for Post-Quantum cryptography.

Multivariate polynomials have been used in cryptography in the early 1980’s with the purpose of designing RSA variants with faster decryption. At this time, Imai and Matsumoto designed the first public-key scheme explicitly based on the hardness of MQ. It made it to the general crypto community a few years later under the name C^* [39].

Several years later, in 1995, Patarin [42] found a devastating attack against C^* , allowing to decrypt and to forge signatures very efficiently. Thereafter many multivariate scheme have been proposed (we counted at least 20 of them), including a plethora of bogus and vainly complicated proposal with a short lifespan. A few constructions stood out and received more attention than the others because of their simplicity and their elegance, such as HFE [43] and UOV [34].

However, the practical break of the first HFE challenge, supposed to offer 80 bits of security, in 2003 [29], and the demise of SFLASH in 2007 [24], just after the NESSIE consortium proposed it to be standardized, shattered the hopes and trust of the cryptographic community at large in multivariate cryptography. This brought the multivariate fashion to a stop.

The main problem in multivariate crypto is that the selection of candidates for the nonlinear layer S is scarce (we will discuss this in Section 4). What remains usually has so strong a structure within, that it can be detected and exploited even in the presence of unknown A layers. A very recent example is the promising matrix-based scheme ABC [48]. In the last years, a few researchers started designing public-key schemes based on the hardness of *random instances* of the MQ problem [46], though no drop-in replacement for conventional public-key encryption schemes has been proposed. Still, they are promising because

there is a consensus that random instances are hard, and all known algorithms are exponential and impractical on random systems.

This overview clearly indicates the need of a larger structure for multivariate cryptosystems, and suggests truly random polynomials in this context, which we use in our schemes.

Retrospective of white-box cryptography

In a parallel development a notion of *white-box cryptography* (WBC) has been introduced in [17]. The initial motivation was to embed symmetric secret keys into the implementation of popular standards like AES or DES in a way that binds the attacker to the specific implementation for DRM purposes. Several proposals have been made [15, 16] with the main idea to obfuscate key-dependent parts of the cipher and publish them as lookup tables, so that the entire encryption routine becomes just a sequence of table lookups. The obfuscation constitutes of wrapping the nonlinear transformation (**S**) with random affine transformations (**A**) so that the affine layers would cancel each other after composition.

As a result, the lookup tables are just instantiations of the **ASA** structure. Moreover, since the nonlinear layers of AES and DES consist of independent S-boxes, the resulting **ASA** structure is very weak and can be attacked by a number of methods [7]. As demonstrated by Biryukov and Shamir [10], even as large structure as **SASAS** is weak if the S-layers consist of smaller S-boxes. Surprisingly overlooked by the designers of white-box schemes, the generic attack [10] exploits multiset and differential properties of **SASAS** and applies to all the published white-box proposals so far. It appears that the mainstream ciphers are just poor choice for white-box implementations due to high diffusion properties and the way how the key is injected.

To formalize the problem, two notions have been suggested [47, 49]. The **weak white-box implementation** of a cryptographic primitive protects the key and its derivatives i.e. aims to prevent the *key-recovery attack*. This ensures that unauthorized users can not obtain any compact information (e.g. the key or the set of subkeys) to decrypt the protected content.

The **strong white-box implementation** of a primitive protects from the *plaintext-recovery attack*, i.e. does not allow to decrypt given the encryption routine with the embedded key. Such an implementation may replace the public-key cryptosystems in many applications, in particular if it is based on an existing symmetric cipher and is reasonably fast for a legitimate user. The existing white-box implementations of AES and DES [16, 17] do not comply with in this notion, since they are easily invertible, which is strikingly different from the black-box implementations of these ciphers. So far the only proposed candidate is the pairing-based obfuscator scheme with poor performance [47].

The **ASASA**-based designs may not only hide the key for the weak white-box implementation, but also provide non-invertibility aiming for the strong white-box construction.

Our contributions

We continue to explore the design space of compact schemes built from layers of affine mappings and S-boxes. We first note that there is no known generic attack on the 5-layered ASASA scheme with injective S-boxes in the flavour of [10], which makes the ASASA structure a promising framework for future white-box, black-box, and public-key schemes. Based on this principle, we propose and analyze the following constructions in this paper:

- Two *public-key / strong white-box* variants of the ASASA symmetric scheme: one is based on Daemen’s quadratic S-boxes [19] (previously used in various hash functions) and another based on random expanding S-boxes. (Section 2). We explore standard cryptanalytic attacks such as differential, linear and others, the recent decomposition attacks [30], and a new interpolation attack on weakened variants of our schemes (Section 3). We demonstrate that our set of parameters offers a comfortable security margin with respect to the existing attacks.
- A concrete instantiation for a fast symmetric ASASA-based blockcipher with secret S-boxes and affine layers and comparable with AES in its encryption/decryption speed (Section 4).
- A concept of *memory-hard white-box implementation* for a symmetric blockcipher and a concrete family of ciphers with tunable memory requirements (Section 5). It prevents key recovery and requires the adversary to share the entire set of lookup tables to allow an unauthorized user to decrypt. Therefore, the cipher solves the problem of *weak white-box* implementation.

Due to the space limits, some references and attacks on the weakened variants of our schemes are not present in this paper and are available at [8].

2 Asymmetric ASASA schemes: strong white-box and public-key

The first ASASA cryptosystem, designed by Patarin and Goubin, was a public-key scheme with non-bijective S-boxes and was easily broken by Biham, exploiting this property in [5]. Shortly afterwards, Biryukov and Shamir explored multi-layer schemes with bijective S-boxes and demonstrated a generic attack on the structure SASAS with two affine layers [10]. The outer S-boxes are recovered with a variant of the Square attack, whereas the inner affine layers are peeled off with linear algebra methods. It was clearly demonstrated that these properties disappear in larger schemes, and no attack on ASASA or other larger structures has been proposed since.

2.1 Strong white-box security

We start with the notion of the strong white-box security that summarizes the discussion in [49].

Definition 1. Let the pair of algorithms (E, D) be a private-key encryption scheme, which takes key K as parameter. Let \mathcal{O}_{E_K} be a function that computes E_K . We say that \mathcal{O}_{E_K} is a secure strong white-box implementation for E_K if it is computationally hard to obtain \mathcal{D}' equivalent to D_K given full access to \mathcal{O}_{E_K} .

In other words, an adversary should be unable to decrypt given the white-box implementation \mathcal{O}_{E_K} of E_K . This notion closely resembles the definition of a trapdoor permutation used to construct a public-key encryption scheme. As we see, our asymmetric proposals are suitable for both notions.

2.2 Outline

We propose several asymmetric instantiations of the ASASA structure, which may serve both in the white-box and public-key setting. We have not found any reasonable use for lookup tables in this framework² and hence look for polynomial-based S-boxes. In order to keep the reasonable size of the description, we restrict to polynomials of degree two over some finite field, so that the resulting scheme has degree four. This approach brings us to the area of *multivariate cryptography*, which aims to design cryptographic primitives based on multivariate polynomials over finite field.

Let us introduce the following notations. The public key/white-box implementation is exposed as a set of polynomials \mathbf{b} , which is constructed out of the following composition:

$$\boxed{\mathbf{b} = \mathcal{U} \circ \mathbf{a}_2 \circ \mathcal{T} \circ \mathbf{a}_1 \circ \mathcal{S},} \quad (1)$$

where $\mathbf{a}_1, \mathbf{a}_2$ are nonlinear transformations, and $\mathcal{U}, \mathcal{T}, \mathcal{S}$ are affine transformations.

There have been many proposals for nonlinear layers in the ASA structure, and various attacks exploited these choices. Most attacks are not evidently translated into degree 4, as they compute, e.g., differentials of the public key, which are linear functions the ASA case. The notable exception is the decomposition attack [30,31], that will be discussed in Section 2.3.

We offer two fresh ideas for the nonlinear layers in ASASA. The first candidate is the so called χ -function. It derives from invertible cellular automata and was brought into symmetric cryptography by Daemen. To the best of our knowledge, it has never been used in multivariate cryptography.

The second candidate is a set of random injective S-boxes of degree 2. Since the families of low degree permutations are small and do not absorb much randomness, we propose to use expanding S-boxes, which can be key-dependent.

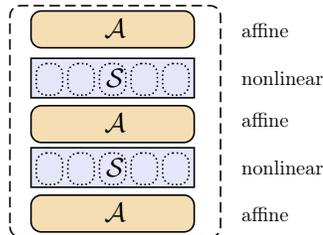


Fig. 1. The ASASA structure: two nonlinear layers surrounded by affine layers.

² So far all attempts to hide a trapdoor in lookup table-based designs failed. We investigated this problem and conjecture that such scheme just does not exist, at least given the state-of-the-art in the design of preimage-resistant functions.

Having the expansion rate of 2, it is rather easy to obtain injective transformations and still keep them quadratic³.

Limitations for expanding schemes. Whatever construction is used, an expanding scheme has a clear limitation in the public-key and white-box setting. It implies that only a tiny subset of potential ciphertexts is decryptable, which makes the encryption and decryption process non-interchangeable. As a result, the expanding scheme can be used for encryption only and can not produce signatures. Also in the white-box context, it can not be used for decrypting the content. On the other hand, it can still be used to ensure tamper-resistance of software [41].

2.3 Defeating decomposition algorithm with perturbations

The authors of recently published decomposition algorithms [30, 31] claim to break ASASA schemes with quadratic nonlinear layers with complexity $O(n^9)$, where n is the number of variables. The decomposition problem is formulated as follows: given a set of polynomials $h = (h_1, \dots, h_u)$ over polynomial ring $\mathbb{K}[x_1, \dots, x_n]$ (\mathbb{K} denoting an arbitrary field) find any $f = (f_1, \dots, f_u)$ and $g = (g_1, \dots, g_n)$ over $\mathbb{K}[x_1, \dots, x_n]$ whose composition is equal to h :

$$h = (h_1, \dots, h_u) = (f_1(g_1, \dots, g_n), \dots, f_u(g_1, \dots, g_n)).$$

and their degree being smaller than h .

In the context of the ASASA structure with quadratic S-boxes, the sets f and g , that are produced by a decomposition algorithm, are linearly equivalent to the internal ASA structures. This does not fully constitute a break, since the adversary still needs to invert both ASA constructions. The proposed algorithms also have not been applied to the parameters and fields that we choose. Nevertheless, it is desirable to find some countermeasure.

Our idea is to introduce some perturbation just after the second S layer in the form of several key-dependent secret polynomials of degree 4. A similar approach has been used by Ding in his modification of C* [21] and HFE [22]. In some cases (notably HFE), the ‘‘perturbation’’ would be identified and removed [25], thanks to a differential attack exploiting properties of the non-linear transformations. The use of perturbation polynomials has been also linked to the LWE (Learning with Error) framework in [32], but the full application of LWE to multivariate cryptography is still to be explored in the future.

Denoting the perturbation polynomials as another nonlinear transformation \mathbf{a}_p we obtain the modified public key \mathbf{b}_p :

$$\mathbf{b}_p = \mathcal{U} \circ [\mathbf{a}_p + (\mathbf{a}_2 \circ \mathcal{T} \circ \mathbf{a}_1 \circ \mathcal{S})], \quad (2)$$

so

$$\mathbf{b}_p(x) = \mathbf{b}(x) + \mathcal{U}\mathbf{a}_p(x).$$

³ Our experiments show that S-boxes with an even smaller rate of 1.75 can be found.

Hence the perturbation polynomials are mixed by the last affine transformation and spread over the public key. The encryption process remains exactly the same, while for decryption we have to guess the values of these polynomials. Suppose that we work over \mathbb{F}_2 so that \mathbf{a}_p is sparse and contains only w polynomials. Let each polynomial be non-zero in $q \cdot 2^n$ points. Then the noise on average consists of qw bit flips, and we guess their positions after about $\binom{w}{qw}$ attempts. For instance, \mathbf{a}_p with 8 non-zero polynomials of weight $\approx 2^{n-1}$ requires 2^6 trial decryptions on average.

We distinguish true plaintexts from false ones either by recomputing the perturbation polynomials or by using expanding S-boxes so that noisy bits prohibit inversion. Padding the plaintexts with zero bits also helps but disallows turning encryption to decryption. The position of noisy bits does not matter much, since it would be concealed by the affine transformation. However, if we filter out noise with expanding S-boxes, it makes sense to spread the noisy bits so that an S-box can still be inverted in the presence of noise.

2.4 χ -scheme

Our first idea was to build the nonlinear transformation out of a popular quadratic S-box χ [19, Section 6.6.2], which has been used in several hash functions including SHA-3/Keccak [4]. The transformation χ can be defined for every odd length $k = 2t + 1$ and has the following features:

- It has degree 2 in the forward direction, but degree $(t + 1)$ in the backward direction.
- It can be efficiently inverted for every size [8].
- Its differential and linear properties have been widely studied [19].

The S-box χ of length k is defined as follows:

$$\chi(x_0, x_1, x_2, \dots, x_{k-1}) = (y_0, y_1, y_2, \dots, y_{k-1}),$$

where

$$y_i = x_i \oplus x_{i+1}x_{i+2} \oplus x_{i+2},$$

and indices are computed modulo k .

Regardless of the χ length (and hence the size of the S-box), we can formulate properties of the whole scheme and its features:

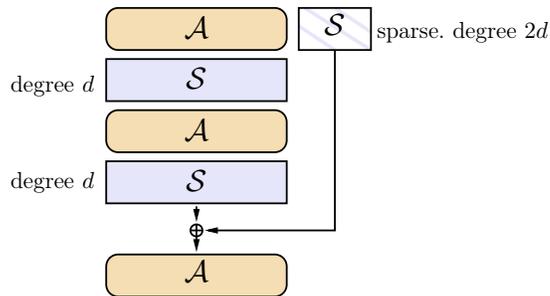


Fig. 2. Small perturbations to defeat decomposition attacks as injection of sparse high-degree polynomials.

1. For the standard block size of 128 bits, we get approximately (since the S-box size might not divide 128) 2^7 input variables. Thus each output coordinate of \mathbf{b} is a polynomial of degree 4 of 2^7 variables, has about $\binom{2^7}{4} \approx 2^{24.5}$ terms, so the full scheme description is about $2^{24.5+7-3} = 2^{28.5}$ bytes, or 300 MBytes.
2. The private key size is much more compact, and is dominated by three matrices with 2^{14} bits each (hence 2^{13} bytes in total). If the matrices are deterministically produced out of some secret key (e.g., 128-bit), the description is even smaller.
3. The inverse polynomial of our schemes has degree $(t+1)^2$ for S-boxes of size $2t+1$.

The S-box size (length of χ) has negligible effect on the performance because of the internal structure of χ and its inverse. Hence it only affects the security of the scheme. We choose a single S-box \mathbf{a} of length 127, so that its inverse has degree 64, and will show later that a system with small S-boxes is insecure.

In order to defeat decomposition algorithms and hide the ASASA structure we suggest using perturbation polynomials. More precisely, we propose 24 random polynomials of degree 4 for the perturbation layer \mathbf{a}_p . We pad each plaintext with 8 zero bits, so that for each guess the probability to fit the padding is 2^{-8} . As a result, we get 2^{16} candidate plaintexts, and then check if we correctly computed the noise. This filters out all wrong plaintexts with probability $1 - 2^{-8}$.

Overall security. We have found a number of attacks on the χ -scheme in different variants (Section 3), so it appears that its algebraic structure yields it vulnerable. Nevertheless, the variant with added perturbation remains unbroken, and we offer it as cryptanalytic challenge, but not for the practical use. We expect the perturbation theory to develop in the near future, which would suggest a more secure set of parameters.

2.5 Scheme with expanding S-boxes

This variant provides a more compact description of the scheme since we may switch to a larger field. First, we want the nonlinear layer be a degree-2 polynomial over $\mathbb{F}_q, q > 2$, and define the linear (affine) transformations over the same field. Though a few examples of bijective transformations of degree 2 over a field not equal to \mathbb{F}_2 exist [36], they appear to be vulnerable to Groebner basis attacks in our own experiments. As a solution, we suggest expanding S-boxes, whose output is twice as big as the input. It is rather easy to design *injective* S-boxes of degree 2 with this property. Indeed, a random function with expansion rate 2 has no collisions with probability around $1/2$, and hence there are enough injective transformations of the desired form.

Here is the summary of the scheme:

- Input length 128 bits (32 variables), output length 512 bits (128 variables);
- All polynomials and affine transformations are defined over \mathbb{F}_{16} ;

- S-boxes map 16 bits to 32 bits and hence are described by 8 degree-2 polynomials over \mathbb{F}_{16} of four variables. The inverse is computed with lookup tables of size 2^{16} .
- The first nonlinear layer has 8 S-boxes and doubles the state size to 256 bits. The second layer has 16 S-boxes and further doubles to 512 bits. Accordingly, the affine transformations S, T, U operate on 128-, 256-, and 512-bit states, respectively.

The output of the scheme is a set of 2^7 degree-4 polynomials over \mathbb{F}_{16} over 32 input variables (each variable is encoded with 4 bits). There are $\binom{2^5}{4} \approx 2^{16.5}$ possible terms, hence, taking 4-bit constants into account, each polynomial is described by $2^{20.5}$ bits, or $2^{20.5+7-3} = 2^{24.5}$ bytes, which is about 24 MBytes.

The private key is smaller: affine layers contain $2^{7+7+1} + 2^{6+6+1} + 2^{5+5+1} \approx 2^{14.2}$ elements of \mathbb{F}_{16} . The 48 S-boxes are described as $2^{5.5+3}$ polynomials of $21 \approx 2^{4.5}$ terms each, hence 2^{13} elements, plus a few noise polynomials. In total, the private key fits into 2^{14} bytes.

We also suggest using perturbation polynomials here. Due to the large expansion rate, we can use rather dense perturbation layer \mathbf{a}_p and still ensure a unique decryption. We use two random polynomials over \mathbb{F}_{16} of degree four at each S-box, hence 32 polynomials in total. While decrypting we face $16^2 = 2^8$ options for each S-box output. As a result, the probability of having non-unique decryption of the last S layer is $2^4 \cdot 2^{16+8-32} = 2^{-4}$, and if this happens the next layer filters out wrong candidates.

As we already mentioned, the expanding character of the scheme allows only public-key and white-box encryption, but not signature generation.

3 Security analysis of our white-box/public-key schemes

In this section we apply various attacks to weakened versions of our schemes, thus demonstrating the design rationale behind them. We demonstrate that the added perturbations are crucial in both schemes, and that they must be secret. We also show that S-boxes in the χ -scheme must be large, that linearity (in contrast to affinity) of \mathbf{A} may weaken the scheme, and that the expanding S-boxes should not be biased (these results are presented mainly in [8]). Our attacks are summarized in Table 1.

These attacks allow us to evaluate the security margin of the unbroken variants of our schemes. Since only the perturbation protects the χ -scheme from a number of practical attacks, we conclude that it is rather fragile, but might become a good candidate for a strong white-box implementation when the complexity of generic algorithms applied to the perturbed version is better understood. In contrast, the expanding scheme appears to be more resistant to generic attacks, and we propose it as a ready-to-use public-key encryption scheme and a strong white-box implementation.

Weakening	Attack complexity	Attack type	Reference
Expanding scheme			
Public perturbation	$2^{45} + D$	Interpolation	Section 3.2
Biased S-boxes (bias= 1/8)	2^{88}	LPN	Section 3.4
χ -scheme			
Public perturbation	$2^{57} + D$	Interpolation	Section 3.2
No perturbation	$\approx 2^{40}$	Groebner-basis	Section 3.3
Small S-boxes	2^{45}	Algebraic	[8]

Table 1. Summary of our attacks on the weakened versions of our schemes. D stands for the complexity of decomposition attacks.

3.1 Generic attacks

Given the public-key of a multivariate scheme, an attacker may directly try to solve the multivariate polynomial equations using a generic algorithm. If the public-key is a vector of m polynomials in n over \mathbb{F}_q , then a plaintext can always be found by exhaustive search in time $\mathcal{O}(q^n)$. The other main family of algorithms to solve systems of polynomial equations are Groebner-basis algorithms, such as Buchberger’s algorithm and all its derivatives [27, 28].

Without going into details (the interested reader is referred to a standard textbook such as [18]), given a system of polynomial equations $f_1 = \dots = f_m = 0$ in x_1, \dots, x_n , a *Groebner basis* of the ideal spanned by the f_i ’s is an equivalent system of equations with nice properties. If the system admits a single solution (a_1, \dots, a_n) , then a Groebner basis is precisely the vector of polynomials: $x_1 - a_1, \dots, x_n - a_n$. It follows that if a Groebner basis can be computed, then the system of equations can be solved.

Groebner basis algorithms work by performing *polynomial elimination*, i.e., by trying to eliminate some terms by summing suitable multiples of other polynomials. The complexities of these algorithms are difficult to analyze [2]. They are essentially exponential in the highest degree reached by the polynomials created and manipulated by the algorithms during their execution. On “generic” systems of n equations in n variables, this degree is typically n . However, in some special cases it can be lower. For instance, the first HFE Challenge could be broken because in HFE, for some ranges of parameters, this degree was roughly $\mathcal{O}(\log n)$.

3.2 Interpolation attack on the ASASA scheme with public perturbation polynomials

We stressed that the perturbation polynomials must be secret. A reader may wonder why this is required, since these polynomials are seemingly mixed by the last affine transformation \mathcal{U} .

In this subsection we outline an attack that peels off the perturbation polynomials and recovers the core ASASA scheme in almost practical time. Suppose we work over a field \mathbb{F}_2 and the scheme adds perturbation polynomials at r bit positions after the nonlinear transformation \mathbf{a}_2 (cf. Eq. (2)), and the total number of variables in the scheme is n . Then we collect N plaintexts x_i such that

$$\mathbf{a}_p(x_i) = 0.$$

Since polynomials of \mathbf{a}_p do not have any structure, finding a common zero is an NP-hard problem, and we expect that 2^r plaintexts must be tried to find a right one. Hence the naive complexity of this step is $N2^r$ evaluations⁴ of \mathbf{a}_p .

Then we evaluate the right plaintexts on the perturbed scheme \mathbf{b}_p . Since \mathbf{a}_p is zero, we have

$$\mathbf{b}_p(x_i) = \mathcal{U} \circ \mathbf{a}_2 \circ \mathcal{T} \circ \mathbf{a}_1 \circ \mathcal{S}(x_i).$$

Therefore, we know the evaluation of the ASASA scheme without perturbations on N plaintexts. Since the scheme has degree 4, the polynomial coefficients can be recovered by the Lagrange interpolation. There are $\sum_{i=0}^4 \binom{n}{i}$ monomials of degree 4 or smaller, hence N must slightly exceed $\sum_{i=0}^4 \binom{n}{i}$ to allow for linear dependencies among plaintexts. For the typical value $n = 2^7$ we need about 2^{25} right plaintexts to fully recover the core ASASA polynomials and then launch the decomposition attack. However, the interpolation itself is not a trivial procedure, since we deal with a multivariate function. Only recently an algorithm with complexity quadratic in the number of monomials has been proposed [1]. Equipped with it, we recover a single polynomial in 2^{50} bit operations, and the entire \mathbf{b}_p in 2^{57} operations. In turn, 2^{25} right plaintexts can be obtained for 16 noisy bits in 2^{41} evaluations of \mathbf{a}_p , and for 24 noisy bits – in 2^{49} evaluations, which is close to 2^{55} bit operations. Therefore, the total complexity of recovering (x) is about 2^{57} bit operations.

This attack clearly shows that the perturbation polynomials must not be public and should not have any structure that would allow the adversary to find their common zeros. We do not see how the attack can be applied to secret polynomials.

3.3 Algebraic attack on the plain χ -scheme

Although χ has been used successfully in the symmetric world, it turns out to be a complete disaster in a multivariate context. An ASA construction where $S = \chi$, with $n = 127$ variables over \mathbb{F}_2 is broken in a few seconds by a direct Groebner basis computation. A two-layer ASASA construction is not more secure, and can be broken in less than two hours using the implementation of the F4 algorithm of the MAGMA computer algebra system [14] (and 100Gbytes of RAM). This happens because a Groebner basis can be computed by manipulating polynomials of small, constant degree (typically 3 or 6).

⁴ Finding subsequent solutions might be easier, but this step is not a dominant in our attack complexity.

Let us give some detailed explanation for the insecurity of the ASA construction. We work within the polynomial ring $R = \mathbb{F}_2[x_0, \dots, x_{n-1}]$, and we consider the ideal of R :

$$\mathcal{I} = \langle f_0, \dots, f_{n-1}, \quad x_0^2 - x_0, \dots, x_{n-1}^2 - x_{n-1} \rangle$$

where $f_i = x_i + x_{i+2} + x_{i+1}x_{i+2} + a_i$ (all indices are taken modulo n), and where the a_i are constants. Any solution (in the x_i 's) making all the polynomials in \mathcal{I} vanish simultaneously, is a solution of $\chi(x_1, \dots, x_n) = (a_0, \dots, a_{n-1})$. Such a solution always exists, and is unique.

We will show that there are many linear polynomials in this ideal, and that they can be “easily” discovered (by manipulating small-degree polynomials). Indeed:

$$x_{i+1} \cdot f_i - x_{i+2} \cdot (x_{i+1}^2 - x_{i+1}) + f_{i-1} = (x_{i-1} + x_{i+1}) - (a_{i-1} + a_{i+1})$$

The expression on the left-hand side is a polynomial combination of elements of \mathcal{I} , therefore it belongs to \mathcal{I} . As a consequence, the linear polynomial on the right-hand side can be found inside \mathcal{I} after performing a few steps of polynomial elimination on polynomials of degree less than 3.

After these n linear relations have been found, another few steps of polynomial elimination allows all the variables but one to disappear. This shows that a Groebner basis of the ideal \mathcal{I} can be computed in polynomial time. Now, performing a (random) linear change of coordinate in \mathcal{I} , or replacing the generators of \mathcal{I} by (random) linear combinations thereof does not change this fact. As a conclusion, the ASA construction, where S is the χ -function, falls victim to a direct algebraic attack, by running any Groebner basis algorithm on the equations defining the “white-box”.

This reasoning extends to the ASASA construction where both non-linear layers are χ (however, this time the degree is 6). It is an open question of how much the added perturbation slows the Groebner-basis attacks (our implementation does not break the selected noise parameters in reasonable time).

3.4 Attack on the expanding scheme with biased S-boxes

If S-box output bits are biased, an attack exploiting this bias can be applied. The last affine transformation can be viewed as affine over \mathbb{F}_2 , so the further analysis without loss of generality applies to any field of characteristic two.

We target a single biased bit b after the second layer of expanding S-boxes: the probability $\mathbb{P}[b = 1]$ of its equality to 1 is equal to $p \neq \frac{1}{2}$. If y is a ciphertext, then following previous notations, the biased bit is the b -th component of $U^{-1} \cdot y$. In other terms, if \bar{u} denotes the b -th line of U^{-1} , then $\langle \bar{u}, \bar{y} \rangle = b$.

Now, assume we collect a large number (say N) of ciphertexts. We stack them vertically into a matrix C , which thus has N rows. Let us also assume that b is biased towards zero. Then we have the “noisy linear system”:

$$\bar{u} \cdot C = \bar{e},$$

where \bar{e} is a vector of i.i.d. random variables following the Bernoulli distribution with mean p . Recovering \bar{u} is exactly an instance of the Learning Parity with Noise (LPN) problem.

The best known algorithms to solve LPN are variants of the BKW algorithm [11], whose complexity is of order $\mathcal{O}(2^{n/\log n})$. The only actual implementation (along with algorithmic improvements) is described in [37], and some more tweaks are given in [3]. The actual complexities of these algorithms depend on the bias (their efficiency decreases when the bias gets closer to zero).

With $n = 512$ variables, and if $\mathbb{P}[b = 1] = 1/8$, then the implementation of [37] is said to require 2^{80} bits of memory (plus the time needed to sort this much memory 80 times). However, time-memory tradeoffs, plus algorithmic improvements, allow [3] to conclude that the same problem can be solved in 2^{59} bits memory and less than 2^{100} bits operations. If 2^{80} bits of memory are available, then the running-time could be decreased to 2^{88} bit operations.

This beats more naive approaches, such as, for instance, enumerating all the possible sparse possibilities for the first n components of \bar{e} , and solving the corresponding linear system for each trial. The above instance would require more than 2^{120} operations to be solved using the naive approach.

Of course, the attack has to be repeated for each row of U^{-1} , and possibly twice for each row (assuming that the targeted bit is biased towards zero, or towards one). Note that the above estimates are extremely pessimistic; in random expanding S-boxes of degree 2, the biases we observed experimentally are much lower than what was used above (we observed $\mathbb{P}[b = 1] \approx 0.49$).

After U is recovered, we can view the output of expanding S-boxes, and are likely to recover them by interpolation due to low degree.

4 Black-box ASASA schemes

Given rather low performance and large key size of the public-key ASASA schemes, a reader may wonder if significant performance increase can be achieved with lower security goals. We answer this question twofold. First, we propose a generic black-box symmetric cipher based on the five layer ASASA. The cipher is expected to have a very fast software implementation thanks to vector instructions in modern processors. Secondly, we use a small version of this cipher as a building block in achieving weak white-box security (Section 5).

4.1 Design

We propose a symmetric cipher with a classical set of parameters, widely used in AES and other designs. It has a block of $n = 128$ bits with $m = 8$ bit S-boxes and a choice of key-sizes 128 – 256 bits. Let us outline specific parameters for linear and nonlinear layers.

Affine layers. A key-dependent $n \times n$ affine transformation can be produced out of the master key K by any secure key derivation function H_K (for example a fast

stream cipher, or a block cipher in the counter mode (more details in [8]), and checking that the resultant matrix is invertible, this can be done in $O(n^3)$ steps, and we also generate an n -bit constant ⁵. The branch number of the matrix [20] determines the minimum number of active S-boxes in a differential trail, and thus the upper bound on the trail probability. Since the matrix is random, we expect the branch number to be close to the maximum possible (the number of S-boxes n/m plus one). Note that for each affine layer a new matrix is generated.

Nonlinear layers. Typically, nonlinear layers of symmetric ciphers consist of several small S-boxes, which have a compact description [12, 20]. For the ASASA scheme we propose to use 32 randomly generated 8-bit invertible S-boxes, which are all different and key-dependent. We note that efficiency of generic attacks on the SASAS structure [10] increases if smaller S-boxes are used, and thus it may be interesting in the future to explore full block size non-linear layers, for which such attack would not work.

Large S-box alternatives. The choice for large block algorithmic S-boxes is surprisingly limited. Unless the S-boxes are themselves multi-layer permutations (e.g., fixed-key ciphers), a compact description is typically delivered in the algebraic form as a function over an appropriate finite field. The resulting *permutation polynomials* have become an active research topic in the recent years. The well known example is X^{2^k+1} over \mathbb{F}_{2^n} (scheme C* [39]); the more recent and interesting include $(X^{2^k} + X + a)^{-1} + X$ over \mathbb{F}_{2^n} by Zeng et al. [50] (derived from Helleseith-Zinoviev polynomials) (more references in [8]). It thus can be an interesting second challenge to break the symmetric ASASA scheme with known block-wide non-linear layers. Note however that fixed S-boxes do not offer implementation advantage and thus we would keep S-boxes secret and randomly generated in the main variant of our scheme. Implementation details

Implementation. The implementation details can be found in the full version of the paper [8].

4.2 Security analysis

Differential and linear attacks. We expect the secret linear layers to hide all differential [6] and linear [38] properties of the cipher, since it becomes impossible to figure out any high-probability differential or linear trail. It can be argued, however, that the existence of high-probability characteristics may lead to efficient distinguishers. For instance, Dunkelman and Keller showed in [26] that if for every α the differential probabilities $\{\alpha \rightarrow \beta\}$ are much higher (or much lower) than for a random permutation, then this can be used as a distinguisher. The

⁵ Non-anonymous final version of this paper will link to implementations of our schemes and challenges gradually increasing complexity for the interested crypt-analysts.

authors further suggested the parameter of effective linearity that essentially measures the average probability of the boomerang difference quartet (α, K) over all possible α, K and is supposed to take even unknown characteristics into account. In the full version of this paper [8] we show that all these methods do not lead to attacks much faster than exhaustive key search.

Algebraic attacks. We expect the random S-box of width m to have the algebraic degree $m - 1 = 7$. As a result, the entire scheme can be represented by a polynomial of degree 49 over \mathbb{F}_2 . As observed by Meier [40], the low degree can be detected by applying differentials of the same order to the ciphertext. Therefore, an attacker can distinguish the ASASA construction from random given 2^{49} chosen data and time. However, this does not lead to the disclosure of the plaintext, and it is unclear how this property can be exploited.

Other attacks. The boomerang and impossible differential attack can be also of concern. We have tried basic and improved versions of these attacks, and in all cases the randomness of the affine layers prevented us from mounting an attack. However, it is possible to build boomerang quartets in the known-key setting by activating a single S-box at both sides of the boomerang. Whether such properties can be carried out to the secret-key setting is the object of the future research. Impossible differential attacks typically rely on truncated differentials with probability 1 which exist in some ciphers due to incomplete diffusion. Since in our case the random affine layers provide complete diffusion and since the entrance into and the exit from the scheme are both guarded by these affine layers, chosen plaintext attacks have little chance of predicting truncated values somewhere inside the scheme.

Our scheme should be more secure than a two-round Even-Mansour cipher [13], where the subkeys are simply XORed to the internal state (as opposed to applying a full-blown secret affine transformation to the internal state). The recent attack on the 2-round Even-Mansour [23] explicitly requires the access to the internal permutation and thus can not be immediately used in our setting.

The meet-in-the-middle attacks [33] do not apply to our scheme, because the amount of key material used to compute any matching variable is too large (several S-boxes and a large part of the affine transformation). The cube attacks do not apply, since there is no compact polynomial representation of the scheme.

Structural attack. Finally, we investigate the structural attack from [10]. We will see that even though it does not apply to the 128-bit ASASA cipher, it allows to bound the security level of schemes with smaller block, which are used in Section 5. First, we recall the main property preserved by the SASA structure with m -bit S-boxes:

Theorem 1 ([10]). *Let $\{x_1, x_2, \dots, x_{2^m}\}$ be the inputs to the SASA structure such that the input bits to one S-box take all possible combinations whereas the other bits are constant. Then the XOR of all outputs is the all-zero bit vector:*

$$\bigoplus_i \text{SASA}(x_i) = (0, 0, \dots, 0).$$

Now consider the input y to some m -bit S-box in the first S layer of the ASASA scheme E . It is an affine function of the scheme input y :

$$y = M \cdot x \oplus c,$$

where M is an $(m \times n)$ -matrix and c is a constant. Let L be an $(n \times n)$ -matrix such that

$$M \times L = (M' \ 0 \ 0 \ \dots \ 0,) \tag{3}$$

where M' is a $(m \times m)$ - submatrix. If we apply E to $L \cdot x$, then y depends only on the first m bits of x . Thus we compose inputs $\{x_1, x_2, \dots, x_{2^m}\}$ as in Theorem 1, multiply them by L and apply E . The output bits must sum to 0 bit-wise. This property allows to recover the outer affine layer and eventually all the components of E . Equation (3) holds with probability $2^{-(n-m)m}$, which makes the attack impractical for large n . However, for small n it might be efficient. Therefore, the maximum security level of the ASASA scheme with n variables and m -bit S-boxes does not exceed $(n - m)m$ bits.

For our design, this gives the upper bound of 960 bits, which is far larger than the key length which is used to generate the affine and nonlinear layers. As a result, we claim the 128-bit security level for our design, even though a small factor over exhaustive key search might be saved by biclique attacks or by exploiting the full codebook. This is still higher security than what is offered by 2048 bit RSA.

5 Proposal for weak white-box security: ASASA-based block cipher

5.1 Weak white-box security

Definition 2. Let the pair of algorithms (E, D) be a private-key encryption scheme, which takes key K as a parameter. We call $\mathfrak{F}(K)$ the equivalent key set for key K , if from any element from $\mathfrak{F}(K)$ it is easy to get an algorithm equivalent to E_K , i.e. there is an (efficient) algorithm

$$\mathcal{A}(\mathcal{K}) \rightarrow \mathcal{E}',$$

where \mathcal{E}' equivalent to E_K .

Definition 3. The function \mathcal{O}_{E_K} is a T -secure weak white-box implementation for E_K if it is computationally hard to obtain $\mathcal{K} \in \mathfrak{F}(K)$ of length less than T given full access to \mathcal{O}_{E_K} .

In other words, an adversary who gets a secure weak white-box implementation is unable to find out any compact (shorter than T) equivalent representation of it. In the practical sense, an adversary who wants to share a protected implementation of the encryption routine, would have to share the entire code. Such ciphers are motivated by DRM applications, which aim to prohibit the users of

protected content from sharing the information needed to decrypt it. Clearly, in this context there is little practical difference between sharing the key and sharing, say, the set of subkeys as long as the other cipher operations are independent of the key. Therefore, naive methods of key protection, e.g. transforming it with a preimage-resistant hash function, would not prevent an attack. Ideally, the adversary would have to isolate and extract the entire decryption routine, which might be hard per se.

5.2 Weak white-box cipher proposal

In this section we propose a blockcipher family, which conforms to the weak white-box security notion, so that it is computationally infeasible to derive a key or any other compact secret information from the white-box implementation.

We further say that the white-box implementation is *memory-hard* if it requires a pre-specified and large enough amount of memory in the spirit of memory-hard key-derivation functions [44]. This concept is even stronger than the T -secure weak WB implementations, as an adversary is unable to reduce the implementation size at all and thus would have to publish the entire set of lookup tables. In contrast to earlier white-box designs, we offer a set of ciphers with a wide range of memory requirements.

Our memory-hard cipher consists of a number of smaller components, which are exposed as lookup tables in the white-box implementation. Each component is either a small-block ASASA cipher, adapted from the construction in Section 4, or just a single S-box. The S-boxes are minimum 8-bit wide to avoid equivalence problems [9], but 10-, and 12-bit ones are also used. All S-boxes and affine layers are derived in a deterministic way from the secret key. In

fact, it is enough to have linear, not affine, layers, since the constant can be kept in the S-boxes. To estimate memory requirements, we assume for simplicity that each table output fits an integer number of bytes (e.g., 2 bytes for 10-bit S-boxes).

We propose the SPN structure for the cipher, i.e. we alternate layers of smaller ciphers (denote their number by R) with a public linear transformation \mathcal{L} . Any transformation with good diffusion shall be fine. Recalling that AES can be partitioned into 5 rounds with 4 32-bit Super S-boxes in each, we propose R layers for similar security margin. The cipher’s pseudocode is as follows (Figure 3):

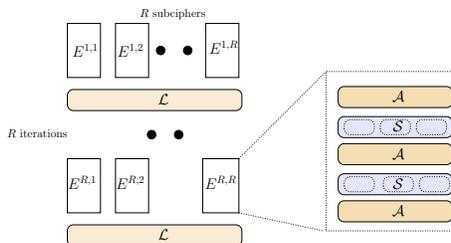


Fig. 3. Blockcipher family for weak white-box security.

- Repeat R times;
 - Apply R parallel ASASA-based distinct blockciphers;
 - Apply the linear transformation \mathcal{L} to the entire state.

We outline specific parameters and memory requirements for 64- and 128-bit blockciphers in Table 2. The 16-bit S layer has two 8-bit S-boxes, the 18-bit – 8-bit and 10-bit S-boxes, the 20-bit – two 10-bit S-boxes, and the 24-bit – three 8-bit S-boxes. We see that whereas the black-box implementation is a few dozen KBytes, the white-box implementation can be made large enough in the range from 2 MBytes to several GBytes.

Security analysis. Our ASASA components have very small block and only a few S-boxes in the S layer. Some attacks that are infeasible on the 128-bit block, may have practical complexity on the 16-bit block. The best attack we could find was presented in Section 4.2 and has complexity $2^{(n-m)m}$ for m -bit S-boxes and the n -bit block. As a result, the 16-bit ASASA components with 8-bit S-boxes have maximum security level of 64 bits, the 20-bit components — 100 bits, and 24-bit components with 8-bit S-boxes — 128 bits.

An easy way to increase the security level is to add two more layers, thus producing ASASASA components. This yields a 50% increase of the private key size, but no increase in the white-box implementation size. Since we have not found a way to expand our attack to this structure, we conjecture its security level to 128 bits. In Table 2 we provide both variants so that a protocol designer may choose between them according to his own requirements.

Rows	Component	Components in row	Security level (bits)	White-box memory	Black-box memory
64-bit block					
4	ASASA	4×(16-bit)	64	2MB	16 KB
4	ASASA	3×(18-bit) + 10-bit	64	9 MB	32 KB
128-bit block					
8	ASASASA	8×(16-bit)	128	8 MB	96 KB
8	ASASA	24-bit + 6×(16-bit) +8-bit	64	384 MB	64 KB
5	ASASASA	4×(28-bit) + (16-bit)	128	20 GB	130 KB

Table 2. Parameters and memory requirements of white-box and black-box implementations for the 128-bit blockcipher. We assume that n -bit component occupies $\lceil \frac{n}{8} \rceil 2^n$ bytes of memory in the white-box implementation.

Acknowledgements

We thank Willi Meier for fruitful comments and the discussion about the results of this paper. We also thank anonymous reviewers for comments, which helped to improve the paper.

6 Conclusion

We have explored deeply the state of the art in black-box, white-box, and multivariate public key cryptography, and concluded that the ASASA structure is the minimal generic construction which is still unbroken. We constructed cipher candidates for all these settings. We designed two ASASA schemes for public-key cryptography based on multivariate polynomials. We showed how to avoid existing attacks on multivariate schemes, including the recent powerful decomposition algorithms by adding appropriate perturbation functions. In the traditional black-box setting we offered a cryptanalytic challenge of a fast cipher with small random S-boxes and random affine layers.

We proposed several solutions for white-box cryptography, both in weak and strong security notions. In the weak model, we designed a memory-hard cipher, which prohibits key extraction and requires an adversary to spend a large, pre-defined amount of memory. It is based on small ASASA components. We showed how our multivariate schemes can be used as strong white-box implementations, as they are not invertible without the key and allow fast encryption and decryption for legitimate users. We compare the implementation size of our schemes with other unbroken MQ-systems in Table 3.

Our findings indicate a number of future research directions. First, it would be interesting to explore algorithmic large S-boxes in the black-box ASASA structure, e.g. instantiated with recently found permutation polynomials. Secondly, a theory of perturbation layers as a countermeasure to generic decomposition algorithms needs to be developed, possibly along the concept of LWE (Learning with Error). Thirdly, we suggest investigating the actual security level of small-block (16-,20-, 24-bit) ASASA schemes to figure out which components are suitable for weak white-box implementations. Finally, open question is to develop constructions with smaller descriptions (e.g., within 1 MByte), which are bijective, suitable for digital signatures, and allow strong white-box implementations.

Scheme	Field	# vars	# polys	Degree	Private key	PK/white-box	Ref.
Black-box ASASA	\mathbb{F}_2	128	-	-	14 KB	196 KB	Sec. 4
χ -scheme*	\mathbb{F}_2	127	127	4	8 KB	300 MB	Sec. 2
Expanding scheme	\mathbb{F}_{16}	32	128	4	16 KB	24 MB	Sec. 2
Memory-hard cipher	\mathbb{F}_2	64	-	-	16-32 KB	2-8 MB	Sec. 5
Memory-hard cipher	\mathbb{F}_2	128	-	-	64-130 KB	8 MB – 20 GB	Sec. 5
HFE	\mathbb{F}_{16}	64	64	2	48 KB	520 KB	[43]
UOV	\mathbb{F}_{256}	78	26	2	71 KB	80 KB	[34]

* — several variants broken in this paper.

Table 3. Our schemes in comparison, along with (presumably) secure parameters for UOV and HFE.

References

1. Frederik Armknecht, Claude Carlet, Philippe Gaborit, Simon Künzli, Willi Meier, and Olivier Ruatta. Efficient computation of algebraic immunity for algebraic and fast algebraic attacks. In *EUROCRYPT'06*, volume 4004 of *Lecture Notes in Computer Science*, pages 147–164. Springer, 2006.
2. Magali Bardet, Jean-Charles Faugère, and Bruno Salvy. On the complexity of Gröbner basis computation of semi-regular overdetermined algebraic equations. In *Proc. International Conference on Polynomial System Solving (ICPSS)*, pages 71–75, 2004.
3. Daniel J. Bernstein and Tanja Lange. Never trust a bunny. In *RFIDSec*, volume 7739 of *Lecture Notes in Computer Science*, pages 137–148. Springer, 2012.
4. Guido Bertoni, Joan Daemen, Michael Peeters, and Gilles Van Assche. The Keccak reference, version 3.0, 2011. <http://keccak.noekeon.org/Keccak-reference-3.0.pdf>.
5. Eli Biham. Cryptanalysis of Patarin’s 2-round public key system with S-Boxes (2R). In *EUROCRYPT'00*, volume 1807 of *Lecture Notes in Computer Science*, pages 408–416. Springer, 2000.
6. Eli Biham and Adi Shamir. *Differential Cryptanalysis of the Data Encryption Standard*. Springer, 1993.
7. Olivier Billet, Henri Gilbert, and Charaf Ech-Chatbi. Cryptanalysis of a white box AES implementation. In *Selected Areas in Cryptography'04*, volume 3357 of *Lecture Notes in Computer Science*, pages 227–240. Springer, 2004.
8. Alex Biryukov, Charles Bouillaguet, and Dmitry Khovratovich. Cryptographic schemes based on the ASASA structure: Black-box, white-box, and public-key. Cryptology ePrint Archive, Report 2014/474, 2014. <http://eprint.iacr.org/>.
9. Alex Biryukov, Christophe De Cannière, An Braeken, and Bart Preneel. A toolbox for cryptanalysis: Linear and affine equivalence algorithms. In *EUROCRYPT'03*, volume 2656 of *Lecture Notes in Computer Science*, pages 33–50. Springer, 2003.
10. Alex Biryukov and Adi Shamir. Structural cryptanalysis of SASAS. In *EUROCRYPT'01*, volume 2045 of *Lecture Notes in Computer Science*, pages 394–405. Springer, 2001.
11. Avrim Blum, Adam Kalai, and Hal Wasserman. Noise-tolerant learning, the parity problem, and the statistical query model. *J. ACM*, 50(4):506–519, 2003.
12. Andrey Bogdanov, Lars R. Knudsen, Gregor Leander, Christof Paar, Axel Poschmann, Matthew J. B. Robshaw, Yannick Seurin, and C. Vikkelsoe. PRESENT: An ultra-lightweight block cipher. In *CHES'07*, volume 4727 of *Lecture Notes in Computer Science*, pages 450–466. Springer, 2007.
13. Andrey Bogdanov, Lars R. Knudsen, Gregor Leander, François-Xavier Standaert, John P. Steinberger, and Elmar Tischhauser. Key-alternating ciphers in a provable setting: Encryption using a small number of public permutations. In *EUROCRYPT'12*, volume 7237 of *Lecture Notes in Computer Science*, pages 45–62. Springer, 2012.
14. Wieb Bosma, John J. Cannon, and Catherine Playoust. The Magma Algebra System I: The User Language. *J. Symb. Comput.*, 24(3/4):235–265, 1997.
15. Vincent Carlier, Hervé Chabanne, and Emmanuelle Dottax. Grey box implementation of block ciphers preserving the confidentiality of their design. *IACR Cryptology ePrint Archive*, 2004:188, 2004.

16. Stanley Chow, Philip A. Eisen, Harold Johnson, and Paul C. van Oorschot. White-box cryptography and an AES implementation. In *Selected Areas in Cryptography'02*, volume 2595 of *Lecture Notes in Computer Science*, pages 250–270. Springer, 2002.
17. Stanley Chow, Philip A. Eisen, Harold Johnson, and Paul C. van Oorschot. A white-box DES implementation for DRM applications. In *Digital Rights Management Workshop*, volume 2696 of *Lecture Notes in Computer Science*, pages 1–15. Springer, 2002.
18. David A. Cox, John Little, and Donal O'Shea. *Ideals, Varieties, and Algorithms: An Introduction to Computational Algebraic Geometry and Commutative Algebra, (Undergraduate Texts in Mathematics)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 1991.
19. Joan Daemen. *Cipher and Hash Function Design Strategies based on linear and differential cryptanalysis*. PhD thesis, Katholieke Universiteit Leuven, Leuven, Belgium, March 1995.
20. Joan Daemen and Vincent Rijmen. *The Design of Rijndael. AES — the Advanced Encryption Standard*. Springer, 2002.
21. Jintai Ding. A new variant of the Matsumoto-Imai cryptosystem through perturbation. In *Public Key Cryptography'04*, volume 2947 of *Lecture Notes in Computer Science*, pages 305–318. Springer, 2004.
22. Jintai Ding and Dieter Schmidt. Cryptanalysis of HFEv and internal perturbation of HFE. In Serge Vaudenay, editor, *Public Key Cryptography'05*, volume 3386 of *Lecture Notes in Computer Science*, pages 288–301. Springer, 2005.
23. Itai Dinur, Orr Dunkelman, Nathan Keller, and Adi Shamir. Key recovery attacks on 3-round Even-Mansour, 8-step LED-128, and full AES². In *ASIACRYPT'13*, volume 8269 of *Lecture Notes in Computer Science*, pages 337–356, 2013.
24. Vivien Dubois, Pierre-Alain Fouque, Adi Shamir, and Jacques Stern. Practical cryptanalysis of SFLASH. In *CRYPTO*, volume 4622, pages 1–12. Springer, 2007.
25. Vivien Dubois, Louis Granboulan, and Jacques Stern. Cryptanalysis of HFE with internal perturbation. In Tatsuaki Okamoto and Xiaoyun Wang, editors, *Public Key Cryptography'07*, volume 4450 of *Lecture Notes in Computer Science*, pages 249–265. Springer, 2007.
26. Orr Dunkelman and Nathan Keller. A new criterion for nonlinearity of block ciphers. *IEEE Transactions on Information Theory*, 53(11):3944–3957, 2007.
27. Jean-Charles Faugère. A new efficient algorithm for computing Gröbner bases (F4). *Journal of Pure and Applied Algebra*, 139(1-3):61–88, June 1999.
28. Jean-Charles Faugère. A New Efficient Algorithm for Computing Gröbner Bases Without Reduction to Zero (F5). In T. Mora, editor, *ISSAC '02: Proceedings of the 2002 International Symposium on Symbolic and Algebraic Computation*, pages 75–83, New York, NY, USA, July 2002. ACM Press. isbn: 1-58113-484-3.
29. Jean-Charles Faugère and Antoine Joux. Algebraic cryptanalysis of Hidden Field Equation (HFE) cryptosystems using Gröbner bases. In Dan Boneh, editor, *Advances in Cryptology - CRYPTO 2003*, volume 2729 of *LNCS*, pages 44–60. Springer, 2003.
30. Jean-Charles Faugère and Ludovic Perret. An efficient algorithm for decomposing multivariate polynomials and its applications to cryptography. *J. Symb. Comput.*, 44(12):1676–1689, 2009.
31. Jean-Charles Faugère, Joachim von zur Gathen, and Ludovic Perret. Decomposition of generic multivariate polynomials. In *ISSAC'10*, pages 131–137. ACM, 2010.

32. Yun-Ju Huang, Feng-Hao Liu, and Bo-Yin Yang. Public-key cryptography from new multivariate quadratic assumptions. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, *Public Key Cryptography*, volume 7293 of *Lecture Notes in Computer Science*, pages 190–205. Springer, 2012.
33. Takatori Isobe. A single-key attack on the full GOST block cipher. *J. Cryptology*, 26(1):172–189, 2013.
34. Aviad Kipnis, Jacques Patarin, and Louis Goubin. Unbalanced Oil and Vinegar Signature Schemes. In *EUROCRYPT*, pages 206–222, 1999.
35. Aviad Kipnis and Adi Shamir. Cryptanalysis of the HFE public key cryptosystem by relinearization. In *CRYPTO'99*, volume 1666 of *Lecture Notes in Computer Science*, pages 19–30. Springer, 1999.
36. Yann Laigle-Chapuy. A note on a class of quadratic permutations over f_{2^n} . In *AAECC*, volume 4851 of *Lecture Notes in Computer Science*, pages 130–137. Springer, 2007.
37. Eric Leveil and Pierre-Alain Fouque. An improved lpn algorithm. In *SCN'06*, volume 4116 of *Lecture Notes in Computer Science*, pages 348–359. Springer, 2006.
38. Mitsuru Matsui. Linear cryptanalysis method for DES cipher. In *EUROCRYPT'93*, volume 765 of *Lecture Notes in Computer Science*, pages 386–397. Springer, 1993.
39. Tsutomu Matsumoto and Hideki Imai. Public quadratic polynomial-tuples for efficient signature-verification and message-encryption. In *EUROCRYPT'88*, volume 330 of *Lecture Notes in Computer Science*, pages 419–453. Springer, 1988.
40. Willi Meier. Personal communication, 2014.
41. Wil Michiels and Paul Gorissen. Mechanism for software tamper resistance: an application of white-box cryptography. In *Digital Rights Management Workshop*, pages 82–89. ACM, 2007.
42. Jacques Patarin. Cryptanalysis of the matsumoto and imai public key scheme of eurocrypt'88. In *CRYPTO'95*, volume 963 of *Lecture Notes in Computer Science*, pages 248–261. Springer, 1995.
43. Jacques Patarin. Hidden fields equations (HFE) and isomorphisms of polynomials (IP): Two new families of asymmetric algorithms. In *EUROCRYPT'96*, volume 1070 of *Lecture Notes in Computer Science*, pages 33–48. Springer, 1996.
44. Colin Percival. Stronger key derivation via sequential memory-hard functions. *Self-published*, 2009.
45. Vincent Rijmen and Bart Preneel. A family of trapdoor ciphers. In *FSE'97*, volume 1267 of *Lecture Notes in Computer Science*, pages 139–148. Springer, 1997.
46. Koichi Sakumoto, Taizo Shirai, and Harunaga Hiwatari. Public-key identification schemes based on multivariate quadratic polynomials. In *CRYPTO*, volume 6841 of *Lecture Notes in Computer Science*, pages 706–723. Springer, 2011.
47. Amitabh Saxena, Brecht Wyseur, and Bart Preneel. Towards security notions for white-box cryptography. In *ISC'09*, volume 5735 of *Lecture Notes in Computer Science*, pages 49–58. Springer, 2009.
48. Chengdong Tao, Adama Diene, Shaohua Tang, and Jintai Ding. Simple matrix scheme for encryption. In *PQCrypto'13*, volume 7932 of *Lecture Notes in Computer Science*, pages 231–242. Springer, 2013.
49. Brecht Wyseur. *White-Box Cryptography*. PhD thesis, Katholieke Universiteit Leuven, Leuven, Belgium, March 2009.
50. Xiangyong Zeng, Xishun Zhu, and Lei Hu. Two new permutation polynomials with the form $(x^{2^k} + x + d)^s + x$ over \mathbb{F}_2^m . *Appl. Algebra Eng. Commun. Comput.*, 21(2):145–150, 2010.