

On the Efficiency and Flexibility of Signature Verification

Cecilia Boschini¹, Dario Fiore², and Elena Pagnin³

¹ Università della Svizzera Italiana, Lugano, Switzerland
cecilia.boschini@usi.ch

² IMDEA Software Institute, Madrid, Spain
dario.fiore@imdea.org

³ Lund University, Lund, Sweden
elena.pagnin@eit.lth.se

Abstract. For decades signature verification has been regarded as a unique, monolithic process. Here, we want to look at it with fresh eyes and pose two fundamental questions: (1) *is it possible to extract meaningful information from a partial signature verification?* (flexibility); and (2) *is it possible to speed up the verification process without impacting unforgeability?* (efficiency). We answer both questions in a positive way for specific classes of post-quantum secure schemes.

In detail, we develop formal frameworks for signatures with efficient verification, flexible verification and combinations of the two. Crucially, we regard these as features that may enhance existing constructions. Flexibility is of particular interest as standard verification cannot provide any meaningful information about the validity of a given signature if interrupted *in media res*. We exhibit generic transformations to realize efficient (and) flexible verification for schemes that involve matrix-vector multiplications among the verification checks.

In addition, we present concrete instantiations of efficient (and) flexible verification for Rainbow [ACNS05] (as representative of schemes based on multivariate quadratic equations), MP [EC12] and GVW [STOC15] (as representative of lattice-based constructions). Interestingly, we are able to *efficiently verify* Rainbow signatures using 50% of the original computational cost, and as little as 0.4% for GVW homomorphic signatures, provided a one-time preprocessing and with only negligible impact on security.

Keywords: Digital signatures, amortized efficiency, flexible verification, post quantum signatures.

1 Introduction

Digital technologies gained a fundamental role in our society, affecting the security of crucial systems such as autonomous vehicles, healthcare, payments, e-voting, and access control. Digital signatures are among the cryptographic primitives employed to safeguard such systems from misuse [11,13,26,33]. Concretely, digital signatures allow one party, the signer, to use her secret key to authenticate a message in such a way that anyone holding the corresponding public key, the verifiers, can check its validity at any later point in time. Among many properties, digital signatures serve the purpose of securely establishing the source and integrity of the information. In many applications, the outcome of a signature verification determines what decisions to take, e.g., whether to accept a financial transaction (Bitcoin protocol), install software updates (Android OS), or deliver e-services (e-Health, electronic tax systems).

While the above examples are not time-critical, digital signatures may be employed in real-time cyber-physical systems as well, where the speed at which verification is performed plays a crucial role. This is the case, e.g., for automatic safety measures in connected vehicles or nuclear power plants. In real-time systems, verification speed (due to resource constraints) is not the only problem though. For a variety of reasons, a computation can get arbitrarily interrupted, leaving the verifier with an unsolved answer about the validity of a given signature. This has to do with

the fact that standard verification procedures provide a binary outcome (0 or 1, reject or accept), which is established only at the very end of the execution. One may wonder: *is it possible to extract meaningful information from a partial signature verification?* Le et al. [25] proposed to address unexpected interruptions using signatures with flexible verification. In a nutshell, such schemes admit a verification algorithm that increasingly builds confidence on the validity of the signature while it performs more steps. In this way, at the moment of an interrupt, the verifier is left with a value $\alpha \in [0, 1] \cup \perp$ that probabilistically quantifies the validity of the signature. In particular, flexible signatures identify tradeoffs between the amount of computation performed and the integrity of a signature. A different way to approach this problem would be to limit the verification to few, quick steps, as to reduce the chance of interruptions in media res. One may wonder: *is it possible to speed up the verification process without impacting unforgeability?* This is the aim of efficient verification: to leverage special tradeoffs to quickly reach accurate conclusions about the validity of a signature.

1.1 Our contribution

In this work, we address two main challenges: (1) speeding up the verification of digital signatures; and (2) extracting meaningful information from a partial signature verification. We do so by designing *new* verification methods for *existing* signature schemes. Namely, our goal is to keep the signatures (i.e., key generation, sign and verify algorithms) as they are, and to devise techniques for alternative verification methods that are significantly more efficient and/or that retain usefulness even when interrupted *in media res*.

We introduce formal security models for the notions of efficient verification and flexible verification. For both settings, we provide generic compilers and show how these directly apply to a wide range existing signature schemes. Moreover, our compilers can be combined to simultaneously achieve flexible and efficient verification for which we present a formal model, generic realizations, security arguments and implications. Finally, we remark that our models can easily be extended to include signatures with advanced properties including: ring, threshold, homomorphic, attribute-based and constrained.

We focus our realizations on two families of digital signatures: lattice-based and multivariate-polynomials-based. Both are interesting examples thanks to their plausibly post-quantum security. The multivariate-polynomials-based scheme we consider is Rainbow [15,14], one of the NIST candidates for standardization, and LUOV [5]. Our main representatives for lattice-based constructions are GPV [20] (hash & sign), MP [28] (Boyen/ BonsaiTree) and GVW [21] (homomorphic).

It is worth noting that finding faster verification for (families of) existing and very well-studied signature schemes may be a too ambitious goal (e.g., it may require major algorithmic breakthroughs such as finding faster matrix multiplication algorithms). In our paper, we therefore address this problem by introducing new verification models based on a probabilistic secret-coins preprocessing.

Efficient Verification. We build on the well-known offline/online paradigm to define an alternative verification that consists of two phases. An offline probabilistic preprocessing `offVer` that, given only the public key and a security parameter, outputs a concise (short) verification key `svk`. An online algorithm `onVer` that uses such verification key to establish the validity of a signature on a given message. In our paradigm, (i) the preprocessing is one-time, namely its output `svk` can be reused to verify an unbounded number of signatures, (ii) the running time of `onVer` is required to be asymptotically smaller than that of the standard verification algorithm, and (iii) the preprocessed

verification key svk must be kept secret by the verifier. While the first two properties (i)-(ii) are clearly the ones that provide the desired efficiency boost, one may wonder if property (iii) is really needed. To this end, we observe that devising algorithms $\text{offVer}, \text{onVer}$ with efficient verification (i)-(ii) and a *public svk* for an existing signature scheme would immediately lead to finding a new version of the signature scheme with faster verification. We therefore consider keeping svk secret and trade some usability for more efficiency. We formalize this secrecy requirement in a rigorous security model in which the standard unforgeability experiment is extended by asking the adversary to produce a forgery (μ, σ) that verifies under $\text{onVer}(\text{svk}, \cdot, \cdot)$ and by giving it oracle access to the online verification with the same preprocessed svk . This experiment essentially models that learning outcomes of verification do not leak enough information about svk and do not help the adversary in finding a forgery. We stress that *in our model signatures are still publicly verifiable*; svk is only a secret that the verifier generates for itself, in order to speedup its computation.

In terms of realizations, we propose efficient verification algorithms that work for a broad class of digital signatures in which the verification includes a matrix-vector multiplication. This can model the $\mathbf{A} \cdot \sigma = \mathbf{u}$ check of several lattice-based signatures (LBS) where matrix \mathbf{A} may be a public-key constant and \mathbf{u} be message-dependent, as in [16,20,21], or \mathbf{A} may be message-dependent and \mathbf{u} be constant [8,9,28]. Or it can model the check of a system of multivariate quadratic polynomials $\{f_i(\mathbf{s}) = h_i\}_i$, where \mathbf{s} depends on the signature, h_i is message-dependent and $f_i(\cdot)$ is a public-key-dependent polynomial. For example, for the LBS instantiations we consider, assuming a lattice of dimension $n = 128$, our method offers an online verification that is between $97\times$ to $99\times$ faster.

Flexible Verification. We consider flexible verification as an *add-on* property to existing signature schemes. Concretely, we devise flexible verification via an algorithm flexVer and a confidence function α_{flex} with the following properties. First, flexVer is stateful, uses private-coins, and consists of at most $N + 1$ inner steps, so that at the end of step i one is either sure about rejection or expects the signature to be valid with confidence $\alpha_{\text{flex}}(i)$. Second, $\alpha_{\text{flex}} : \{0, \dots, N\} \rightarrow [0, 1]$ is a non-decreasing function (that depends both on the signature scheme and on flexVer) satisfying $\alpha_{\text{flex}}(0) = 0$ and $\alpha_{\text{flex}}(N) = (1 - \varepsilon(\lambda))$. In particular, reaching the last flexible verification step provides the same security guarantees as standard verification (except for a negligible chance of error). Third, we model secure flexible verification by letting the adversary decide at which step to interrupt each execution of flexVer , even in the forgery check. Since in such a case a wrong signature may be marked as ‘maybe valid’ with non negligible probability (as the adversary may choose to verify it with not enough steps to reach overwhelming confidence), we require the adversary to find a forgery that is not rejected with probability non-negligibly *higher* than the expected confidence level at the chosen interruption step i , i.e., $\alpha_{\text{flex}}(i)$.

In terms of realizations, we again focus on the aforementioned class of signatures with a matrix-vector multiplication and we propose flexible verification algorithms for them. In a nutshell, for signatures that work over \mathbb{Z}_q or \mathbb{F}_q , our flexible verification interrupted at step i can achieve confidence $(1 - q^{-i})$. In the case of signature schemes where $q = 2^{\text{poly}(\lambda)}$, this confidence level is always overwhelming, even with minimal computation, i.e., $i = 1$. Interestingly, in this case our flexVer algorithm also satisfies efficient verification, namely, there is a minimum number N so that N executions of flexVer are faster than N executions of Ver . For signature schemes where $q = \text{poly}(\lambda)$, we show a variant of our compiler that achieves both efficient and flexible verification. This though holds in a weaker security model, in which the confidence function α_{flex} depends, and degrades, with the number of verifications.

1.2 An Overview of Our Techniques

Our Compiler for Efficient Verification. We consider the class of digital signature schemes for which the bulk of computation in the verification procedure consists of a matrix-vector product (what we call ‘ $\mathbf{M}\mathbf{v}$ ’-style check). Several lattice-based [6,21,16,20,28] and multivariate-polynomials-based signatures [15,5,29] fall in this category. For simplicity, here we give a technical overview for LBS constructions only, for the case $\mathbf{M}\mathbf{v} = \mathbf{u}$, where $\mathbf{M} \in \mathbb{Z}_q^{n \times m}$ is fixed (e.g., the signer’s public key), \mathbf{v} is a vector (e.g., a signature) and \mathbf{u} depends on the message [16,20,21]. The key observation to efficient verification is that if a signature \mathbf{v} verifies for a given \mathbf{M} (and \mathbf{u}), then it also verifies for any linear combination of the rows of \mathbf{M} (against the same combination to the entries of \mathbf{u}). Let $\mathbf{c} \xleftarrow{\$} \mathbb{Z}_q^n$ be a random vector; denote $\mathbf{z} \leftarrow \mathbf{c} \cdot \mathbf{M} \in \mathbb{Z}_q^m$, and $w \leftarrow \mathbf{c} \cdot \mathbf{u} = \sum_{i=1}^n \mathbf{c}[i]\mathbf{u}[i] \in \mathbb{Z}_q$. Then

$$\mathbf{M}\mathbf{v} = \mathbf{u} \pmod q \implies \mathbf{c} \cdot (\mathbf{M}\boldsymbol{\sigma}) = \mathbf{c} \cdot \mathbf{u} \pmod q \Leftrightarrow \mathbf{z} \cdot \mathbf{v} = w \pmod q \quad (1)$$

Our compiler is based on the fact that the first implication in (1) holds left-to-rightwards (always) *and also* right-to-leftwards with *good enough* probability. Moreover, one can adjust this probability by increasing the number of vectors \mathbf{c} and \mathbf{z} to use in the signature verification. Now, we can precompute a set of $k \geq 1$ ‘random combinations of rows of \mathbf{M} ’ during an ‘offline verification’ phase; and re-use the output pairs $(\mathbf{c}_j, \mathbf{z}_j)$, $j \in \{1, \dots, k\}$ to efficiently verify an unbounded number of signatures. Concretely, the efficiency gain in the online verification comes by replacing the check $\mathbf{M}\mathbf{v}$ (that requires nm multiplications modulus q) with k checks of the form $\mathbf{z}_j \cdot \mathbf{v} = w_j$ (that require at most km multiplications). Remarkably, the leftwards implication in (1) holds with all but negligible probability with very small values of k , and even with $k = 1$ for schemes where $q = 2^{\text{poly}(\lambda)}$, e.g., [16,21].

Our Compiler for Flexible Verification. The idea of the compiler presented above can be used to achieve flexible verification in the following way. For a set of freshly sampled vectors $\{\mathbf{c}_j\}_{j=1}^J$, a verification that uses a subset of $k < J$ such vectors is correct with probability $(1 - q^{-k})$; hence this value can be the confidence, i.e., $\alpha_{\text{flex}}(k) = (1 - q^{-k})$.

1.3 Related Work

The problem of trading security for less computation during a signature verification has been considered by Le et al. [25] who introduced the notion of flexible signatures and a construction based on the Lamport-Diffie one-time signature [24] with Merkle trees [27]. A similar idea in the context of message authentication codes (MACs) has been considered by Fischlin [17] who put forth progressive verification for MACs and presented two concrete constructions. One main difference between our model and those of [17,25] is that we aim to capture flexible verification as an independent feature that can enhance existing schemes, rather than a standalone primitive. This is in a way more challenging as it leaves less design freedom when crafting these algorithms. Therefore we decided to define flexible verification as a *stateful* algorithm in contrast to stateless [17,25]: although this makes our model slightly more involved, it is comparably more general and can capture more (existing) schemes.

Our model for efficient verification is close the offline-online paradigm used in homomorphic authentication [2,10] and verifiable computation [19]; where a preprocessing is done with respect to a function f , and its result can be used to verify computation results involving the same f .

A recent work by Sipasseuth et al. [31] investigates how to speed up lattice-based signature verification while reducing the memory (storage) requirements. The overall idea in [31] is similar

to ours (and inspired to Freivalds’ Algorithm): to replace the inefficient matrix multiplication in the verification with a probabilistic check via an inner product computation. However, [31] focuses on a concrete construction, the DRS signature [30], and investigates the trade-off between pre-computation time for verification and memory storage for this scheme only. Moreover, the work lacks a formal, abstract analysis of the security impact of such a shift in the verification procedure. In contrast, we devise a general framework to model ‘more efficient’ and ‘partial’ signature verification. Albeit we developed our approach independently of [31], our techniques can be seen as a generalization of what presented in [31].

Independently from us, Taleb and Vergnaud recently published a paper about speeding up signature verification [32]. The paper analyzes three kinds of signatures, the RSA and ECDSA signatures and the lattice-based signature from [20]. While the first two are complementary to our work, the latter construction is a different take on achieving efficient verification for this class of LBS (they do not have a comparable approach for the flexible case). Their approach exploits a particular type of error correcting codes that do not have words with small Hamming weight. Given a generator matrix \mathbf{G} for the code, their verification algorithm checks that the codeword $\mathbf{G}(\mathbf{A}\boldsymbol{\sigma} - \mathbf{u}) \bmod q$ has null components. The sparsity property of the code ensures that if $\mathbf{A}\boldsymbol{\sigma} - \mathbf{u} \neq 0 \bmod q$ (i.e., $\boldsymbol{\sigma}$ is not a valid signature), the corresponding codeword has enough nonzero components that a random choice of them would contain a nonzero one with overwhelming probability. The drawback is that such construction is not a compiler, as ours, but requires to modify both the key generation and the verification algorithm, and requires the public key to be longer, thus slowing down the full verification. Moreover, their progressive verification still requires a number of vector multiplications linear in the security parameter, while ours only requires a logarithmic number of linear products.

Remark on terms. We use the term ‘flexible signature’ to be consistent with previous work [25], however, this wording may be misleading in that there is no flexibility in the signature scheme, but rather the verification procedure is carried out in a progressive way (à la Fischlin [17]) in the sense that the more computation one performs the more confident one becomes of accepting or rejecting a signature (gradually approaching 100% certainty). We believe that an algorithm should be able to read (and understand) its input, and should not have an accessory input that cannot use (the $[[k]]$ of [25]).

2 Preliminaries

We denote the set of real values by \mathbb{R} , integers by \mathbb{Z} , natural numbers by $\mathbb{Z}_{\geq 0}$, and finite fields of integers by \mathbb{Z}_q , where q is a (power of a) prime number. We denote vectors by bold, lower-case letters, and matrices by bold, upper-case letters. We use $\mathbf{v}[i]$ to identify the i -th entry of a vector \mathbf{v} , and $\mathbf{A}[i, j]$ to identify the entry in the i -th row and j -th column of a matrix \mathbf{A} . The norm of a vector is denoted as $\|\mathbf{v}\|$ and unless otherwise specified, it is assumed to be the infinity norm, i.e., $\|\mathbf{v}\| = \max_i\{\mathbf{v}[i]\}$. \mathbf{A}^T denotes the transposed of a matrix. We use $rows(\mathbf{A})$, $cols(\mathbf{A})$, and $rk(\mathbf{A})$ to respectively refer to the number of rows, the number of columns, and the rank of a matrix \mathbf{A} ; $\mathbf{1}_{1 \times n}$ (resp. $\mathbf{0}_{1 \times n}$) denotes the row vector of length n that has all entries equal to 1 (resp. 0); while \mathbf{I}_n denotes the n by n identity matrix of dimension n . We omit the explicit dimensions when they are clear from the context. We denote by $L_1|L_2$ the result of appending a list of elements L_2 to L_1 . Given two values $a < b$, we denote a continuous interval as $[a, b] \subseteq \mathbb{R}$, and a discrete interval as $\{a, \dots, b\} \subseteq \mathbb{Z}$. A function $\varepsilon : \mathbb{Z}_{\geq 0} \rightarrow [0, 1]$ is negligible if $\varepsilon(\lambda) < 1/\text{poly}(\lambda)$ for every univariate polynomial $\text{poly} \in \mathbb{R}[X]$ and a large enough integer $\lambda \in \mathbb{Z}_{\geq 0}$. Throughout the paper,

$\lambda \in \mathbb{Z}_{\geq 0}$ denotes the security parameter of a cryptographic scheme. The abbreviation PPT refers to algorithms that are probabilistic and run in polynomial time.

Definition 1 (Signature Scheme). A signature scheme Σ a tuple of PPT algorithms $\Sigma = (\text{KeyGen}, \text{Sign}, \text{Ver})$ defined as follows.

$\text{KeyGen}(1^\lambda)$: The key generation algorithm takes in input the security parameter and outputs a pair of keys (sk, pk) .

$\text{Sign}(\text{sk}, \mu)$: The sign algorithm takes in input a secret key sk and a message μ ; it outputs a signature σ .

$\text{Ver}(\text{pk}, \mu, \sigma)$: The verification algorithm takes in input a public key, a message μ and signature σ . It outputs 0 (reject) or 1 (accept).

Correctness For a given security parameter λ , for any key pair $(\text{sk}, \text{pk}) \leftarrow \text{KeyGen}(\lambda)$, for any $\mu \in \mathcal{M}$, for any $\sigma \leftarrow \text{Sign}(\text{sk}, \mu)$, it holds that

$$\text{Prob}[\text{Ver}(\text{pk}, \mu, \sigma) = 1] = 1.$$

We use the expression ‘valid (resp. invalid) signature’ to identify signatures that are accepted (resp. rejected) by the verification procedure.

Security The basic security requirement for a digital signature scheme is unforgeability. In a nutshell this notion states that an adversary should not be able to produce a valid signature without knowledge of the secret signing key.

2.1 Lattices

Let $n, m \in \mathbb{Z}_{>0}$, and $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ be a uniformly chosen matrix. We define the lattice that has \mathbf{A} as parity check matrix as

$$\Lambda_q^\perp(\mathbf{A}) = \{\mathbf{z} \in \mathbb{Z}^m : \mathbf{Az} = 0 \pmod{q}\}.$$

One of the most used lattice-based hardness assumptions is the Short Integer Solution (SIS), defined in the following.

Definition 2 (Short Integer Solution problem). Given a matrix $\mathbf{A} \leftarrow_{\mathbb{S}} \mathbb{Z}_q^{n \times m}$, the $\text{SIS}_{n,q,m,\beta}$ problem requires to find a vector $\mathbf{s} \in \mathbb{Z}^m$ such that $\mathbf{As} = 0 \pmod{q}$ and $\|\mathbf{s}\| \leq \beta$.

The inhomogeneous version of the SIS problem (ISIS) requires, given a matrix $\mathbf{A} \leftarrow_{\mathbb{S}} \mathbb{Z}_q^{n \times m}$ and a vector $\mathbf{u} \in \mathbb{Z}_q^n$ to find a vector $\mathbf{s} \in \mathbb{Z}^m$ such that $\mathbf{As} = \mathbf{u} \pmod{q}$ and $\|\mathbf{s}\| \leq \beta$.

3 Efficient Verification for Digital Signatures

In this section, we introduce the concept of *efficient verification* for digital signatures and a suitable formal security model that allows to estimate the impact efficient verification has on the unforgeability of the scheme. Then, we describe a generic compiler to obtain efficient verification for a wide class of signatures, prove its security, show realizations from lattices and multivariate polynomials and finally discuss its concrete efficiency against full-fledged verification. In addition we discuss concrete instantiations and their corresponding concrete efficiency gains. This section ends with a generalization of our results to signatures with properties.

Efficient Verification In a nutshell, the core idea of efficient signature verification is to split the verification process into two steps. The first step is a one-time and signature-independent setup called ‘offline verification’. Its purpose is to produce randomness to derive a (short, secret) verification key svk from the signer’s public key pk . Note that the offline verification does not change the signature, which remains publicly verifiable; instead it ‘randomizes’ pk to obtain a concise verification key svk that essentially enables one to verify signatures with (almost) the same precision as the standard verification, but in a more efficient way. The second verification step consists of an ‘online verification’ procedure. It takes as input svk and can verify an unbounded number of message-signature pairs performing significantly less computation than the standard verification algorithm.

Definition 3 (Efficient Verification). *A signature scheme $\Sigma = (\text{KeyGen}, \text{Sign}, \text{Ver})$ admits efficient verification if there exist two PPT algorithms $(\text{offVer}, \text{onVer})$ with the following syntax:*

$\text{offVer}(\text{pk}, k)$: *on input a verification key pk , a positive integer $k \in \{1, \dots, \lambda\}$ (referred to as confidence level), the offline verification algorithm returns a secret verification key svk .*

$\text{onVer}(\text{svk}, \mu, \sigma)$: *on input a secret verification key svk , a message μ , and a signature σ , the efficient online verification algorithm outputs 0 (reject) or 1 (accept).*

The standard properties of an efficient verification scheme are described below.

Correctness. For a given security parameter λ , for any honestly generated key pair $(\text{sk}, \text{pk}) \leftarrow \text{KeyGen}(\lambda)$, for any message $\mu \in \mathcal{M}$, for any signature σ such that $\text{Ver}(\text{pk}, \mu, \sigma) = 1$, and for any confidence level $k \in \{1, \dots, \lambda\}$; the following condition holds:

$$\Pr[\text{onVer}(\text{svk}, \mu, \sigma) = 1 \mid \text{svk} \leftarrow \text{offVer}(\text{pk}, k)] = 1.$$

This guarantees that a valid signature σ is always accepted by online verification. We remark that any signature rejected by onVer , would be rejected by Ver as well.

Amortized Efficiency. Informally, the notion of amortized efficiency captures the fact that the global computational cost of doing the preprocessing (running offVer once) and running onVer a given number times is considerably smaller than the cost of running the standard signature verification Ver the same number of times. More formally, let λ be the security parameter and $\text{cost}(\cdot)$ be a function that, given in input an algorithm returns its computational cost (in some desired computational model). The pair $(\text{offVer}, \text{onVer})$ satisfies amortized efficiency if for any key pair $(\text{sk}, \text{pk}) \leftarrow \text{KeyGen}(\lambda)$, for any message-signature pair (μ, σ) , for any confidence level $k \in \{1, \dots, \lambda\}$ and $\text{svk} \leftarrow \text{offVer}(\text{pk}, k)$, there exists a ‘small’ positive real constant $\epsilon \in]0, 1[= [0, 1] \setminus \{0, 1\}$ such that

$$\frac{\text{cost}(\text{onVer}(\text{svk}, \mu, \sigma))}{\text{cost}(\text{Ver}(\text{pk}, \mu, \sigma))} < \epsilon.$$

Note that in this definition we do not include the cost needed to compute svk (i.e., $\text{cost}(\text{offVer})$). This is justified by the fact that, in our framework, offVer is a polynomial-time *one-time* set up, and its output (svk) can be reused in an unbounded number of executions of onVer . Namely, the computational cost of offVer is amortized over sufficiently many online verifications.

The concrete formula to quantify the amortized efficiency gain compares the computational cost of verifying r signatures with the offline/online approach over running r standard verifications.

Definition 4 (Concrete Amortized Efficiency). We say that a pair $(\text{offVer}, \text{onVer})$ satisfies (r_0, e_0) -concrete amortized efficiency if r_0 is the smallest, non-negative integer such that for any $r \geq r_0$ there exists a small, real constant $0 < e_0 < 1$ for which the following holds true:

$$\frac{\text{cost}(\text{offVer}(\text{pk}, k)) + r \cdot \text{cost}(\text{onVer}(\text{svk}, \mu, \sigma))}{r \cdot \text{cost}(\text{Ver}(\text{pk}, \mu, \sigma))} < e_0. \quad (2)$$

Security Our definition of efficient verification lets the verifier set the confidence level $k \in \{1, \dots, \lambda\}$ at which she wishes to carry out the signature verification. Notably k also determines the amount of computation to be performed by offVer and onVer and thus, the efficiency gain. Our security notion aims at measuring the potential security loss that comes with performing less checks than the full-fledged verification. Intuitively, we say that an efficient verification is secure if the probability that onVer accepts a signature that would be rejected by Ver is negligible. We cannot expect a more efficient verification to detect more forgeries than the full verification. Since the output of offVer is a possibly secret verification key, in the security game we allow the adversary to interact polynomially many times (in the security parameter λ) with both the sign oracle OSign and an additional efficient verification oracle OonVer . As one would expect, the adversary can query the oracles in an adaptive and parallel way. This is formalized by the following definition.

$\text{cmvEUF}(\lambda, \Sigma, k)$	$\text{Exp}_{\mathcal{A}, \Sigma}^{\text{cmvEUF}}(\lambda, k)$	$\text{OSign}_{\text{sk}}(\mu)$
1: $L_S \leftarrow \emptyset$	1: $(\mu^*, \sigma^*) \leftarrow \text{cmvEUF}(\lambda, \Sigma, k)$	1: $L_S \leftarrow L_S \cup \{\mu\}$
2: $(\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(1^\lambda)$	2: if $\mu^* \in L_S$	2: $\sigma \leftarrow \text{Sign}(\text{sk}, \mu)$
3: $\text{svk} \leftarrow \text{offVer}(\text{pk}, k)$	3: return 0	3: return σ
4: $(\mu^*, \sigma^*) \leftarrow \mathcal{A}^{\text{OSign}, \text{OonVer}}(\text{pk}, k)$	4: if $\text{Ver}(\text{pk}, \mu^*, \sigma^*) = 1$	<hr/> $\text{OonVer}_{\text{svk}}(\mu, \sigma)$ <hr/>
5: return (μ^*, σ^*)	5: return 0	1: $b \leftarrow \text{onVer}(\text{svk}, \mu, \sigma)$
	6: $b \leftarrow \text{onVer}(\text{svk}, \mu^*, \sigma^*)$	2: return b
	7: return b	

Fig. 1: Security model for efficient verification of signatures: unforgeability under adaptive chosen message and verification attack (security game, experiment and oracles).

Definition 5 (Security of Efficient Verification). Let Σ be a signature scheme that admits efficient verification. For a given security parameter λ and for any confidence level $k \in \{1, \dots, \lambda\}$, the pair of algorithms $(\text{offVer}, \text{onVer})$ realizes a secure efficient verification for Σ if it is existentially unforgeable under adaptive chosen message and verification attack. In other words, if for all PPT adversaries \mathcal{A} the success probability in the cmvEUF experiment in Figure 1 is negligible, i.e.:

$$\text{Adv}_{\mathcal{A}, \Sigma}^{\text{cmvEUF}}(\lambda, k) = \Pr \left[\text{Exp}_{\mathcal{A}, \Sigma}^{\text{cmvEUF}}(\lambda, k) = 1 \right] \leq \varepsilon = \varepsilon(\lambda, k).$$

Line 5 of the cmvEUF experiment excludes forgeries against the original signature scheme. This is justified by the correctness of efficient verification and by the fact that Σ is supposed to be an unforgeable signature scheme in the usual sense. In other words, the above security experiment checks whether the adversary outputs invalid signatures that are however accepted by the efficient verification mechanism.

Notably, the security definition depends on the confidence level k , and we do not consider a solution to be secure if k is such that $\varepsilon(\lambda, k)$ is actually a non-negligible value. Yet, the latter scenario

may occur in cases where, e.g., the execution of `onVer` is *prematurely interrupted*. A meaningful definition of security in such scenarios requires a new model as we show in Section 4.

3.1 A Generic Compiler for Mv-style Verifications

Until now we presented a framework for efficient verification of digital signatures. In what follows we show how to apply the efficient verification paradigm to a wide class of existing schemes. To this end, we begin with abstracting the verification procedure of a signature scheme into two types of verification checks: a matrix-vector multiplication (referred to as $\mathbf{M}\mathbf{v} = 0$, for appropriate matrix \mathbf{M} and vector \mathbf{v}) and other generic checks (collected in the `Check` subroutine). Notably, our compiler leads to efficient verification whenever the computational complexity of the verification procedure is dominated by the matrix-vector multiplication, i.e. $\text{cost}(\text{Check}) \ll \text{cost}(\mathbf{M}\mathbf{v}) \sim mn$ field multiplications (for $\mathbf{M} \in \mathbb{Z}_q^{n \times m}$).

We present a generic way to achieve efficient verification (as of Definition 3) for any signature scheme in which the verification procedure includes a matrix-vector ($\mathbf{M}\mathbf{v}$) multiplication. For simplicity, we call these **Mv-style** check verifications (see Figure 2 for details). Schemes that fall in this category include some seminal lattice-based signatures [8,20,9,28], homomorphic signatures [6,21,16], and multivariate signatures [15,5].

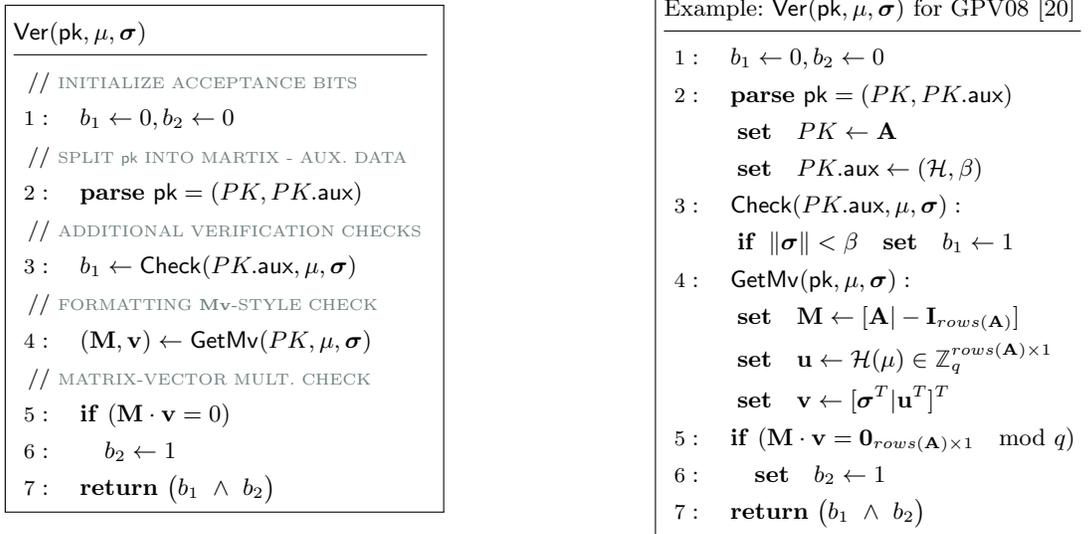


Fig. 2: General structure of a verification with an **Mv-style** check (on the left); an instructive example: the GPV08 signature verification (on the right) [20].

Our generic compiler for efficient verification is detailed in Figure 3 with a sketch of instantiation for the LBS scheme GPV08 [20] as a running example (more details on all the proposed instantiations can be found in Section 3.2). The key observation is that for any pair of vectors σ and \mathbf{u} , and for any matrix \mathbf{A} (of opportune dimensions) if $\mathbf{A} \cdot \sigma = \mathbf{u}$ then for any random vector \mathbf{c} (of opportune dimension) it holds that $\mathbf{c} \cdot (\mathbf{A} \cdot \sigma) = \mathbf{c} \cdot \mathbf{u}$. By collecting variables on the left hand side we get: $\mathbf{c} \cdot [\mathbf{A} | -\mathbf{1}] \cdot (\sigma, \mathbf{u}) = 0$. Thus one can precompute the vector $\mathbf{z} \leftarrow \mathbf{c} \cdot [\mathbf{A} | -\mathbf{1}]$ and perform efficient ‘online’ verification by checking whether $\mathbf{z} \cdot \mathbf{u} = 0$, where $\mathbf{u} \leftarrow (\sigma, \mathbf{u})$ (note that the matrix-vector

multiplication is now replaced by a vector-vector multiplication, aka inner-product). Correctness and efficiency are immediate. Soundness comes from the fact that if $\mathbf{z} \cdot \mathbf{u} = 0$ then with all but negligible probability also the original system of linear equations $\mathbf{A} \cdot \boldsymbol{\sigma} = \mathbf{u}$ is satisfied, as proven in Theorem 1.

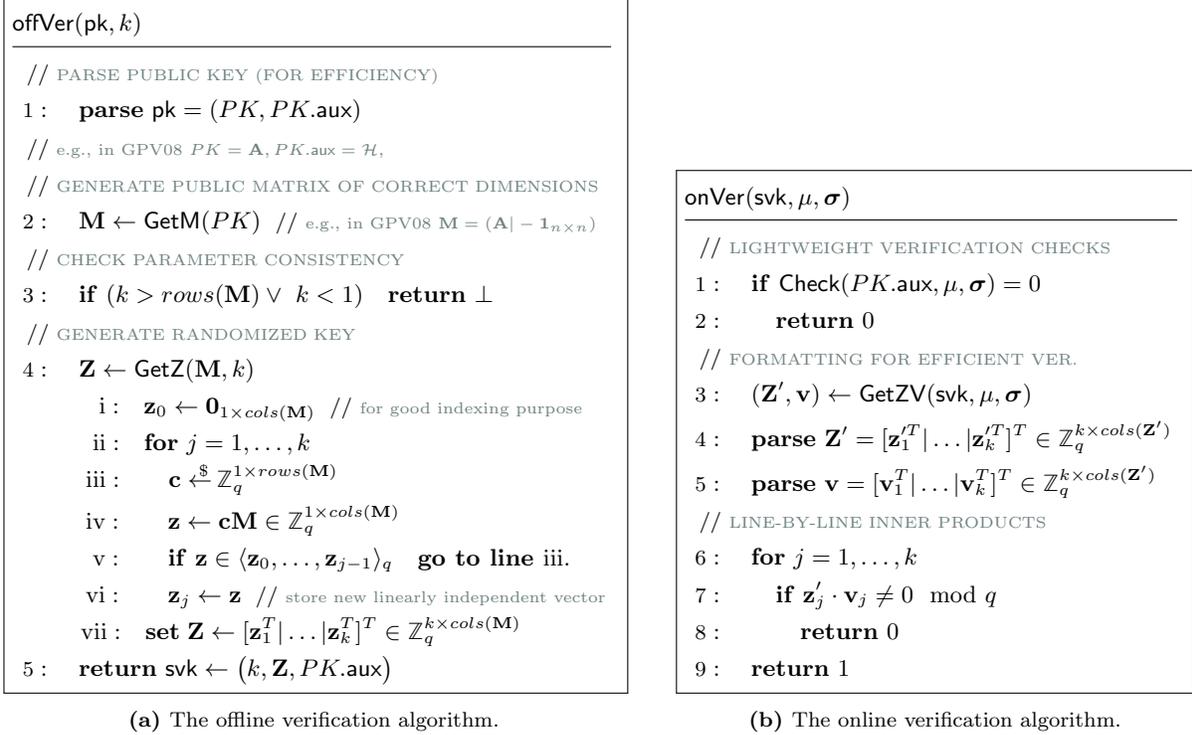


Fig. 3: Our compiler for efficient verification of signatures with **Mv**-style verification. The four scheme-dependent subroutines are: **parse pk** and **GetZ** (in **offVer**); **Check** and **GetZV** (in **onVer**). The complexity of the **onVer** is linear in k , the chosen confidence level.

Security Analysis of Our Compiler. Next, we prove the security of our generic compiler. We discuss the tradeoffs between $q(\lambda)$ and k in the next section.

Theorem 1. *Let Σ be a signature scheme with the structure of Figure 2. Then our compiler for Σ (depicted in Figure 3) generates a secure efficient verification realization of Σ —i.e., (**offVer**, **onVer**) are existentially unforgeable under adaptive chosen message and verification attacks—with $\varepsilon = \frac{q_V + 1}{q^k}$, where $k \in \{1, \dots, rk(\mathbf{M})\}$ denotes the chosen confidence level and $q_V = \text{poly}(\lambda)$ is a bound on the total number of verification queries.*

Proof. Let us note that the winning condition of the experiment requires \mathcal{A} to produce a message-signature pair (μ^*, σ^*) such that μ^* has not been queried to the signing oracle during the game, the signature is invalid under standard verification, i.e., $\text{Ver}(\text{pk}, \mu^*, \sigma^*) = 0$, but it is accepted by the online verification, i.e., $\text{onVer}(\text{svk}, \mu^*, \sigma^*) = 1$. The goal of the proof is to bound the probability that this event occurs.

First, let us define \mathbf{G}_0 to be the cmvEUF experiment of Figure 1. Next, we define a series of q_V hybrid games where, for every $i = 1$ to q_V , \mathbf{G}_i is the same as \mathbf{G}_{i-1} except that the i -th query to onVer is answered by using the standard verification algorithm Ver . Clearly, for every i , the view of the adversary in \mathbf{G}_i is the same as in \mathbf{G}_{i-1} unless, at the i -th verification query (μ_i, σ_i) , this is answered with a reject in \mathbf{G}_i (i.e., $\text{Ver}(\text{pk}, \mu_i, \sigma_i) = 0$) whereas it would have been accepted in \mathbf{G}_{i-1} (i.e., $\text{onVer}(\text{svk}, \mu_i, \sigma_i) = 1$). Let us define this event as

$$\text{bad}_i = \{\text{Ver}(\text{pk}, \mu_i, \sigma_i) = 0 \wedge \text{onVer}(\text{svk}, \mu_i, \sigma_i) = 1\} . \quad (3)$$

Then, it is clear that, $\forall i \in \{1, \dots, q_V\}$,

$$|\Pr[\mathbf{G}_{i-1} = 1] - \Pr[\mathbf{G}_i = 1]| \leq \Pr[\text{bad}_i].$$

Moreover, observe that the adversary wins in game \mathbf{G}_{q_V} if the same event occurs for the message-signature pair (μ^*, σ^*) (that we also denote below with $(\mu_{q_V+1}, \sigma_{q_V+1})$):

$$\text{bad}_{q_V+1} = \{\text{Ver}(\text{pk}, \mu^*, \sigma^*) = 0 \wedge \text{onVer}(\text{svk}, \mu^*, \sigma^*) = 1\}.$$

So, by a union bound it follows that

$$\text{Adv}_{\mathcal{A}, \Sigma}^{\text{CMVA}}(\lambda, k) = \Pr[\mathbf{G}_0 = 1] \leq \sum_{i=1}^{q_V+1} \Pr[\text{bad}_i] \quad (4)$$

In the following claim we show that for every $i \in [q_V + 1]$, $\Pr[\text{bad}_i] \leq 1/q^k$ over the choice of the \mathbf{c}_j 's, which concludes the proof of the theorem.

Claim. $\Pr[\text{bad}_i] \leq 1/q^k$, where q is the modulus (or the size of the field) and k is the chosen confidence level to run the efficient verification.

Let \mathbf{M} and \mathbf{v} be the matrix and vector returned by $\text{GetMv}(\text{pk}, \mu_i, \sigma_i)$ during the standard verification. Since this is a bad forgery the online verification returns 1. This means that (1) the signature σ_i must satisfy all of the ‘consistency checks’ determined by $\text{Check}(\mu_i, \sigma_i)$ (otherwise it would be rejected also by the online verification); and (2) $\mathbf{M}\mathbf{v} \neq 0$ but $(\mathbf{c}_j \mathbf{M})\mathbf{v} = 0$ for all $j \in \{1, \dots, k\}$. We can split $\Pr[\text{bad}_i]$ as:

$$\Pr[\text{bad}_i] = \Pr[\text{Check}(\text{PK.aux}, \mu_i, \sigma_i) = 1] \cdot \Pr[\mathbf{c}_j(\mathbf{M}\mathbf{v}) = 0 \text{ for all } j \in \{1, \dots, k\}].$$

We remark that $\mathbf{z}'_j = \mathbf{c}_j \mathbf{M}$ is (part of) the output of GetZV in the online verification. Assuming the worst case scenario where \mathcal{A} can trivially generate a pair (μ_i, σ_i) that passes the consistency checks in Check , we bound only the second factor in the equation above. We argue that:

$$\Pr[\mathbf{c}_j(\mathbf{M}\mathbf{v}) = 0 \text{ for all } j \in \{1, \dots, k\}] = \frac{1}{q^k}.$$

Let $\mathbf{w} = (\mathbf{M}\mathbf{v}) \neq \vec{0}_{\text{rows}(\mathbf{M}) \times 1}$; thus there exists at least one entry $\ell \in [\text{rows}(\mathbf{M})]$ such that $\mathbf{w}[\ell] \neq 0$. The online verification accepts the adversary’s output if and only if $\mathbf{c}_j \cdot \mathbf{w} = 0$ for all j , i.e., if and only if, for all k (linearly independent constraints) it holds that:

$$\mathbf{c}_j[\ell] = \mathbf{w}[\ell]^{-1} \cdot \sum_{i=1, i \neq \ell}^n \mathbf{c}_j[i] \mathbf{w}[i]$$

At this point we observe that all online verification queries before this i -th one were answered *without* using the \mathbf{c}_j vectors. Therefore, for the sake of bounding $\Pr[\text{bad}_i]$ we can assume that all the vectors \mathbf{c}_j are randomly sampled at this stage, namely they are freshly random, and unknown to \mathcal{A} . So, for every j , the equality above holds with probability $1/q$. Since we have k independent such relations (i.e., the \mathbf{c}_j and the \mathbf{z}'_j are linearly independent), it follows that: $\Pr[\text{bad}_i] \leq \prod_{j=1}^k \frac{1}{q} = \frac{1}{q^k}$. Such probability is negligible for appropriate choices of $q = q(\lambda)$ and $k \in \mathbb{Z}_{>0}$. Concretely, if $q = 2^{\text{poly}(\lambda)}$ as in [21,16] our compiler defines a secure generalized efficient verification already for any $k = 1$. \square

3.2 Concrete Instantiations of Our Compiler

Any instantiation of our compiler is completely determined by the four subroutines **parse pk**, **GetM**, **Check**, and **GetZV**.

From Lattices We present concrete instantiations of our compiler for two categories of LBS: ‘hash & sign’ with representative the GPV08 signature [20]; and ‘Boyer/BonsaiTree’ style with representative MP12 [28].

Efficient Verification for GPV08 [20]. The **parse pk** procedure splits the public key into $PK = \mathbf{A} \in \mathbb{Z}_q^{n \times m}$ (the matrix identifying the signer’s public key), and the auxiliary public information $PK.\text{aux} = \mathcal{H}$, i.e., a description of a full-domain hash function $\mathcal{H} : \{0, 1\}^* \rightarrow \mathbb{Z}_q^n$. The **Check** procedure is exactly as in the original verification (enforcing the norm bound on the signature). The **GetM** algorithm takes in input the public matrix $PK = \mathbf{A}$, and appends to it the identity matrix to obtain $\mathbf{M} = [\mathbf{A} \mid -\mathbf{I}_n]$. The **GetZV** routine returns the matrix \mathbf{Z}' and the vector \mathbf{v} . The matrix \mathbf{Z}' is made up of the same ‘randomized key’ vectors produced by **GetZ** during the offline verification, i.e., $\mathbf{z}'_j = \mathbf{z}_j \leftarrow \mathbf{c}_j \mathbf{M} = [\mathbf{c}_j \mathbf{A} \mid -\mathbf{c}_j]$, i.e., $\mathbf{v} = [\boldsymbol{\sigma} \mid \mathcal{H}(\mu) \cdot \mathbf{1}_{1 \times n}]$. Thus the core verification check (line 7 in **onVer**) is actually ensuring that $\mathbf{z}'_j \mathbf{v}_j = 0$, i.e., $\mathbf{c}_j \cdot \mathbf{A} \cdot \boldsymbol{\sigma} = \mathbf{c}_j \mathcal{H}(\mu)$ which is the probabilistic check of the original verification equality.

Efficient Verification for MP12 [28]. The **parse pk** procedure assigns $PK \leftarrow \mathbf{A} = [\tilde{\mathbf{A}} \mid \mathbf{A}_0 \mid \dots \mid \mathbf{A}_\ell] \in \mathbb{Z}_q^{n \times (\bar{m} + n \lceil \log q \rceil \ell)}$ (the matrix identifying the signer’s public key), where $\bar{m} = O(n \lceil \log q \rceil)$, and ℓ denotes the number of bits in the message, i.e., $\mu \in \{0, 1\}^\ell$. The auxiliary public information is $PK.\text{aux} = \mathbf{u}$. The **Check** procedure is exactly as in the original verification (enforcing the norm bound on the signature). The **GetM** algorithm takes in input the public matrix $PK = \mathbf{A}$, and appends to it the identity matrix to obtain $\mathbf{M} = [\mathbf{A} \mid -\mathbf{I}_{n \times n}]$. The **GetZV** routine returns the matrix \mathbf{Z}' and the vector \mathbf{v} . The matrix \mathbf{Z}' is made up of vectors of the form $\mathbf{z}'_j = [\tilde{\mathbf{z}}_j \mid \mathbf{z}_j^0 + \sum_{i=1}^\ell \mu[i] \mathbf{z}_j^i \mid \mathbf{c}_j]$ that identify a message-dependent lattice (called \mathbf{A}_μ in [28]). The vector \mathbf{v} is the concatenation of the signature with the auxiliary vector, i.e., $\mathbf{v} = [\boldsymbol{\sigma} \mid \mathbf{u}]$. Note that \mathbf{u} is the same for all messages; thus, one could further optimize the online verification by computing (once and for all) the k inner products $\mathbf{z}_j[\bar{m} + n \lceil \log q \rceil + 1] = \mathbf{c}_j \cdot \mathbf{u}$ during the offline phase. To conclude we notice that the online verification ensures that $\mathbf{z}'_j \mathbf{v}_j = 0$, i.e., $\mathbf{c}_j \cdot \mathbf{A}_\mu \cdot \boldsymbol{\sigma} = \mathbf{c}_j \cdot \mathbf{u}$ which is the probabilistic check of the original verification equality.

From Multivariate Equations For signatures schemes based on multivariate equations we take Rainbow [15,14] as representative example as this is one of the NIST candidates for standardization. For completeness, we also show how to apply our compiler to the LUOV scheme [5].

Efficient Verification for Rainbow [14]. In the description below we consider the standard Rainbow verification. A similar approach can be used to speed up the verification also in the “cyclic” and the “compressed” Rainbow variants as in those cases the verification includes an additional initial phase to reconstruct the full public key. We recall that in this scheme the public key contains a system of m multivariate quadratic polynomials in n variables. For convenience, let $N = n(n+1)/2$ and $\mathbb{F} = \mathbb{F}_{2^r}$. Using a Macaulay matrix representation we can visualize this system as a wide matrix composed of a quadratic term \mathbf{Q} (actually a $m \times N$ submatrix), a linear term \mathbf{L} ($m \times n$ submatrix) and a constant term \mathbf{C} (a $m \times 1$ vector). The **parse pk** procedure extracts from the public key PK this matrix $\mathbf{pk} = [\mathbf{Q}|\mathbf{L}|\mathbf{C}] \in \mathbb{F}^{m \times (N+n+1)}$ and a description of a full-domain hash function $\mathcal{H} : \{0, 1\}^* \rightarrow \mathbb{F}^m$ as the auxiliary public information $PK.aux = \mathcal{H}$. The **Check** procedure is trivial and always returns 1. This is because the whole verification can be written as a matrix-vector multiplication. The **GetM** algorithm extracts from PK the matrix representing the system of quadratic multivariate equations $[\mathbf{Q}|\mathbf{L}|\mathbf{C}]$. Finally, it appends to this the identity matrix, so $\mathbf{M} \leftarrow [\mathbf{Q}|\mathbf{L}|\mathbf{C} | \mathbf{I}_m]$. We remark that \mathbf{M} can be seen as a matrix of blocks, where any block has the same height ($m =$ number of rows), but different length (number of columns). The **GetZV** routine reads the matrix $\mathbf{Z}' = \mathbf{Z}$ made up of the rows $\mathbf{z}'_j = \mathbf{z}_j \leftarrow \mathbf{c}_j \mathbf{M} = [\mathbf{c}_j \cdot \mathbf{Q} | \mathbf{c}_j \cdot \mathbf{L} | \mathbf{c}_j \cdot \mathbf{C} | -\mathbf{c}_j] \in \mathbb{F}^{1 \times (N+n+1)}$. In addition, this algorithm parses the signature as $\sigma = (\mathbf{s}, \text{salt})$, computes the (salted) hash of the message \mathbf{d} as $\mathbf{h} \leftarrow \mathcal{H}(\mathcal{H}(\mathbf{d})|\text{salt})$ and outputs the vector $\mathbf{v} = [\tilde{\mathbf{s}}|\mathbf{s}|1|\mathbf{h}]$, where \mathbf{s} is part of the signature and $\tilde{\mathbf{s}}$ is the ‘quadratic vector’ obtained by computing all products of pairs of elements in \mathbf{s} (with monomials ordered lexicographically), i.e., $\tilde{\mathbf{s}} \leftarrow [s[1]^2, s[1]s[2], \dots, s[n-1]s[n], s[n]^2]$. Clearly $\mathbf{z}'_j \cdot \mathbf{v} = 0$ if and only if $\mathbf{c}_j \cdot (\mathbf{Q}\tilde{\mathbf{s}} + \mathbf{L}\mathbf{s} + \mathbf{C}) = \mathbf{c}_j \cdot \mathbf{h}$, which is a probabilistic check of the original system of verification equations in Rainbow.

Efficient Verification for LUOV [5]. The **parse pk** procedure splits the public key into $PK = (\text{public.seed}, \mathbf{Q}_2)$ (the concise information needed to retrieve the full signer’s public key), and the auxiliary public information $PK.aux = \mathcal{H}$, i.e., a description of a full-domain hash function $\mathcal{H} : \{0, 1\}^* \rightarrow \mathbb{F}^m$, where $m = \text{rows}(\mathbf{Q}_2)$ and $\mathbb{F} = \mathbb{F}_{2^r}$. The **Check** procedure is trivial and always returns 1. This is because the whole LUOV verification can be written as a matrix-vector multiplication. The **GetM** algorithm takes in input $PK = (\text{public.seed}, \mathbf{Q}_2)$ and derives the full public key as done in the original verification: it runs $[\mathbf{C}|\mathbf{L}|\mathbf{Q}_1] \leftarrow \mathcal{G}(\text{public.seed})$ to get the constant constant (vector), the linear (matrix) and the first quadratic (matrix) parts of the verification equation; and then it reconstructs the full quadratic term as $\mathbf{Q} \leftarrow [\mathbf{Q}_1|\mathbf{Q}_2]$. Finally it appends to the public key the identity matrix $\mathbf{M} \leftarrow (\mathbf{C}, \mathbf{L}, \mathbf{Q}, -\mathbf{I}_{\text{rows}(\mathbf{Q})})$, we remark that \mathbf{M} can be seen as a matrix of blocks, where any block has the same height (number of rows), but different length (number of columns). The **GetZV** routine reads the matrix $\mathbf{Z}' = \mathbf{Z}$ made up of the rows $\mathbf{z}'_j = \mathbf{z}_j \leftarrow \mathbf{c}_j \mathbf{M} = (\mathbf{c}_j \cdot \mathbf{C}, \mathbf{c}_j \cdot \mathbf{L}, \mathbf{c}_j \cdot \mathbf{Q}, -\mathbf{c}_j)$. It also outputs the vector $\mathbf{v} = [1|\mathbf{s}|\tilde{\mathbf{s}}|\mathbf{h}]$, where \mathbf{s} is part of the signature $\sigma = (\mathbf{s}, \text{salt})$, $\tilde{\mathbf{s}}$ is the ‘quadratic vector’ obtained by computing all products of pairs of elements in \mathbf{s} , i.e., $\tilde{\mathbf{s}} \leftarrow [s[1]^2, s[1]s[2], \dots, s[n-1]s[n], s[n]^2]$, finally \mathbf{h} is the hash of the message and the salt, i.e., $\mathbf{h} \leftarrow \mathcal{H}(\mu||0x0||\text{salt})$. Clearly $\mathbf{z}'_j \cdot \mathbf{v} = 0$ if and only if $\mathbf{c}_j \cdot (\mathbf{C} + \mathbf{L}\mathbf{s} + \mathbf{Q}\tilde{\mathbf{s}}) = \mathbf{c}_j \cdot \mathbf{h}$, which is a probabilistic check of the original verification equation in LUOV.

3.3 Concrete Efficiency Estimates for Our Compiler

In what follows, we evaluate the efficiency gains provided by our compiler using the (r_0, e_0) -concrete efficiency notion of Equation (2). In brief, a realization of our compiler provides (r_0, e_0) -concrete amortized efficiency if r_0 is the smallest, non-negative integer for which it holds that $e_0 < 1$,

where e_0 is an upperbound on the ratio of running r_0 offline/online verifications versus r_0 standard signature verifications. For convenience, we estimate only the cost of the most expensive ‘steps’ in the verification, namely the ones involving several field element multiplications (matrix-vector products), and disregard the cost of adding elements, generating random values, reading algorithm inputs or evaluating hash functions. Moreover, we do not consider ad-hoc optimizations of matrix multiplication due to probabilistic checks using, e.g., Freivalds’ Algorithm or its variant [31]. We remark that in this work, we develop a more general and full-fledged approach that is not limited to the specific technique in [31]. Table 1 collects the common notation, while Table 2 displays a summary of our findings, that we motivate below.

Table 1: Parameters involved in the efficiency estimate of our compiler.

q	Modulus of the lattice or size of the field
n	Number of rows in the public key
$m \in \Omega(n \log q)$	Number of columns in the public key
β	Bound on the noise / size of signatures
σ or \mathbf{U}	Vector or matrix signatures
k	Number of steps in the online verification (confidence level)
r	Number of signatures verified (repetitions of onVer)
$\text{cost}(\text{alg})$	Number of field multiplications needed to compute alg

The computational complexity of **Ver** in **Mv**-style verification, e.g., [8,20,28,5,14] is dominated by a matrix-vector multiplication. Let $n = \text{rows}(\mathbf{M})$ and $m = \text{cols}(\mathbf{M})$, with $m \geq n$. Thus the cost of computing $\mathbf{M} \cdot \mathbf{v}$ is, in the worst case, nm field multiplications. Our offline verification algorithm executes k vector-matrix multiplications (one for each \mathbf{z}'_j in \mathbf{Z}'), resulting in knm multiplications in the worst case. The computational complexity of our online verification is dominated by the k vector-vector (inner) products $\mathbf{z}_i \cdot \mathbf{v}$, resulting in km multiplications in the worst case. Thus, the compiler presented in Section 3.1 outputs an efficient verification for signature with **Mv**-style verification that has the following concrete amortized efficiency:

$$\frac{\text{cost}(\text{offVer}) + r \cdot \text{cost}(\text{onVer})}{r \cdot \text{cost}(\text{Ver})} = \frac{knm + rkm}{rnm} = \frac{k}{r} + \frac{k}{n}. \quad (5)$$

Clearly the first addend in Equation (5) comes from amortizing the cost of **offVer** (over verifying r signatures), while the second term is the fix trade-off between the computational costs of **onVer** and **Ver** (at each and every verification). Table 2 collects the main results described in the reminder of the section. In detail, k_0 determines the length of the **svk**, the computational complexity and the unforgeability of the efficient verification; r_0 states how many signatures one should verify in order to have a concrete efficiency gain of e_0 (lower values of e_0 correspond to better efficiency gains). The last column displays how ‘cheaper’ the online verification is with respect to the fullfledged verification (ignoring the one-time cost of running the offline verification). Again, lower values in this column correspond to better efficiency; for instance, a ratio of 0.4% means that the computational cost of **Ver** is $99.6\times$ higher than the one of **onVer** (in other words, **onVer** is expected to be about $99\times$ faster).

For convenience, in our analysis we categorize signatures according to the size of their underlying algebraic structure (q exponential or polynomial in the security parameter).

The modulus q is exponential in λ : To the best of our knowledge, the only LBS constructions that fall in this category are the homomorphic signatures by Gorbunov et al. [21] and by Fiore et al. [16]. In this case, using our compiler (with some caveat, as we show in the next section) yields that the advantage in the `cmvEUF` experiment (in Equation (4)) is negligible in the security parameter for $k = 1$. However, in [21,16] the complexity of `Ver` is dominated by the matrix-matrix multiplication $\mathbf{A}\mathbf{U}$ where $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ is the fixed public key, and $\mathbf{U} \in \mathbb{Z}_q^{m \times m}$ is the signature⁴. We have computed parameters for this scheme according to the methodology by Albrecht et al. [1]. Setting $\lambda = 128$, $q = 2^\lambda$ and $n = 256$ yields that reduction algorithms (in particular, the optimized BKZ algorithm) would have runtime 2^{128} and would solve at most $\text{SIS}_{256, 2^{128}, 65536, 2^{80}}$, while the security of the scheme relies on a SIS instance with norm bound $\beta = 2^{49d}$, where d is the depth of the circuit. We can now use this set of parameters to determine the concrete amortized efficiency (as of Definition 4) reached by our compiler for the LBS in [16,21]. Let $k = 1$ and $n = 256$ in Equation (3.3), we want to extract the minimum r_0 for which $1/r + 1/256$ is smaller than 1. It is easy to see that $r_0 = 2$ suffices and we get $e_0 = 0.504 < 1/2 + 1/256$. In other words, the cost of setting up the online verification (running `offVer`) *plus* performing $r = 2$ online verifications is about half of the cost of $r = 2$ standard verifications, while preserving the security level. Moreover, for this set of parameters $\frac{\text{cost}(\text{onVer})}{\text{cost}(\text{Ver})} = \frac{k}{n} = \frac{1}{256} < 0.004$, i.e., our online verification requires about 0.4% of the computational cost of running the standard verification algorithm; alternatively, we can read this results as our `onVer` is $99\times$ faster than `Ver`.

The modulus q is polynomial in λ : This is the most common setting given the ‘small’ size of q . In this category fall the schemes by Gentry et al. [20], Boyen [8] and its improved version by Micciancio and Peikert [28], Boneh Franklin linearly homomorphic signature [6], Rainbow [15] and LUOV [5]. For the LBS constructions, in order to guarantee a negligible advantage in the `cmvEUF` experiment (see Equation (4)) we need to set an appropriate value of $k > 1$. We argue that ‘appropriate’ values of k are still ‘small’ in comparison to n and lead to a ‘good’ amortized efficiency even for ‘few’ verifications. We recall that for these constructions `Ver` computes a product $\mathbf{A}\boldsymbol{\sigma}$ where $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and the signature is just a vector $\boldsymbol{\sigma} \in \mathbb{Z}_q^{m \times 1}$. To guarantee the security of our efficient verification, the value k should be set so that q^{-k} be negligible. In other words, for the advantage in Equation (4) to be negligible it must hold that $q^{-k} = 2^{-\lambda}$. Hence, to estimate k , one need to first fix the value of λ , compute the corresponding q that can guarantee such level of security, and then extract k from the relation above.

Computing parameters for lattice-based schemes is not straightforward, as so far there is no unique way to derive the parameters from a given λ . However, a good measure of the security of a set of parameters can be extracted computing a value δ that was introduced by Gama and Nguyen [18]. Concretely, δ provides an indication of how reduction algorithms would perform against the hardness assumption underlying the lattice-based construction. Generally, the ‘smaller’ the δ , the ‘more secure’ the scheme.

For Boyen’s signature [7] and its variant by Micciancio and Peikert [28], we use the parameters provided in Figure 2 in [28]. Since in [28] they set $\delta = 1.007$, to ensure a fair comparison, we use the same value to compute the parameters Gentry et al.’s signature [20]. As a result, we observed that for this δ all of the schemes require about the same modulus $q \sim 2^{30}$ (for $n \sim 256$). For this set of parameters, our efficient verification provides 80 (resp. 250) bits of security with just only $k = 4$ (resp. $k = 9$). Thus our compiler achieves a (5, 0.81)-concrete amortized efficiency (resp. (10, 0.94)),

⁴ In [16] the dimension m additionally depends on the number $t \geq 1$ of distinct identities (users) involved in labeled program. For simplicity, in what follows we consider $t = 1$.

and a concrete tradeoff between `onVer` and `Ver` of $k/n = 1.5$ (resp. 0.0352), i.e., `onVer` is about $98\times$ faster than `Ver`. The take away is that for any $r > k$ there is a concrete efficiency gain.

We remark that $\delta = 1.007$ was ‘enough’ back in 2012, but it is now obsolete. As of now, this value of δ probably guarantees less than $\lambda = 60$ bits of security. This does not make our analysis meaningless. Indeed, the value of δ decreases for larger q and n (intuitively, this is because it is harder to find small vectors in larger lattices). In particular, the decrease in the value of δ is more dramatic for larger q than for larger n . Hence, a value of δ that would guarantee 128 bits of (post-quantum) security today implies a larger q than what we consider in this work, thus a k smaller than, e.g., $128/30 \sim 4.3 < 5$.

For **Rainbow**, we follow the latest guidelines provided during the second round of the NIST competition. We recall that the Rainbow signature scheme is based on the unbalanced oil and vinegar approach, and it is fully determined by the field on which it operates and a sequence of integer numbers indicating the amount of vinegar variables and oil variables per each layer. One of the current settings utilizes $\mathbb{F} = \mathbb{F}_{2^4}$ and a two-layer oil variable setting with $(v_1, o_1, o_2) = (32, 32, 32)$, which lead to $m = 96$ and $n = 64$ (for consistency in this paper we set n to be the number of rows of a matrix and m to denote the number of columns, classically the variables are swapped for multivariate signatures). A suitable k to achieve NIST security category level II is $k = 32$, since $q^{-k} = 2^{-4 \cdot 32} = 2^{-128}$. In regards to the concrete efficiency estimates, plugging in $k = 32$ and $n = 64$ in Equation (3.3) shows that the minimum number of repetitions r_0 to achieve amortized efficiency (i.e., for which we have $k/r + k/n < 1$) is $r_0 = 65$, the corresponding amortization factor is $e_0 = 0.9923 = 32/65 + 32/64$. For this set of parameters we have $\frac{\text{cost}(\text{onVer})}{\text{cost}(\text{Ver})} = \frac{k}{n} = \frac{32}{64} = 0.5$, in other words, our compiler produces an online verification that is $2\times$ *faster* than the standard verification. For $\mathbb{F} = \mathbb{F}_{2^8}$ and $(v_1, o_1, o_2) = (68, 36, 36)$, we have $m = 140$ and a suitable k in this case would be $k = 16$, since $q^{-k} = 2^{-8 \cdot 16} = 2^{-128}$. For the highest NIST security category, [14] suggests to use $\mathbb{F} = \mathbb{F}_{2^8}$ and $(v_1, o_1, o_2) = (92, 48, 48)$. As a result we have $m = 188$, $n = 96$ and again $k = 16$ but a better amortize efficiency factor $e_0 = 0,9666$ already for $r_0 = 20$. We remark that for this set of parameters $\frac{\text{cost}(\text{onVer})}{\text{cost}(\text{Ver})} < 0.166$, i.e., our compiler produces an online verification procedure that is $6\times$ *faster* than the standard verification. The verification is dominated by the matrix-vector multiplication, which requires nm field multiplications in the worst case. Our offline verification algorithm executes k vector-matrix multiplications (one for each $(\mathbf{c}_j, \mathbf{z}_j)$), resulting in knm multiplications. The online verification is dominated by the k vector-vector (inner) products $\mathbf{z}_i \cdot \sigma$, resulting in kn multiplications. Thus, for [8,20,28], the concrete amortized efficiency can be computed as $\frac{\text{cost}(\text{offVer}) + r\text{cost}(\text{onVer})}{r\text{cost}(\text{Ver})} = \frac{knm + rkm}{rnm} = \frac{k}{r} + \frac{k}{n}$. The first addend in Equation (3.3) comes from amortizing the cost of `offVer` (over r repetitions), while the second term is the fixed trade-off between the computational costs of `onVer` and `Ver` (at every verification).

3.4 Generalization to Signatures with Properties

Our model can be easily generalized to the case of signatures with properties (e.g., homomorphic signatures, group signatures, threshold signatures etc.). Such a generalization is particularly important in the case of LBS. Indeed, lattice-based hardness assumptions allow to build a wide range of LBS with properties that often have a verification algorithm very similar to standard **Mv**-style verification (cf. Figure 2).

Signatures with properties require more complex security definitions than standard signatures. Typically, they include (some specific flavor of) unforgeability. In Figure 4 we present a very general description of the unforgeability experiment for signatures with properties.

Table 2: A summary of the concrete efficiency achieved by various instantiations of our compiler. In the table, k_0 denotes the minimum accuracy level that securely realizes efficient verification, i.e., for which $\Pr[\text{bad}_1] = \frac{1}{q^{k_0}}$ is negligible (cf. proof of Theorem 1); r_0 is the smallest positive integer for which $e_0 < 1$ where e_0 is a (tight) upperbound on the ratio $e_0 > \frac{\text{cost}(\text{offVer}_{k_0}) + r \cdot \text{cost}(\text{onVer})}{r \cdot \text{cost}(\text{Ver})}$.

Signature Scheme	Ring-Field size	Min. Accuracy Level	Concrete Amortized Efficiency	Online Efficiency $\frac{\text{cost}(\text{onVer})}{\text{cost}(\text{Ver})}$
FMNP16 [16] GVW15 [21]	$q = 2^{\text{poly}(\lambda)}$	$k_0 = 1$	$(r_0 = 2, e_0 = 0.5)$	$< 0.4\%$
Boyen10 [8] GPV08 [20] MP12 [28]	$q = \text{poly}(\lambda)$	$k_0 = 3$	$(r_0 = 4, e_0 = 0.77)$	$< 1.1\%$
Rainbow[15] with $\mathbb{F}_{2^4} - (32, 32, 32)$	$q = \text{poly}(\lambda)$	$k_0 = 15$	$(r_0 = 18, e_0 = 0.99)$	$< 2.4\%$

Essentially, all of them require:

1. A setup phase, where a probabilistic routine (denoted **Setup** in Figure 4) generates a set of secret values **sval** and other auxiliary values **pval** (that includes the verification key **pk**). The latter are given as input to the adversary, while the former are used by the oracles.
2. A challenge phase, where the adversary is given access to some, possibly stateful, oracles (usually, at least an oracle that returns signatures by honest users), and has to output a message and a forged signature on it. We model this by defining two oracles: $OK(\cdot; \text{sval}, \text{st}_K)$: Returns signing/secret keys (of users or other entities that \mathcal{A} may corrupt). $O(\cdot; \text{st})$: Encompasses all the other possible oracles (signing, opening for group signatures, etc.).
3. A check phase, where the experiment checks whether the signature output by \mathcal{A} is valid and if \mathcal{A} won the experiment. The former requires an execution of the verification algorithm. The latter includes a variety of other checks that essentially are supposed to ensure that the signature was not trivially produced by the adversary (e.g., the signature has not been output by the signing oracle). We model this second check with the **WinCond** predicate. Clearly, the specification of **WinCond** depends on each primitive.

Adapting the syntax and security experiment of efficient verification to signatures with properties and to the generic security experiment described above is rather straightforward. Similarly, our

$\text{Exp}_{\mathcal{A}, \Sigma}^{\text{generic-unf}}(\lambda)$
1 : $(\text{pval}, \text{sval}) \leftarrow \text{Setup}(1^\lambda)$
2 : $\text{st}_K \leftarrow \emptyset, \text{st} \leftarrow \emptyset$
3 : $(\mu^*, \sigma^*, \text{aux}^*) \leftarrow \mathcal{A}^{OK(\cdot; \text{sval}, \text{st}_K), O(\cdot; \text{st})}(\text{pval})$
4 : if $\text{Ver}(\text{pk}, \mu^*, \sigma^*, \text{aux}^*) = 1 \wedge 1 \leftarrow \text{WinCond}(\mu^*, \sigma^*, \text{aux}^*, \text{pval}, \text{st}_K, \text{st})$
5 : return 1
6 : else return 0.

Fig. 4: Generic description of the unforgeability under adaptive chosen message attacks experiment for signatures with properties.

compiler of Section 3.1 can be easily adapted to signatures with properties that have a similar \mathbf{Mv} -style verification. In the following we analyze the impact of our compiler on the efficiency of some schemes whose verification is structured as in Figure 2: the constrained LBS in [34], the (indexed) attribute-based LBS in [23], the homomorphic LBS in [16], the threshold LBS in [4], and the multivariate-based ring signature RingRainbow [29].

Constrained Signatures (CS). CS allow a signer to sign a message only if either the message or the key satisfies certain preset constraints. The verification algorithm of the lattice-based instantiation of CS by Tsabary [34] is syntactically equivalent to the ‘standard’ verification: it includes a ‘ \mathbf{Mv} ’-style check (where the matrix has n rows) and a norm check. Hence, our compiler applies directly to this scheme. Unforgeability requires that $n \geq \lambda$ and $q \leq 2^\lambda$, so for an average value $q \sim 2^{32}$, we can set $k = 8 \ll \lambda$ so that $1/q^k = 1/2^{256}$. Remark that larger values of q (that could be required to have higher security guarantees) imply smaller values of k . Therefore, for this less conservative choice of parameters the efficiency gain is $\frac{\text{cost}(\text{onVer})}{\text{cost}(\text{Ver})} = \frac{k}{n} = \frac{8}{256} < 0.032$, i.e., the online verification requires about 3.2% of the computational cost of running the standard verification algorithm.

Indexed Attribute-based Signatures (iABS) and Homomorphic Signatures (HS). iABS allow a signer to generate a valid signature on a message only if the signer holds a set of attributes that satisfy some policy (represented by a circuit C). HS allow a signer to sign messages μ_i so that it is possible to publicly derive a valid signature for a message μ that corresponds to the output of a computation on the original messages, i.e., $\mu = C(\mu_1, \dots, \mu_r)$. According to the type of homomorphism supported by the scheme, the circuit C can encode only linear functions, polynomial functions, or any function of bounded multiplicative degree. In both iABS in [23] and HS in [16] the signature verification is composed by three steps:

1. Computation of the public matrix \mathbf{M} from the circuit C (either the policy, or the homomorphic computation specified by the labelled program);
2. An ‘ \mathbf{Mv} ’-style check;
3. A norm check on the signature.

The first step is critical because the public matrix \mathbf{M} can be generated through a non-linear transformation, i.e., it might include multiplications of the public matrix by itself (or by a gadget matrix). This would not allow to compute the first step online from the \mathbf{z}_i ’s, but the verifier would have to use \mathbf{M} and the \mathbf{c}_i ’s instead, defying the purpose of our compiler. Hence, our compiler can be applied to these signatures in an efficient way only if either (1) C involves solely linear operations on the public matrix, or (2) C is fixed, or (3) C is known before running verification. In these cases, we achieve efficient verification by letting `offVer` take as (additional) input C and compute \mathbf{M} using the algorithm `PubEval` from [22]. The vectors $(\mathbf{Z}', \mathbf{v})$ used in the verification might (as in [16]) or might not (as in [23]) depend on the message. In the latter case the subroutine `GetZV` in `onVer` simply returns the input.

The impact of the compiler on the efficiency of HS was already analyzed in Section 3.3. Regarding the iABS, the suggested value of the modulus q is such that $q \geq n^8$. The standard requirement $n \geq 2$ already implies that $1/q^k \leq 1/(2^8)^k = 1/256^k$. However, to guarantee the hardness of lattice-based problems usually n needs to be at least $n = 128$. In this case $q \geq 2^{56}$, hence $k = 5$ already guarantees the unforgeability of this iABS. As $n = O(d \log d)$ (where d is the depth of C) and the efficiency gain can be bounded as follows: $\frac{\text{cost}(\text{onVer})}{\text{cost}(\text{Ver})} \leq \frac{k}{O(d \log d)} = \frac{5}{O(d \log d)}$. From this

inequality is clear that already for a circuit of depth 4 the online verification only requires 62.5% of the computation required by standard verification; the impact of our compiler increases for larger size of the circuit.

Threshold Signatures (TS). TS allow h out of ℓ parties to produce a signature on a message. Unforgeability is guaranteed for up to t colluding parties. Tsabary [34] introduced a compiler that allows to distribute the signature generation step of the GPV08 signature, and convert it into a TS. The idea is to share the signing trapdoor among the parties using a h -out-of- ℓ secret sharing scheme. Signing requires at least h parties to come together to generate a signature satisfying a \mathbf{Mv} -type equation (where \mathbf{M} is the public verification key). Verification is composed by the standard \mathbf{Mv} equation and norm checks. Therefore, the thresholdizing compiler is composable with our compiler for efficient verification. As neither of them change the parameters of the underlying GPV08 scheme, the efficiency gain is the same (cf. Section 3.3).

RingRainbow [29]. RingRainbow is a ring signature scheme – i.e., a signature that allows a user to sign a message anonymously on behalf of a group – based on multivariate equations. This scheme is a hash-and-sign type of signature built as a modification of Rainbow. Verification requires to check whether the signature satisfies a multivariate quadratic system, and can be converted in a \mathbf{Mv} -style verification with the same technique used for Rainbow (cf. Section 3.3). Therefore, our compiler can be applied to RingRainbow as well. To evaluate the efficiency gain due to our compiler, we consider the efficient version of RingRainbow, (whose parameters can be found in Table 2 in [29]). For $\lambda = 128$ and a group of 5 users the authors set $\mathbb{F} = \mathbb{F}_{2^8}$ and $(v_1, o_1, o_2) = (36, 21, 22)$, which yield $m = 5 * (v_1 + o_1 + o_2) = 395$ and $n = 43$. Theorem 1 requires at least $1/q^k = 1/2^{256}$ for 128 bits of post-quantum security, hence $k = 32$. Plugging this in Equation (3.3) yields that the minimum number of repetitions r_0 to achieve amortized efficiency is $r_0 = 126$, that results in an amortization factor of $e_0 = 32/126 + 32/43 = 0.744$. In this case, our compiler produces an online verification that is $\frac{\text{cost}(\text{onVer})}{\text{cost}(\text{Ver})} = \frac{k}{n} = \frac{32}{43} < 0.75$, in other words, our compiler produces an online verification that requires only 75% of the computation required by the standard verification.

4 Modeling Stateful and Flexible Verification

We introduce our framework for *flexible* verification of digital signatures. The basic idea is to design a flexible verification algorithm, `flexVer`, that instead of returning a binary answer (accept/reject) as done by `Ver`, outputs a value α expressing the confidence on the validity of the signature. The confidence α is proportional to the amount of computation invested in the verification: the more verification steps performed the higher the confidence. In this section we show how to realize flexible verification disregarding performance impacts. In the next section we address the question to securely achieve efficient and flexible verification simultaneously.

Differently from Le et al. [25], who define flexible signatures as a standalone primitive, we prefer to treat flexible verification as an add-on feature to enhance existing schemes. Compared to [25] our model is more general, yet slightly more involved. In particular, we let `flexVer` be a *stateful* algorithm; the state `st` is maintained by the verifier and should not be disclosed to the adversary. Albeit `st` can be seen as a secret, updatable, compact verification key, we remark that it is always possible to verify signatures using the signer’s public key and the standard verification procedure. Moreover, any verifier can generate a `st` from the signer’s `pk`, but `st` does not allow the verifier to forge signatures in the name of the signer.

4.1 Syntax

The core idea behind flexible verification is to derive from the original verification procedure of a signature scheme (Ver) an alternative algorithm flexVer with certain properties. Concretely, this task boils down to identifying a sequence of $K_{\text{flex}} + 1$ atomic instructions (that we informally call ‘steps’) that ensure that, if interrupted at any point, flexVer returns a meaningful value about the correctness of the signature. In more detail, if one of the steps fails, flexVer returns $\alpha = \perp$ and rejects the signature. On the other hand, if none of the initial t steps fails, it is possible to derive a real value $\alpha_{\text{flex}}(t) \in [0, 1]$ stating the probability that the given signature is valid. Intuitively, we want that the more steps flexVer executes successfully, the higher the value of $\alpha_{\text{flex}}(t)$.

Definition 6 (Flexible Verification). *Let $\Sigma = (\text{KeyGen}, \text{Sign}, \text{Ver})$ be a signature scheme. We say that Σ admits flexible verification if there exists a non-negative integer $K_{\text{flex}} \in \mathbb{Z}_{\geq 0}$, an efficiently computable confidence function $\alpha_{\text{flex}} : \{0, \dots, K_{\text{flex}}\} \rightarrow [0, 1]$, a set of admissible states⁵ \mathcal{S} (that depends on K_{flex} and a public key pk), and an algorithm flexVer consisting of $K_{\text{flex}} + 1$ steps $\text{flexVer}_0, \dots, \text{flexVer}_{K_{\text{flex}}}$, such that flexVer has the following syntax:*

$\text{flexVer}(\text{st}, \text{pk}, \mu, \sigma, t)$: *This a randomized algorithm takes in input a state $\text{st} \in \mathcal{S}$, a public key pk (generated by KeyGen), a message μ , a signature σ and an interruption position $t \in \mathbb{Z}_{\geq 0}$. It outputs a value $\alpha \in [0, 1] \cup \perp$. Concretely, flexVer works as shown below: it is made of $K_{\text{flex}} + 1$ algorithms $\text{flexVer}_j(\text{st}, \text{pk}, \mu, \sigma)$ that update the state and return a bit b . If $b = 0$ the flexible verification outputs $\alpha = \perp$ (σ is rejected); otherwise, $b = 1$ and the flexible verification upgrades the confidence level to $\alpha \leftarrow \alpha_{\text{flex}}(j)$ (σ is valid at step j).*

$\text{flexVer}(\text{st}, \text{pk}, \mu, \sigma, t)$	
1:	$\alpha \leftarrow \perp$
2:	if $t < 0$: return \perp
3:	if $t > K_{\text{flex}}$: set $t \leftarrow K_{\text{flex}}$
4:	for $j = 0, \dots, t$
5:	$(b, \text{st}) \leftarrow \text{flexVer}_j(\text{st}, \text{pk}, \mu, \sigma)$
6:	if $(b = 0)$: return \perp
7:	else $(b = 1)$: $\alpha \leftarrow \alpha_{\text{flex}}(j)$
8:	return α

Additionally flexVer satisfies **flexible correctness**: For a given security parameter λ , for any key pair $(\text{sk}, \text{pk}) \leftarrow \text{KeyGen}(\lambda)$, for any admissible state $\text{st} \in \mathcal{S}$, for any $\mu \in \mathcal{M}$ and signature σ such that $\text{Ver}(\text{pk}, \mu, \sigma) = 1$ and for any value $t \in \{0, \dots, K_{\text{flex}}\}$, it holds:

$$\Pr[\text{flexVer}(\text{st}, \text{pk}, \mu, \sigma, t) = \alpha_{\text{flex}}(t)] = 1.$$

Flexible correctness as defined above essentially states that on valid signatures, i.e., signatures that would be accepted by the standard verification, flexVer accepts with confidence $\alpha_{\text{flex}}(t)$. We follow the approach of [25] and let the flexible verification algorithm output a value α that either rejects the signature ($\alpha = \perp$), or accepts it with certainty α in the real interval $[0, 1]$. We use the same interruption variable t as in [25] to model runtime interruptions of the algorithm execution.⁶ We

⁵ The set \mathcal{S} includes the initial state $\text{st} = \emptyset$ and any possible state output by some flexVer_i .

⁶ Our $\alpha_{\text{flex}}(\cdot)$ is essentially the inverse of the function $\text{iExtract}_{\Sigma}(\cdot)$ in [25].

point out that while t is input to flexVer , it is *not* given to each flexVer_j ; this syntax models the fact that the algorithm must work without knowing where it will stop, which is essential to capture arbitrary interruptions. In the case of an interruption, $t < K_{\text{flex}}$, the algorithm would return a *premature* decision α on the validity of the signature: if $\alpha = \perp$ the signature is invalid (and this is certain, there are no false negatives); if $\alpha \in [0, 1]$ this means that according to the t steps performed, the signature seems valid.

Defining the confidence function $\alpha_{\text{flex}}(\cdot)$ is fundamental to showing a construction of flexible verification for a given signature scheme. Any signature scheme Σ admits trivial flexible verification and confidence function. One way to see it is to consider Ver as a sequence of K_Σ instructions (e.g., computation or verification steps); set $K_{\text{flex}} = K_\Sigma$ and flexVer_j be the j -th step of Ver . So the confidence function can be defined as

$$\alpha_{\text{flex}}(t) = \begin{cases} 0 & \text{for } t \in \{0, \dots, K_\Sigma - 2\} \\ 1 & \text{for } t = K_\Sigma - 1 \end{cases}$$

A trivial confidence function entails un-interesting flexible verification: on correct signatures, flexVer

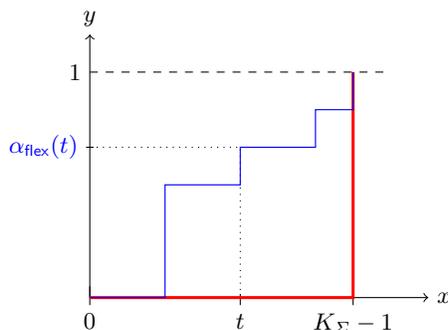


Fig. 5: Examples of confidence functions for flexible verification: the trivial case (thick, red); a non-trivial case (thin, blue).

returns 0 if interrupted at any step $t < K_\Sigma - 1$. While all signature schemes admit a trivial confidence function, not all support a non-trivial one, (e.g., a step function as shown in Figure 5). We remark that the confidence function $\alpha_{\text{flex}}(\cdot)$ depends both on the scheme Σ and on how each flexVer_i is defined. In Section 4.3 we show non-trivial realizations for post-quantum secure signature schemes based either on multivariate quadratic equations or on lattices. It is easy to see that, without loss of generality, $\alpha_{\text{flex}}(\cdot)$ has a non-decreasing trend, $\alpha_{\text{flex}}(0) = 0$ and $\alpha_{\text{flex}}(K_{\text{flex}}) \leq 1$. Looking ahead, we will define $\alpha_{\text{flex}}(t)$ to be $1 - \text{bad}(t)$, where $\text{bad}(t)$ represents the probability of accepting a forgery after t verification steps. While $\text{bad}(t)$ decreases, $\alpha_{\text{flex}}(t)$ increases and so does the confidence in the correctness of the signature.

Flexible Verification vs. Offline/Online Signature Verification. While both flexible verification and efficient verification may be seen as possible ways of reducing the computational cost of a signature verification, the way this objective is achieved in the two models is quite different. One can think that these are two sides of the same coin. In flexible verification, the verification algorithm (and thus the verifier) is unaware of when the computation will be interrupted, and *its execution is independent of t* . On the other hand, the offline verification (and thus the verifier) sets precisely what is the confidence level k – which may be seen as an interruption value for onVer – that otherwise

would check all of the verification checks imposed by `Ver`. The fundamental difference is that in the latter setting, the (online) verification is aware of the confidence level k (seen as interruption value), and *adapts its execution to k* .

Stateful vs. Stateless Verification We choose to define flexible signatures as stateful algorithms to keep the framework as general as possible. The same model could capture stateless flexible verification by considering the special case in which the state `st` is always \emptyset . In the latter case, there is no need for a verification oracle in the security experiment (Figure 6). In this sense, our approach is more general than the one by Le et al. [25] as we can potentially capture more schemes.

4.2 Security Model

Our security model revisits the notion of existential unforgeability under chosen message attack for flexible signatures proposed by Le et al. [25] with three twists.

First, our security game needs to take into account that `flexVer` maintains a possibly non trivial state. We handle this by allowing the adversary to interact with the flexible verification oracle `OflexVer` during the query phase (similarly to what we did in the efficient verification model of Section 3) and with the signing oracle `OSign`, in a concurrent manner.

Second, queries to `OflexVer` have the form (μ, σ, t') , where t' is the desired interruption value submitted by \mathcal{A} (and chosen adaptively). To make the model generic and reduce the assumptions on \mathcal{A} , we include an interruption oracle `OInt` that takes as input t' and returns the actual value t to be used in `flexVer`. We discuss in detail the role of `OInt` in a dedicated paragraph after stating our security notion. For the purpose of this work, we consider the strongest security model in which the interruption oracle returns the adversary's value, i.e., $t \leftarrow \text{OInt}(t')$ with $t = t'$. This resembles side-channel attack settings, where \mathcal{A} may try to freeze the execution of the verification.

Finally, a minor change: instead of outputting a single bit, our experiment returns a pair (b, t^*) . The bit $b \in \{0, 1\}$ flags the absence or the potential presence of a forgery, while $t^* \in \{0, \dots, K_{\text{flex}}\}$ reports the interruption position used in the final flexible verification. Including t^* in the output of the experiment allows us to measure security in terms of *how far* is the probability of \mathcal{A} ‘winning’ the experiment, from the expected value $\alpha_{\text{flex}}(t^*)$. Concretely, we consider a flexible verification to be *secure* if, for any PPT adversary, the probability that the security experiment returns $(1, t^*)$ is only negligibly higher than the expected probability of non detecting a forgery after t^* verification steps. For the purpose of this work, t^* is the interruption parameter chosen by \mathcal{A} for its forgery.

Our security game and experiment for existential unforgeability under adaptive chosen message *with flexible verification* attack (flexEUF) are reported in Figure 6.

Definition 7 (Security of Flexible Signatures (flexEUF)). *Let Σ be a signature scheme that admits a non-trivial realization of flexible verification `flexVer` with corresponding confidence function α_{flex} . For a given security parameter λ , $(\text{flexVer}, \alpha_{\text{flex}})$ realize a secure flexible verification for Σ if for all PPT adversaries \mathcal{A} the success probability in the flexEUF experiment in Figure 6 is negligible, i.e.,:*

$$\text{Adv}_{\mathcal{A}, \Sigma}^{\text{flexEUF}}(\lambda) = \Pr \left[\text{Exp}_{\mathcal{A}, \Sigma}^{\text{flexEUF}}(\lambda) = (1, t^*) \right] - (1 - \alpha_{\text{flex}}(t^*)) = \varepsilon \leq \varepsilon(\lambda).$$

Intuitively, Definition 7 states that an adversary has only negligible probability to make the flexible verification output a confidence value α^* higher than the expected one. Let $\text{bad}(t)$ denote the probability of accepting a forgery after t verification steps. Then by setting $\alpha_{\text{flex}}(t) = 1 - \text{bad}(t)$, we get $\text{Adv}_{\mathcal{A}, \Sigma}^{\text{flexEUF}}(\lambda) = \Pr \left[\text{Exp}_{\mathcal{A}, \Sigma}^{\text{flexEUF}}(\lambda) = (1, t^*) \right] - \text{bad}(t^*) \leq \varepsilon(\lambda)$.

$\text{flexEUF}(\Sigma, \lambda)$	$\text{Exp}_{\mathcal{A}, \Sigma}^{\text{flexEUF}}(\lambda)$	$\text{OSign}_{\text{sk}}(\mu)$
1: $L_S \leftarrow \emptyset$	1: $(\mu^*, \sigma^*, t') \leftarrow \text{flexEUF}(\Sigma, \lambda)$	1: $L_S \leftarrow L_S \cup \{\mu\}$
2: $\text{st} \leftarrow \emptyset$	2: $\beta \leftarrow \text{Ver}(\text{pk}, \mu^*, \sigma^*)$	2: $\sigma \leftarrow \text{Sign}(\text{sk}, \mu)$
3: $(\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(1^\lambda)$	3: $t^* \leftarrow \text{OInt}(t')$	3: return σ
4: $(\mu^*, \sigma^*, t') \leftarrow \mathcal{A}^{\text{OSign}, \text{OflexVer}}(\text{pk}, \lambda)$	4: $\alpha \leftarrow \text{flexVer}(\text{st}, \text{pk}, \mu^*, \sigma^*, t^*)$	$\text{OflexVer}_{\text{st}, \text{pk}}(\mu, \sigma, t')$
5: return (μ^*, σ^*, t')	5: if $\mu^* \in L_S \vee \alpha = \perp \vee \beta = 1$	1: $t \leftarrow \text{OInt}(t')$
	6: return $(0, t^*)$	2: $\alpha \leftarrow \text{flexVer}(\text{st}, \text{pk}, \mu, \sigma, t)$
	7: return $(1, t^*)$	3: return α

Fig. 6: Security game, experiment and oracles for stateful and flexible signature verification: unforgeability under adaptive chosen message with flexible verification attack.

In this work, we prove security in the strongest model where $t' = t$, i.e., \mathcal{A} has the power to choose when to stop the verification. Since we put no restriction on the values t queried by \mathcal{A} to OflexVer during the game, we will see that by running OflexVer on ‘too few’ steps, \mathcal{A} may learn information about the internal state st .

Modelling Interruptions. In [25], unexpected interruptions are modeled via an interruption oracle $\text{iOracle}(\lambda)$ that returns a value $t \in \{0, \dots, K_{\text{flex}}\}$ used by the flexible verification. However, it is not clear whether \mathcal{A} may control iOracle . We overcome these ambiguities by letting \mathcal{A} output t' at every flexible verification query. It is possible to relax and generalize our model by introducing an interruption oracle OInt , programmed at the beginning of the game. At each verification query, OInt takes as input the adversary’s suggestion for an interruption position t' and outputs the value t to be used by the flexible verification. In case $t = t'$, we are modelling side channel attacks, but we can also let t be independent of t' . A realistic definition of OInt is outside the scope of this work.

4.3 A compiler for flexible ‘Mv’-style Verifications

We exhibit a compiler to obtain flexible verification for the class of post quantum signature schemes with ‘Mv’-style verification introduced in Section 3.1. Specifically, our compiler for flexible verification builds on our previous compiler for efficient verification.

Two Insecure Solutions. Let us warm up with an instructive example. For schemes with ‘Mv’-style verification there is a natural way to divide the verification algorithm into atomic steps: the initial step, flexVer_0 , includes the consistency checks (e.g., the norm bound on the signature); while for $i = 1$ to $K_{\text{flex}} = \text{rows}(\mathbf{M})$, flexVer_i performs the check $\mathbf{M}[i, *] \cdot (\sigma, \mathbf{u}) = 0 \pmod q$, where $\mathbf{M}[i, *]$ denotes the i -th row of the matrix \mathbf{M} . Notably this solution does not require a secret state. Checking only $t < K_{\text{flex}}$ constraints, however, may not be a secure choice. Indeed, an adversary that knows \mathbf{M} and \mathbf{u} can efficiently produce a signature that verifies the first $t \ll K_{\text{flex}}$ linear constraints and trivially break unforgeability (flexEUF). A naive way to bypass this issue is to randomize the order in which the steps are performed (assuming flexVer_0 is always run). In this case, setting $t^* = 1$, the adversary can produce a valid forgery for say the i^* -th row of \mathbf{M} and win the flexEUF game with probability $1/K_{\text{flex}}$ (which corresponds to the probability that flexVer executes only flexVer_0 and flexVer_{i^*}). Notably, $1/K_{\text{flex}}$ is not negligible, so neither of the two presented approaches is secure.

$\text{flexVer}_0(\text{st}, \text{pk}, \mu, \sigma)$	$\text{flexVer}_i(\text{st}, \text{pk}, \mu, \sigma)$
1 : $\text{svk} \leftarrow \text{offVer}(\text{pk}, K_{\text{flex}})$	1 : $b \leftarrow 0$
2 : $\text{parse svk} = (K_{\text{flex}}, \mathbf{Z}, PK.\text{aux})$	2 : $\text{parse st} = (\mathbf{Z}', \mathbf{v})$
3 : $b \leftarrow \text{Check}(PK.\text{aux}, \mu, \sigma)$	3 : if $\mathbf{Z}'[i, *] \cdot \mathbf{v}[*, i] = 0 \pmod q$
4 : $\text{st} \leftarrow \text{GetZV}(\text{svk}, \mu, \sigma)$	4 : return $(b \leftarrow 1, \text{st})$
5 : return (b, st)	5 : return $(b \leftarrow 0, \text{st})$

$\alpha_{\text{flex}} : \{0, \dots, K_{\Sigma}\} \rightarrow [0, 1], \quad \alpha_{\text{flex}}(t) = (1 - \frac{1}{q^t})$

Fig. 7: Generic compiler for flexible verification of \mathbf{Mv} -style signatures. The algorithms offVer , Check and GetZV are defined in Section 3.1, Figure 3. Here $K_{\text{flex}} = \text{rows}(\mathbf{M})$.

Our Compiler. We define the algorithm flexVer by describing its atomic steps flexVer_i . Analogously to Section 3.1, we present a generic compiler for signatures with \mathbf{Mv} -style verification in Figure 7.

To formally define flexible verification we need to specify the value K_{flex} (total verification steps), the set \mathcal{S} (valid states), and the confidence function α_{flex} . The value K_{flex} sets the upper bound on the number of linear constraints the verifier wants to check, hence $K_{\text{flex}} = \text{rows}(\mathbf{M})$, where \mathbf{M} is the matrix employed in the matrix-vector multiplication check in the signature verification. The set \mathcal{S} includes \emptyset and any possible state output by some flexVer_i , specifically $\mathcal{S} = \{0, 1\} \times \mathbb{Z}_q^{\text{rows}(\mathbf{Z}') \times \text{cols}(\mathbf{Z}')} \times \mathbb{Z}_q^{\text{rows}(\mathbf{v}) \times \text{cols}(\mathbf{v})} \times \{0, 1\}^\lambda \cup \emptyset$.

We extract the confidence level from the probability of a flex forgery (as motivated by the proof of security given in Theorem 1). Alike efficient verification, the probability that an adversary creates a flex forgery with interruption step t is $1/q^t$. If the size of the underlying algebraic structure is $q = 2^{\text{poly}(\lambda)}$ this probability is negligible even for $t = 1$, thus for signatures with large q efficient verification and flexible verification coincide, trivially. The interesting case is $q = \text{poly}(\lambda)$, as the adversary could create a flex forgery with non negligible probability. We remark that in this section we are not targeting efficiency, and our instantiations of flexible verification refresh the svk produced by offVer at every verification query. This way, \mathcal{A} cannot exploit the information possibly leaked by a flex forgery in future forgery attempts.

Theorem 2. *Let Σ be a signature scheme with \mathbf{Mv} -style verification (as of Fig. 2). Then our compiler for flexible verification (in Figure 7) returns $(\text{flexVer}, \alpha_{\text{flex}})$ that realize secure flexible verification for Σ .*

Proof. Recall that an adversary \mathcal{A} wins the security experiment in Definition 7 if it outputs a pair message-signature (μ^*, σ^*) and an interruption t' such that (1) (μ^*, σ^*) is rejected by Ver , but accepted by the flexible verification flexVer when it is interrupted at step $t^* \leftarrow \text{OInt}(t')$, and (2) the flexible verification algorithm outputs a too high confidence level $\alpha_{\text{flex}}(t^*)$. Following Definition 7 we can realize secure flexible verification by setting $\alpha_{\text{flex}}(t) = 1 - \Pr[\text{Exp}_{\mathcal{A}, \Sigma}^{\text{flexEUF}}(\lambda) = (1, t)] + \varepsilon(\lambda) \quad \forall t$. The core part of the proof is to estimate this probability.

Recall that our compiler for efficient \mathbf{Mv} -style verification (in Figure 7) runs offVer at every verification query (step 1 in flexVer_0). This means that every verification query is answered using a freshly generated svk . In particular, the final verification (line 4 in the $\text{Exp}_{\mathcal{A}, \Sigma}^{\text{flexEUF}}(\lambda)$ in Figure 8) checks \mathcal{A} 's output using independent randomness from the previous queries. So, whatever information the adversary may collect from previous queries is useless to win the experiment. As

a consequence, the probability that the adversary wins the game equals the probability that the adversary outputs a valid forgery *without querying* `OflexVer`. The latter is precisely the probability of a bad forgery we investigated in Theorem 1, hence $\Pr[\text{Exp}_{\mathcal{A},\Sigma}^{\text{flexEUF}}(\lambda) = (1, t^*)] = \frac{1}{t^*}$ and:

$$\begin{aligned} Adv_{\mathcal{A},\Sigma}^{\text{flexEUF}}(\lambda) &= \Pr[\text{Exp}_{\mathcal{A},\Sigma}^{\text{flexEUF}}(\lambda) = (1, t^*)] - (1 - \alpha_{\text{flex}}(t^*)) \\ &= \frac{1}{q^{t^*}} - \left(1 - \left(1 - \frac{1}{q^{t^*}}\right)\right) = 0. \end{aligned}$$

□

4.4 Flexible and Efficient Verification

Given the two notions of efficient verification (Section 3) and flexible verification (Section 4), the reader might be wondering now whether it would be possible to build a compiler for flexible *and* efficient verification of **Mv**-style signatures. Our approach is to start from our compiler for flexible verification (depicted in Fig. 7), and skip line 1 of `flexVer0` (where it runs `offVer` to generate `svk`) in all verifications except for the first one. This means that `svk` is created at the first verification and reused in all subsequent runs of `flexVer`. While this obviously improves the (amortized) efficiency of the compiler, it also impacts security.

We now verify the unforgeability of our compiler for q exponential. From Theorem 1 we already know that in this case verification only checks the linear equation for one \mathbf{c} . Hence adding flexibility to the model does not impact it, as the adversary can only interrupt verification either before or after the linear equation has been checked.

Theorem 3 (informal). *For signatures with **Mv**-style verification relying on algebraic structures of size $q = 2^{\text{poly}(\lambda)}$ our compiler outlined for efficient flexible verification is unforgeable according to Definition 7 and achieves $(2, 1/2 + 1/\text{rows}(\mathbf{M}))$ -concrete amortized efficiency as per Definition 4.*

Proof. (sketch) The proof is based on the same argument used in the proof of Theorem 1. Using the same notation as in Theorem 1, take the same sequence of q_V hybrid games and replace any appearance of `onVer` by `flefVer` (essentially each game delays the query on which the first flexible verification is run). Since q is exponential in the security parameter, the interruption values t' can be all set to 1. We can bound the adversary's success probability by the occurrence of **bad** events.

$$\Pr[\text{Exp}_{\mathcal{A},\Sigma}^{\text{flexEUF}}(\lambda) = (1, t^*)] \leq \sum_{i=1}^{q_V} \Pr[\text{bad}_i] + \Pr[\mathbf{G}_{q_V} = 1] \leq \frac{q_V}{q} + \frac{1}{q}.$$

We used $\Pr[\mathbf{G}_{q_V} = 1] \leq \frac{1}{q}$ since in that game the vectors `svk` is freshly generated. To conclude, $q = 2^{\text{poly}(\lambda)}$ and $q_V \leq \text{poly}(\lambda)$ it holds that,

$$\begin{aligned} Adv_{\mathcal{A},\Sigma}^{\text{flexEUF}}(\lambda) &= \Pr[\text{Exp}_{\mathcal{A},\Sigma}^{\text{flexEUF}}(\lambda) = (1, t^*)] - (1 - \alpha_{\text{flex}}(t^*)) \\ &\leq \frac{q_V}{q} + \frac{1}{q^{t^*}} - \left(1 - \left(1 - \frac{1}{q^{t^*}}\right)\right) = \frac{q_V}{q} = \varepsilon(\lambda). \end{aligned}$$

In this setting the concrete efficiency of our compiler is achieved already with $r_0 = 2$ repetitions, and the corresponding amortized efficiency value is $\mathbf{e}_0 = \frac{\text{cost}(\text{offVer}(\text{pk}, 1))}{r_0 \cdot \text{cost}(\text{Ver}(\text{pk}, \mu, \sigma))} + \frac{r_0 \cdot \text{cost}(\text{onVer}(\text{svk}, 1, \mu, \sigma))}{r_0 \cdot \text{cost}(\text{Ver}(\text{pk}, \mu, \sigma))} = \frac{1}{2} + \frac{1}{\text{rows}(\mathbf{M})}$. □

In case $q = \text{poly}(\lambda)$, the freshness argument on svk can no longer be used to bound the probability of finding a forgery: some non-negligible amount of information about the secret combinators \mathbf{c}_i in svk is leaked at every flex-forgery submitted during verification queries. To formally handle these cases, in the next Section we introduce the notion of r -bounded randomness reuse.

5 Modeling Efficient & Flexible Signature Verification with r -Bounded Randomness Reuse

We introduce the concept of *flexible* and *efficient* signature verification (flefs) with r -bounded randomness reuse (or r -flefs in short). Similarly to Definition 6 (flexible signatures), this *sustainable* variant is defined for a given value k , that determines the maximum desired confidence level achievable by the verification. In addition to k , we need a second parameter, r , that determines the maximum number of times the verification randomness can be reused. For correctness and security, both k and r are input to the confidence function, which now is named α_{flef} .

Definition 8 (r -flefs). *A signature scheme Σ admits a (k, r) -efficient and flexible verification realization if there exist two positive integers k and r ; an efficiently computable confidence function $\alpha_{\text{flef}} : \{0, \dots, k\} \times \{0, \dots, r\} \rightarrow [0, 1]$; a set of admissible sequences of states $\mathcal{S} = \{st^{(1)}, st^{(2)}, \dots\}$ (each sequence $st^{(j)}$ contains $r+1$ states, i.e., $st^{(j)} = (\mathbf{st}_i)_{i=0}^r, \mathbf{st}_0 = \emptyset$) and a flexible verification algorithm flefVer consisting of $k+1$ steps $\text{flefVer}_0, \dots, \text{flefVer}_k$ satisfying the properties of correctness, efficiency and security stated below.*

r -Reuse k -Flexible Correctness For a given security parameter λ , for any key pair $(\text{sk}, \text{pk}) \leftarrow \text{KeyGen}(\lambda)$, for any one sequence of admissible states $st \leftarrow \mathcal{S}$, $st = (\mathbf{st}_i)_{i=0}^r$, for any choice of r message-signature pairs (μ_i, σ_i) with $i \in \{1, \dots, r\}$, $\mu_i \in \mathcal{M}$ and σ_i such that $\text{Ver}(\text{pk}, \mu_i, \sigma_i) = 1$ and for any sequence of interruption values $t_i \in \{1, \dots, k\}$, it holds that:

$$\Pr[\text{flefVer}(\mathbf{st}_i, \text{pk}, \mu_i, \sigma_i, t_i) = \alpha_{\text{flef}}(t_i, i)] = 1 \quad \forall i \in \{1, \dots, r\}.$$

r_0 -Concrete Amortized Efficiency For a given security parameter λ , for any key pair $(\text{sk}, \text{pk}) \leftarrow \text{KeyGen}(\lambda)$, for any one sequence of admissible states $st \leftarrow \mathcal{S}$, there exists a minimal positive integer number $r_0 \leq r$ for which it holds that:

$$\sum_{i=0}^{r_0-1} \text{cost}(\text{flefVer}(\mathbf{st}_i, \text{pk}, \mu_i, \sigma_i, k)) < r_0 \cdot \text{cost}(\text{Ver}(\text{pk}, \mu, \sigma)) \quad (6)$$

r -bounded flexible security Figure 8 collects a description of our security game and experiment for existential unforgeability under adaptive chosen message attack for signatures with *flexible and efficient verification* (r -flefEUF).

Definition 9 (r -Bounded Flexible Security (r -flefEUF)). *Let Σ be a signature scheme that admits a non-trivial realization of (k, r) -efficient and flexible verification flefVer . Then for a given security parameter λ the derived verification is existential unforgeable under adaptive chosen message attack with flexible and efficient verification (r -flefEUF) if for all efficient PPT adversaries \mathcal{A} the success probability in the r -flefEUF experiment is:*

$$\Pr \left[\begin{array}{l} \text{Exp}_{\mathcal{A}, \Sigma, r}^{\text{r-flefEUF}}(\lambda, k, r) = (\text{ctr}^*, t^*) \\ \wedge (\text{ctr}^*, t^*) \neq (0, 0) \end{array} \right] \leq (1 - \alpha_{\text{flef}}(t^*, \text{ctr}^*)) + \varepsilon(\lambda).$$

<p>r-fleEUf(Σ, λ, k)</p> <hr/> <p>1: $\text{ctr} \leftarrow 0, \text{st}_0 \leftarrow \emptyset, \text{L}_S \leftarrow \emptyset$</p> <p>2: $(\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(1^\lambda)$</p> <p>3: $(\mu^*, \sigma^*, t^*) \leftarrow \mathcal{A}^{\text{OSign, fleVer}}(\text{pk}, \lambda)$</p> <p>4: return $(\text{ctr}, \mu^*, \sigma^*, t^*)$</p> <p>OfleVer$_k(\text{st}_{\text{ctr}}, \text{pk}, \mu, \sigma, t')$</p> <hr/> <p>1: if $(\text{ctr} \geq r)$ return \perp</p> <p>2: $t \leftarrow \text{OInt}(t')$</p> <p>3: $\alpha \leftarrow \text{fleVer}(\text{st}_{\text{ctr}}, \mu, \sigma, t)$</p> <p>4: if $(\alpha \neq \perp)$</p> <p>5: $\text{ctr} \leftarrow \text{ctr} + 1$</p> <p>6: return α</p>	<p>Exp$_{\mathcal{A}, \Sigma, r}^{\text{r-fleEUf}}(\lambda, k)$</p> <hr/> <p>1: $(\text{ctr}, \mu^*, \sigma^*, t') \leftarrow \text{r-fleEUf}(\Sigma, \lambda, k)$</p> <p>2: $\beta \leftarrow \text{Ver}(\text{pk}, \mu^*, \sigma^*)$</p> <p>3: $t^* \leftarrow \text{OInt}(t')$</p> <p>4: $\alpha \leftarrow \text{fleVer}_k(\text{st}_{\text{ctr}}, \text{pk}, \mu^*, \sigma^*, t^*)$</p> <p>5: if $(\mu^* \in \text{L}_S \vee \alpha = \perp \vee \beta = 1)$</p> <p>6: return $(0, 0)$</p> <p>7: return (ctr, t^*)</p> <p>OSign$_{\text{sk}}(\mu)$</p> <hr/> <p>1: $\text{L}_S \leftarrow \text{L}_S \cup \{\mu\}$</p> <p>2: $\sigma \leftarrow \text{Sign}(\text{sk}, \mu)$</p> <p>3: return σ</p>
---	---

Fig. 8: Security model for signatures with stateful, (k, r) -efficient and flexible verification: security game, experiment and oracles.

5.1 A Generic Compiler

First of all we notice that our compiler for efficient verification described in Section 3.1 is trivially ∞ -fleEUf (i.e., existentially unforgeable against an unbounded polynomial number of verification queries), by defining $\text{fleVer} = (\text{fleVer}_0, \text{fleVer}_1)$ as shown in Figure 9. Recall that by assumption $\text{st}_0 = \emptyset$ and that for q exponential in the security parameter Theorem 1 shows that we can set $k = 1$ and have unforgeability.

We now present a compiler for signatures with **Mv**-style verification and $q = \text{poly}(\lambda)$ that realizes efficient bounded flexible verification. Our r -flefs compiler builds on top of the two compilers presented in Section 3.1 and 4.3. Intuitively, the problem with flexible verification is that if interrupted after $t < k$ steps the process may erroneously accept an invalid signature with a non-negligible probability of $1/q^t$. Moreover, from one such forgery the adversary learns that the employed \mathbf{z}_i 's satisfy a specific linear system. In particular, using the terminology introduced in the proof of Theorem 1, given one ‘bad’ forgery, the adversary has higher chance to derive a new ‘bad’ forgery. In Section 4.3 we mitigate this leakage of information between queries by refreshing the vectors in svk after every verification that outputs $\alpha > 0$. This, however, has a clear negative impact on the efficiency of our solutions. Here we want to prioritize efficiency at the cost of accuracy, and investigate how the confidence function degrades when the same set of vectors \mathbf{z}_i is used to perform r flex-verifications.

Our r -flefs compiler works essentially as the efficient verification compiler in Figure 7, except that the offVer algorithm (that generates a fresh randomness for the secret verification key svk) is run only *once every r accepting verifications*, i.e., verifications that return an outcome $\alpha \neq \perp$. To further optimize the scheme, we replace the GetZV algorithm by k algorithms GetZV_i each of which is run by the corresponding fleVer_i . The behavior of GetZV_i depends on the signature scheme and in what follows we define it for each of the three major classes we identified in this paper.

GPV08 [20]: the GetZV_i routine returns $\mathbf{z}'_i = \mathbf{z}_i = \mathbf{c}_i \mathbf{M}$, and $\mathbf{v}_i = [\sigma | \mathcal{H}(\mu)]$.

MP12 [28]: the GetZV_i routine outputs $\mathbf{z}'_i = [\tilde{\mathbf{z}}_i \mid \mathbf{z}_i^0 + \sum_{j=1}^{\ell} \mu[j] \mathbf{z}_i^j | \mathbf{c}_i]$ and $\mathbf{v}_i = [\sigma | \mathbf{u}]$.

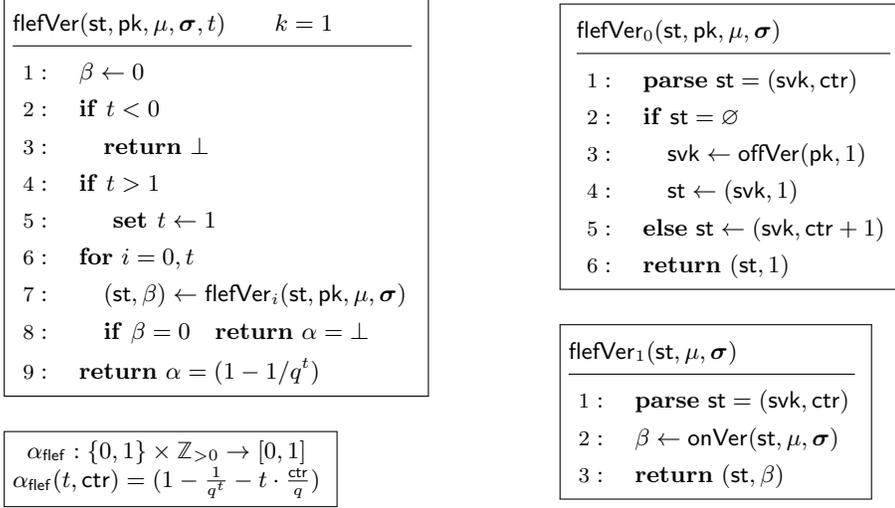


Fig. 9: The trivial two-step flefVer solution based on our compiler for efficient verification (Figure 3).

Rainbow [14]: the GetZV_i routine outputs $\mathbf{z}'_i = \mathbf{z}_i = \mathbf{c}_i \mathbf{M}$, and $\mathbf{v}_i = [\tilde{\mathbf{s}} | \mathbf{s} | \mathbf{h}]$.

Finally, the confidence function $\alpha_{\text{flef}}(\cdot, \cdot)$ is defined as:

$$\alpha_{\text{flef}}(t, \text{ctr}) = \begin{cases} \left(1 - \frac{1}{q^t} - \frac{\text{ctr}}{q}\right) & \text{if } t > 0 \\ 0 & \text{if } t = 0 \end{cases} \quad (7)$$

r₀-concrete amortized efficiency. The cost of flefVer_i varies depending whether $i = 0$ or $i > 0$. When flefVer is run the first time (or with an empty state), the step flefVer₀ generates the state. This includes computing knm multiplications, in the worst case. After that, every step flefVer_i computes only $(n + m)$ multiplications (the first term represents the cost of running GetZV_i). Therefore,

$$\text{cost}(\text{flefVer}(\text{st}_0, \mu_0, \sigma_0, k)) = knm + k(n + m) .$$

However, this is true only for the first execution of flefVer, as when executing the verification $1 < r_0 < r$ times, the algorithm flefVer₀ does not refresh the multipliers. Hence, for $i > 0$

$$\text{cost}(\text{flefVer}(\text{st}_i, \mu_i, \sigma_i, k)) = k(n + m) .$$

This yields $\sum_{i=0}^{r_0-1} \text{cost}(\text{flefVer}(\text{st}_i, \mu_i, \sigma_i, k)) = knm + r_0 k(n + m)$. The cost of a verification is dominated by $\text{cost}(\text{Ver}(\text{pk}, \mu, \sigma)) = nm$ multiplications, in the worst case. Therefore, Equation (6) yields

$$knm + r_0 k(n + m) < r_0 nm \quad \Rightarrow \quad r_0 > \frac{knm}{nm - k(n + m)} .$$

From the above formula we can derive a lower bound on values of r that yield efficiency (recall that by definition $r_0 \leq r$). A concrete security approach should lead to a meaningful upper bound on the value r that can be safely used in realistic applications. Intuitively, lower values of r yield higher accuracy (and unforgeability), higher ones guarantee better amortized efficiency.

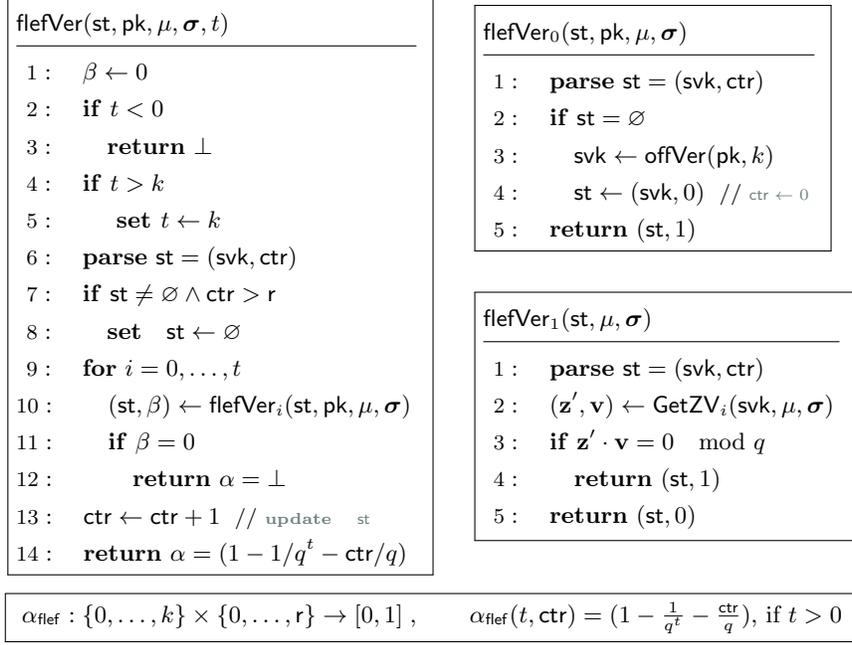


Fig. 10: Generic compiler to obtain efficient and flexible verification of signature schemes with **Mv**-style verification and q polynomial in the security parameter.

5.2 On the Concrete Security of r -flefs

The previous section seems to imply that flexibility and efficiency are in general two mutually exclusive properties of a provably secure signature scheme. This is true when considering asymptotic security definitions for lattice-based parameters that are (only) polynomial in λ . However, this does not imply that efficient flexible signatures are always insecure in practice. On the contrary, in this section we show that by limiting the amount of reuse of a secret verification key we get the best of both worlds. Concretely, we show that it is possible to simultaneously be both efficient and (arguably) accurate even for those signatures with **Mv**-style verification that rely on algebraic structures of size q polynomial in the security parameter λ . The idea is to limit the amount of verification queries the adversary is allowed to make during the security game. This simulates the fact that, after r flef verifications that return a value $\alpha > 0$, the verifier re-sets svk , i.e., the secret randomness used for the efficient verification is refreshed. We keep track of this quantity via a counter ctr , as done in the compiler presented before.

In the following we formally prove the security of our compiler and conclude with remarks on actual values.

Theorem 4. *Let Σ be a signature with **Mv**-style verification. Then our compiler for efficient and flexible verification provided in Figure 10 is r -bounded flexible secure for $r = k$.*

Proof. (sketch) The proof proceeds quite similarly to that of Theorem 2 for the case of $q \approx 2^\lambda$. Let \mathbf{G}_0 be the r -flefEUF experiment in Figure 8, and define a series of q_V hybrid games where, for every $i = 1$ to q_V , \mathbf{G}_i is the same as \mathbf{G}_{i-1} except that the i -th query to flefVer is answered using

Table 3: Concrete accuracy and efficiency values for realistic parameter choices for lattice based signatures with polynomial modulus: size of $q = 2^{30}$, $n = 256$, $k = 5$, and the worst case scenario where $t = 1$. Values are obtained by setting $\alpha_{\text{flef}}(1, r) = 1 - (r + 1)/q$ (cf. Equation (7)) to the stated value, and extracting the corresponding r from the formula. Then, r is plugged in Equation (3.3) to derive the corresponding e .

$\min \alpha_{\text{flef}}$	$1 - 10^{-2}$	$1 - 10^{-4}$	$1 - 10^{-6}$	$1 - 10^{-8}$
r	$\approx 2^{23}$	$\approx 2^{16}$	$1072 \approx 2^{10}$	$9 \approx 2^3$
e	$\approx 10^{-6}$	$\approx 10^{-4}$	$0.0047 \approx 10^{-2}$	$1.94596 \approx 10^{-0.2}$

the standard verification algorithm `Ver`. Let bad_i be the event that the adversary distinguishes \mathbf{G}_i from \mathbf{G}_{i-1} as in Equation (3). Then we have the following

$$\begin{aligned} \Pr\left[\text{Exp}_{\mathcal{A}, \Sigma}^{r\text{-flefEUf}}(\lambda) = (j^*, t^*), j^* > 0\right] &= \Pr[\mathbf{G}_0 = 1] \\ &\leq \sum_{i=1}^{j^*} \Pr[\text{bad}_i] + \Pr[\mathbf{G}_{j^*} = 1] \\ &\leq \sum_{i=1}^{j^*} \frac{1}{q^{t_i}} + \frac{1}{q^{t^*}} \leq \frac{j^*}{q} + \frac{1}{q^{t^*}} \end{aligned}$$

where t_i the interruption value used in the i -th verification, and t^* is the interruption value used in the verification of the forgery. Note that above we used $\Pr[\mathbf{G}_{j^*} = 1] \leq \frac{1}{q^{t^*}}$ since in that game the vectors $\{\mathbf{c}\}$ have never been used until the forgery’s verification, hence they can be considered fresh. Also, $\Pr[\mathbf{G}_i = 1] = 0$ for all $i > j^*$ since j^* is the index of the last verification query. Therefore, assuming $t > 0$ (otherwise, by construction $\alpha_{\text{flef}}(0) = 0$) we have:

$$\begin{aligned} \text{Adv}_{\mathcal{A}, \Sigma}^{r\text{-flefEUf}}(\lambda, k) &= \Pr\left[\text{Exp}_{\mathcal{A}, \Sigma}^{r\text{-flefEUf}}(\lambda) = (j^*, t^*), j^* > 0\right] - (1 - \alpha_{\text{flef}}(t^*, j^*)) \\ &\leq \frac{j^*}{q} + \frac{1}{q^{t^*}} - \left(1 - \left(1 - \frac{1}{q^{t^*}} - \frac{j^*}{q}\right)\right) = 0 \end{aligned}$$

□

Parameters for concrete security depend on the minimum accuracy that the verifier decides to tolerate, i.e., the minimum value of α_{flef} allowed. Such a value comes up in the worst-case scenario, in which the adversary only allows for $t = 1$ verification steps. From the description of α_{flef} given in Equation (7), the bound r on the number of “bad forgeries” (i.e., signatures whose verification with `fleVer` returns $\alpha > 0$ while with `Ver` returns 0) our compiler can tolerate. Hence, r is the maximum number of times the randomness `svk` can be reused while having a confidence level of at least α_{flef} . In Table 3 we report some values of r and $\alpha_{\text{flef}}(1)$ for different parameters of the main LBS considered thus far. We also provide the lower bound e on the amortized efficiency reached by our compiler, if run on r flef verifications. The third column of Table 3 states the following information: if we run our compiler on a lattice-based signature scheme with `Mv`-style verification and modulus $q = 2^{30}$ of polynomial size, $n = 256$ with maximum confidence level $k = 5$, we get:

- A `r-flefs` scheme with minimal accuracy 0.0001%; in other words `fleVer` may accept an invalid signature once every one million verifications, i.e., our compiler is correct 99.9999% of the times.

- For this accuracy level, we can ‘safely’ run at most $r = 1072$ flexible and efficient verifications using the same svk .
- For this accuracy level, running `fleVer` this number of repetitions leads to an efficiency ratio $e = 0.0047$, i.e., our flexible and efficient verification requires 200x less computation than running the standard signature verification about a 1000 times.

Overall this implies the verification yields at least 30 bits of security. This number might seem low when compared to asymptotic security. However, let us compare it with another real-world case, and consider a credit card with a 4 digit PIN code. To mitigate the chance for an attack to be successful the bank tolerates up to 3 unsuccessful trials. Thus, if the PIN is chosen at random, a credit card is secured at a level equal to $3/10000 > 2^{-12}$ bits of security.

Moreover, even for high accuracy levels (where we expect `fleVer` to mistakenly not recognize a forgery among one million signatures), we get astonishing efficiency gain, and can reuse the randomness a surprisingly large number of times.

6 Conclusions and Future Work

We presented a study on how to achieve efficiency and flexibility in a signature verification. In addition to putting forth these notions and formal models for them, we also presented two compilers that allow one to realize efficient (resp. flexible) verification for a wide class of existing lattice-based signatures. While our constructions show the feasibility of the desired properties, they also raise some natural follow up questions. For instance, is it possible to realize a compiler for LBS with $q \sim \text{poly}(\lambda)$ that simultaneously provides efficient *and* flexible verification? We address this question in a positive way in Section 5, albeit in weaker security model. A solution with full fledged flexible security remains an interesting open problem. Another question is, is it possible to generalize our approach to other classes of digital signatures, e.g., code-based or LBS obtained through the Fiat-Shamir heuristic or from ideal lattices? Finally, it would be worth to explore more applications of flexible and efficient verification. On top of the already mentioned applications to real-time systems, another possible venue is parallel and distributed verification of digital signatures. Consider a public bulletin board that stores authenticated (signed) data. For security reasons, one may be tempted to use post quantum signature schemes such as LBS. However, the large sizes of the public keys and signatures and the slow speed of the verification are notorious bottlenecks to deploy them in such scenarios. Using our approach, a pool of parties –acting as verifiers– can be made in charge of running each a single verification check (i.e., `flexVer` includes only `flexVer0` and `flexVer1`). In terms of security, although a single verifier may be wrong with non-negligible probability $1/q$, the probability that k honest verifiers are all wrong becomes negligible already for $k = 5$. Finally, we think that it would be interesting to explore the study of efficient and flexible verification also for more cryptographic primitives, such as commitments and zero-knowledge proofs.

Acknowledgments. This work was partly funded by: ELLIIT, the Swedish Foundation for Strategic Research grant RIT17-0035, SNSF, Swiss National Science Foundation project number 182452. Part of this work was made while C.B. was at IBM Research - Zurich (CH) and visiting the University of Aarhus (DK). This work has also received funding in part from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program under project PICOCRYPT (grant agreement No. 101001283), by the Spanish Government under projects SCUM (ref. RTI2018-102043-B-I00), CRYPTOEPIC (ref. EUR2019-103816), and SECURITAS (ref.

RED2018-102321-T), and by the Madrid Regional Government under project BLOQUES (ref. S2018/TCS-4339).

References

1. M. R. Albrecht, B. R. Curtis, A. Deo, A. Davidson, R. Player, E. W. Postlethwaite, F. Virdia, and T. Wunderer. Estimate all the lwe, ntru schemes! In *Security and Cryptography for Networks SCN*, LNCS, 2018.
2. M. Backes, D. Fiore, and R. M. Reischuk. Verifiable delegation of computation on outsourced data. In *2013 ACM SIGSAC Conference on Computer and Communications Security, CCS.*, pages 863–874. ACM, 2013.
3. M. Bellare, C. Nampreppe, D. Pointcheval, and M. Semanko. The one-more-rsa-inversion problems and the security of chaum’s blind signature scheme. *J. Cryptology*, 16(3):185–215, 2003.
4. R. Bendlin, S. Krehbiel, and C. Peikert. How to share a lattice trapdoor: Threshold protocols for signatures and (H)IBE. In *Applied Cryptography and Network Security ACNS*, LNCS, 2013.
5. W. Beullens, A. Szepieniec, F. Vercauteren, and B. Preneel. Luov: Signature scheme proposal for nist pqc project. 2019.
6. D. Boneh and D. M. Freeman. Linearly homomorphic signatures over binary fields and new tools for lattice-based signatures. In *International Workshop on Public Key Cryptography*, pages 1–16. Springer, 2011.
7. D. Boneh, B. Lynn, and H. Shacham. Short signatures from the Weil pairing. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 514–532. Springer, 2001.
8. X. Boyen. Lattice mixing and vanishing trapdoors: A framework for fully secure short signatures and more. In *International Workshop on Public Key Cryptography (PKC)*, pages 499–517. Springer, 2010.
9. D. Cash, D. Hofheinz, E. Kiltz, and C. Peikert. Bonsai trees, or how to delegate a lattice basis. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 523–552. Springer, 2010.
10. D. Catalano, D. Fiore, and B. Warinschi. Homomorphic signatures with efficient verification for polynomial functions. In *Advances in Cryptology – CRYPTO*, 2014.
11. D. Chaum. Blind signatures for untraceable payments. In *Advances in Cryptology - CRYPTO*, pages 199–203. Springer, 1983.
12. J. Coron and A. Gini. A polynomial-time algorithm for solving the hidden subset sum problem. In D. Micciancio and T. Ristenpart, editors, *Advances in Cryptology - CRYPTO 2020 - 40th Annual International Cryptology Conference, CRYPTO 2020, Santa Barbara, CA, USA, August 17-21, 2020, Proceedings, Part II*, volume 12171 of *Lecture Notes in Computer Science*, pages 3–31. Springer, 2020.
13. R. Del Pino, V. Lyubashevsky, G. Neven, and G. Seiler. Practical quantum-safe voting from lattices. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 1565–1581, 2017.
14. J. Ding, M.-S. Chen, A. Petzoldt, D. Schmidt, and B.-Y. Yang. Rainbow. Available at <https://csrc.nist.gov/projects/post-quantum-cryptography/round-2-submissions>. Accessed: 2020-09-21.
15. J. Ding and D. Schmidt. Rainbow, a new multivariable polynomial signature scheme. In *Applied Cryptography and Network Security, ACNS. Proceedings*, volume 3531 of *Lecture Notes in Computer Science*, pages 164–175, 2005.
16. D. Fiore, A. Mitrokotsa, L. Nizzardo, and E. Pagnin. Multi-key homomorphic authenticators. In *International Conference on the Theory and Application of Cryptology and Information Security - ASIACRYPT*, 2016.
17. M. Fischlin. Progressive verification: The case of message authentication. In *International Conference on Cryptology in India*, pages 416–429. Springer, 2003.
18. N. Gama and P. Q. Nguyen. Predicting lattice reduction. In *Advances in Cryptology - EUROCRYPT. Proceedings*, pages 31–51, 2008.
19. R. Gennaro, C. Gentry, and B. Parno. Non-interactive verifiable computing: Outsourcing computation to untrusted workers. In *CRYPTO*, 2010.
20. C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *ACM Symposium on Theory of Computing. Proceedings*, pages 197–206, 2008.
21. S. Gorbunov, V. Vaikuntanathan, and D. Wichs. Leveled fully homomorphic signatures from standard lattices. In *ACM symposium on Theory of computing. Proceedings*, pages 469–477. ACM, 2015.
22. S. Gorbunov and D. Vinayagamurthy. Riding on asymmetry: Efficient ABE for branching programs. In T. Iwata and J. H. Cheon, editors, *Advances in Cryptology - ASIACRYPT. Proceedings*, Lecture Notes in Computer Science, 2015.

23. S. Katsumata and S. Yamada. Group signatures without NIZK: from lattices in the standard model. In *Advances in Cryptology - EUROCRYPT*, 2019.
24. L. Lamport. Constructing digital signatures from a one-way function. Technical report, Technical Report CSL-98, SRI International, 1979.
25. D. V. Le, M. Kelkar, and A. Kate. Flexible signatures: Making authentication suitable for real-time environments. In *European Symposium on Research in Computer Security, (ESORICS)*, pages 173–193. Springer, 2019.
26. L. P. Malasinghe, N. Ramzan, and K. Dahal. Remote patient monitoring: a comprehensive study. *Journal of Ambient Intelligence and Humanized Computing*, 10(1):57–76, 2019.
27. R. C. Merkle. A certified digital signature. In *Conference on the Theory and Application of Cryptology - ASIACRYPT*, pages 218–238. Springer, 1989.
28. D. Micciancio and C. Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *EUROCRYPT*, 2012.
29. M. S. E. Mohamed and A. Petzoldt. RingRainbow - An efficient multivariate ring signature scheme. In *Progress in Cryptology - AFRICACRYPT. Proceedings*, volume 10239 of *Lecture Notes in Computer Science*, pages 3–20, 2017.
30. T. Plantard, A. Sipasseuth, C. Dumondelle, and W. Susilo. Drs: diagonal dominant reduction for lattice-based signature. In *PQC Standardization Conference*, 2018.
31. A. Sipasseuth, T. Plantard, and W. Susilo. Using freivalds’ algorithm to accelerate lattice-based signature verifications. In *International Conference on Information Security Practice and Experience*, pages 401–412. Springer, 2019.
32. A. R. Taleb and D. Vergnaud. Speeding-up verification of digital signatures. *Journal of Computer and System Sciences*, 2020.
33. S. Tayeb, M. Pirouz, G. Esguerra, K. Ghobadi, J. Huang, R. Hill, D. Lawson, S. Li, T. Zhan, J. Zhan, et al. Securing the positioning signals of autonomous vehicles. In *2017 IEEE International Conference on Big Data*, 2017.
34. R. Tsabary. An equivalence between attribute-based signatures and homomorphic signatures, and new constructions for both. In *Theory of Cryptography TCC*, 2017.