L-BAS: A Lattice-Based Blind Adaptor Signature Scheme

Amin Mohammadali[1][0000−0001−8495−7421] and Riham AlTawy[2][0000−0002−4734−3700]

University of Victoria, BC , Canada
aminmohammadali@uvic.ca raltawy@uvic.ca

Abstract. Lattice-based blind signatures have attracted significant attention in recent years due to the rapid growth of digital currencies, the increasing demand for privacy-preserving digital interactions, and the ongoing transition toward quantum-resistant cryptographic primitives. While blind signatures provide anonymity guarantees, achieving fairness without compromising privacy to a third party remains a challenging problem. Blind adaptor signatures (BAS) address this limitation by enriching blind signatures with conditional-execution functionality, enabling fair exchange while preserving user anonymity. In particular, a BAS scheme allows a user to engage in an atomic swap with a verifier using an adapted blind signature obtained from a signer, thereby maintaining privacy against the signer while ensuring fairness between the user and the verifier. In this work, we observe that the ABDLOP commit-and-prove framework (CRYPTO 2022) exhibits a dichotomic structure that can be leveraged to realize adaptor functionality. Building on this, we propose a lattice-based blind adaptor signature (L-BAS) scheme that simultaneously achieves fairness along with the privacy guarantees of blind signing. Compared with the underlying lattice-based blind signature scheme, our construction incurs only a modest overhead, increasing the signature size by approximately 5.2 KB while largely preserving the efficiency of the original system. We formally analyze the security of the proposed construction and prove that it satisfies extractability, unique extractability, computational pre-verification soundness, one-more unforgeability, and blindness under standard lattice-based assumptions. Our results demonstrate that fairness can be incorporated into lattice-based blind signatures with minimal performance degradation, making the proposed scheme a practical candidate for privacy-preserving and quantum-resistant fair exchange applications.

Keywords: Blind signature · Adaptor signature · Blind adaptor signature · Privacy · Fair exchange · Post-quantum

1 Introduction

Privacy, fairness, and post-quantum security are becoming central requirements in modern digital payment systems. In privacy-preserving payment systems,

users should be able to obtain and spend digital tokens without allowing the issuer or other parties to link the token to their identity [15]. At the same time, practical buyer-seller interactions require fair exchange (a.k.a atomic swap): a seller should not be able to receive payment without releasing the agreed digital asset, and a buyer should not be able to obtain the asset without completing the payment [41]. These requirements become even more important in decentralized and blockchain-based systems, where parties may not trust each other and where disputes are difficult to resolve through a central authority.

Blind signatures protect user privacy by allowing a signer to sign a message without learning its content, while producing a signature valid for the original message [14]. They have been instantiated under various cryptographic assumptions and recently proposed for national-scale applications such as central bank digital currencies [10, 15]. However, blind signatures alone do not ensure fairness between mutually distrustful parties. Adaptor signatures provide this complementary functionality through a pre-signature tied to an NP statement $\mathbf{M}$: a party knowing a witness $\mathbf{m}$ can adapt it into a valid signature, while doing so without $\mathbf{m}$ is infeasible. Their extractability property further allows $\mathbf{m}$ to be recovered from a pre-signature and its corresponding valid signature [24], enabling conditional execution in payment channels and atomic swaps [17, 18], coin mixing [25, 41], and oracle-based contracts [34]. In particular, when the adaptor uses the pre-signature to obtain a valid signature, the signer can recover the witness. Nevertheless, adaptor signatures alone do not hide the signed message from the signer.

Blind adaptor signatures (BAS) combine the functionalities of blind signatures and adaptor signatures into a unified primitive [28, 41]. Conceptually, a BAS scheme involves three entities: a signer, a user, and an adaptor. Through a blind signing protocol, the user obtains an blinded pre-signature from the signer on a hidden message, ensuring that the signer learns nothing about the message being signed. The user can subsequently transfer the unblinded pre-signature to the adaptor, who transforms it into a valid signature by providing a witness corresponding to a public statement. Consequently, BAS simultaneously achieves blindness, which preserves message privacy against the signer, and adaptor functionality, which makes the validity of the final signature contingent upon the release of a witness thus achieving fair exchange.

The use of BAS has been demonstrated in blockchain-based payment scenarios. Qin et al. [41] employed BAS to realize privacy-preserving conditional payments between mutually distrustful parties. However, their construction relies on a weak notion of blindness. Under weak blindness, once the signer obtains the final signature pair, it may be able to associate this pair with a particular execution of the signing protocol. This is weaker than the standard blindness notion of blind signatures, which requires such linkage to be infeasible. Consequently, although the construction provides a certain degree of privacy, it does not fully achieve the anonymity guarantees expected from conventional blind-signature schemes. This limitation motivates the development of BAS constructions satisfying a stronger blindness notion.

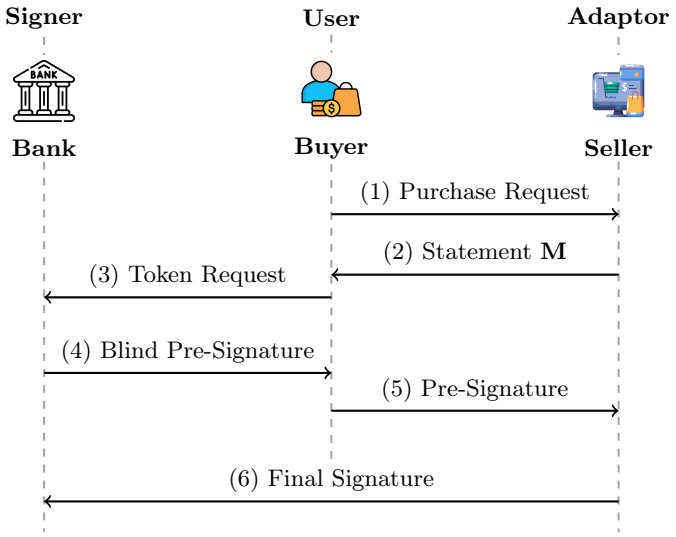


Fig. 1. High-level flow of a blind-adaptor-signature-based anonymous fair payment scenario.

In this work, we consider another important application of BAS: anonymous fair payment in an eCash-style setting. Chaum and Moser [15] introduced eCash 2.0, in which users obtain digital coins from a central bank through a blind-signature protocol. Such systems provide anonymity with respect to the issuing bank but do not inherently address fairness between a buyer and a seller. If the buyer pays first, the seller may refuse to release the promised service or digital item. Conversely, if the seller delivers first, the buyer may refuse to complete the payment. Blind adaptor signatures provide a natural mechanism for overcoming this limitation.

Consider a setting involving three parties: a bank, a buyer, and a seller, as illustrated in Figure 1. The seller generates a public statement $\mathbf{M}$ with witness $\mathbf{m}$ and sends $\mathbf{M}$ to the buyer. The buyer then requests from the bank a payment token associated with a hidden message and the statement $\mathbf{M}$. The bank returns a blind adaptor pre-signature, which the buyer unblinds and forwards to the seller. Using the witness $\mathbf{m}$ corresponding to $\mathbf{M}$, the seller transforms the adaptor pre-signature into a valid signature and redeems the resulting payment token at the bank. The seller can obtain a redeemable payment token only by revealing $\mathbf{m}$ required to complete the signature, while the buyer can recover $\mathbf{m}$ from the completed signature that is redeemed and published by the bank and thereby obtain the corresponding service authorization. At the same time, the blindness property prevents the bank from linking the redeemed token/signature to the original issuance request, thereby preserving the buyer's anonymity.

A second motivation for this work comes from the transition to post-quantum cryptography. A major contemporary shift in cryptography is the move away

from classical hardness assumptions, such as integer factorization and discrete logarithms, toward constructions believed to remain secure against quantum adversaries. The National Institute of Standards and Technology has published post-quantum public-key cryptographic standards [35–37], while the National Security Agency has established migration timelines requiring the adoption of quantum-resistant cryptography by the early 2030s [38]. Consequently, it is essential to design cryptographic primitives, including blind signatures and blind adaptor signatures, that remain secure in the presence of quantum attacks. Although numerous efficient blind signature schemes have been proposed under classical assumptions, comparatively little work has focused on quantum-secure alternatives [1, 5, 10, 12, 39]. Moreover, to the best of our knowledge, no lattice-based blind adaptor signature scheme has previously been proposed, leaving the construction of post-quantum BAS an open problem.

This leaves an important research gap. Blind signatures provide anonymity but not fairness, adaptor signatures provide conditional fairness but not blindness, and existing BAS constructions have been developed exclusively in the classical setting. Future privacy-preserving payment systems, particularly those based on eCash and blockchain-enabled conditional payments, require BAS constructions that simultaneously provide strong blindness, conditional fairness, and post-quantum security.

1.1 Related Works

Since this work presents the first post-quantum blind adaptor signature scheme, we review the existing literature on post-quantum blind signatures and adaptor signatures.

Post-quantum blind signatures. Early lattice-based blind signature schemes were proposed by Rückert [42] and Alkadri et al. [2, 3], building on ideas inspired by the Schnorr blind signature framework [43]. However, the security of this line of work was later questioned by Hauck et al. [27]. Subsequently, Lyubashevsky et al. [32] proposed a lattice-based blind signature scheme based on Gaussian one-time signatures. Although this construction was conceptually important, it suffers from a substantial efficiency limitation, with signatures of approximately 150 KB.

Agrawal et al. [1] proposed the first lattice-based blind signature scheme in the Fischlin framework [22], relying on the one-more ISIS assumption. Their construction achieves a signature size of approximately 45 KB. Later, del Pino [39] introduced the first efficient blind signature scheme based on standard lattice problems. Although this scheme improved the state of the art at the time, its signature size remains relatively large, approximately 100 KB.

Beullens et al. [10] proposed a lattice-based blind signature scheme with significantly smaller signatures than prior lattice-based constructions, achieving a signature size of approximately 22 KB. However, their scheme requires proofs of the correctness of the inputs and outputs of cryptographic hash functions. Although these proofs are excluded from the reported signature size, they introduce

substantial communication overhead, increasing the total communication cost to several hundred kilobytes. In addition, signature generation takes approximately 30 seconds. In the second scheme, Beullens et al. [10] introduced the first lattice-based keyed-verification blind signature scheme, inspired by the work of Chase et al. [13]. Although this construction, the second construction, achieves one of the smallest signature sizes among comparable schemes, its signing procedure is highly inefficient, requiring approximately 30 minutes to generate a signature, which limits its practicality.

Bootle et al. [11] later introduced a general framework for constructing anonymous credentials, blind signatures, and group signatures based on a new hardness assumption, denoted by ISIS_f. Their framework demonstrates several applications of this assumption and, like some previous works, uses the Falcon signature scheme [40] as an underlying primitive for constructing blind signatures.

Baldimtsi et al. [5] proposed the first lattice-based non-interactive blind signature scheme, building on the concept of non-interactive blind signatures introduced by Hanzlik [26]. Although their construction achieves non-interactivity, its signature size remains relatively large compared with several interactive post-quantum blind signature schemes. In particular, their signatures are approximately 68 KB, which may limit scalability in large-scale deployment scenarios.

More recently, Charles et al. [12] introduced a generic framework for constructing post-quantum blind signatures from hash-and-sign signature schemes. Their work revisits Fischlin's round-optimal framework [22], similarly to Agrawal et al. [1] and Beullens et al. [10], but moves beyond purely lattice-based instantiations. In particular, they show how to blind hash-and-sign schemes based on algebraic trapdoor functions, including multivariate schemes such as UOV [9] and MAYO [8], as well as code-based schemes such as Wave [7]. Their most efficient instantiation, based on UOV, achieves signatures ranging from 3.8 KB to 11 KB, thereby outperforming prior post-quantum blind signatures in terms of final signature size.[1]

The key technical ingredient in the framework of Charles et al. [12] is a commit-append-and-prove system, which allows the user to commit to the blinding randomness, append the signer's response, and prove the correctness of the resulting signature without revealing either the message or the hidden randomness. Despite its compact signatures, the required NIZK/SNARK proof introduces substantial communication overhead, ranging from several hundred kilobytes to a few megabytes depending on the instantiation

Post-quantum adaptor signatures. In parallel with the development of post-quantum blind signatures, post-quantum adaptor signatures have also attracted recent attention. Esgin et al. [20] introduced the first lattice-based adaptor signature scheme. Their construction is closely related to ours in terms of the adaptor-signature component. However, their scheme does not provide blind-

[1] The same framework also yields MAYO- and Wave-based instantiations, with signature sizes up to approximately 18 KB and 138 KB, respectively. Our comparison considers their most efficient UOV-based variant.

Table 1. Comparison of post-quantum blind signature schemes

Scheme	Assumption	Signature size	Adapt. fn.
Lyubashevsky et al. [32]	M-LWE + M-SIS	150 KB	No
del Pino [39]	SIS + LWE	100 KB	No
Baldimtsi et al. [5]	rOM-ISIS	68 KB	No
Agrawal et al. [1]	OM ISIS + LWE + NTRU	45 KB	No
Beullens et al. [10]	SIS + LWE + NTRU	22 KB	No
Charles et al. [12]	UOV	3.8–11 KB	No
L-BAS	SIS + LWE + NTRU	27.4* KB	Yes

* : The communication cost of signature generation increases by 14 KB compared with the original blind signature scheme.

ness and therefore cannot be directly used in applications requiring signer-side privacy of the message. In contrast, our construction combines adaptor functionality with blindness, yielding a blind adaptor signature scheme. In another direction, Tairi et al. [44] proposed the first efficient isogeny-based adaptor signature scheme.

1.2 Our Contribution

In this work, we take a first step toward post-quantum blind adaptor signatures. We begin by formalizing the security requirements of blind adaptor signatures under different application scenarios and by identifying the properties necessary to achieve both strong blindness and adaptor functionality. We then show that the ABDLOP commitment scheme [31] possesses a dichotomic structure that can be exploited to construct adaptor mechanisms in both blind and non-blind signature settings.

As a concrete realization of this framework, we transform the first lattice-based blind signature scheme construction of Beullens et al. [10] into a blind adaptor signature scheme (L-BAS). To the best of our knowledge, this results is the first lattice-based blind adaptor signature construction and the first BAS scheme based on post-quantum assumptions. In contrast to the existing BAS construction of Qin et al. [41], L-BAS achieves strong blindness in the standard sense of blind signatures while preserving the adaptor-signature properties necessary for fair exchange and conditional execution. Furthermore, the proposed transformation preserves practicality, introducing only a modest increase in signature size compared to the underlying blind signature scheme. Finally, we prove that the resulting construction satisfies all required security properties under standard lattice assumptions. Table 1 compares L-BAS scheme with existing post-quantum blind signature constructions.

2 Preliminaries

In this section, we introduce the notation and mathematical preliminaries used throughout the paper. Additional background material is provided in Appendix A.

2.1 Notation

We denote by $x \leftarrow S$ the uniform sampling of x from a finite set S, and by $x \leftarrow D$ the sampling of x according to a distribution D. Let $\mathbb{Z}_q$ denote the ring of integers modulo q.

Let d be a power of two and q be a positive integer. We define the polynomial rings $\mathcal{R} = \mathbb{Z}[X]/(X^d + 1), \mathcal{R}_q = \mathbb{Z}_q[X]/(X^d + 1)$. Elements of $\mathcal{R}$ and $\mathcal{R}_q$ are denoted by lowercase letters. Bold lowercase letters and bold uppercase letters denote vectors and matrices over these rings, respectively.

The ring $\mathcal{R}$ admits an automorphism group $\mathrm{Aut}(\mathcal{R})$. For an automorphism $\sigma^i \in \mathrm{Aut}(\mathcal{R})$, its action on X is defined as $\sigma^i(X) = X^i$. We extend this notation componentwise: for any vector $\mathbf{s} = (s_1, \ldots, s_k) \in \mathcal{R}^k$, we define $\sigma^i(\mathbf{s}) = (\sigma^i(s_1), \ldots, \sigma^i(s_k))$, and similarly for matrices, where $\sigma^i(\mathbf{A})$ denotes the entry-wise application of σ^i.

For a polynomial $w = w_0 + w_1 X + \cdots + w_{d-1} X^{d-1} \in \mathcal{R}$, we define the norms

$$\|w\|_\infty = \max_j |w_j|, \qquad \|w\|_t = \left(\sum_{j=0}^{d-1} |w_j|^t \right)^{1/t}.$$

For a vector $\mathbf{w} = (w_1, \ldots, w_k) \in \mathcal{R}^k$, we define

$$\|\mathbf{w}\|_\infty = \max_j \|w_j\|_\infty, \qquad \|\mathbf{w}\|_t = \left(\sum_{j=1}^{k} \|w_j\|_t^t \right)^{1/t},$$

and we write $\|\mathbf{w}\| = \|\mathbf{w}\|_2$ by default. Norms over $\mathbb{Z}_q$ are defined analogously.

Finally, for a bound z, we define $S_z = \{x \in \mathcal{R}_q : \|x\|_\infty \leq z\}$, the set of ring elements with infinity norm at most z.

2.2 M-SIS and M-LWE Problems

The security of our construction relies on two standard hard problems over module lattices: the Module Short Integer Solution (M-SIS) problem and the Module Learning With Errors (M-LWE) problem [29].

Definition 1 ($\mathbf{M\text{-}SIS}_{n,m,\beta}$). *Let $\mathbf{A} \leftarrow \mathcal{R}_q^{n \times m}$. The M-SIS problem is, given $\mathbf{A}$, to find a non-zero vector $\mathbf{z} \in \mathcal{R}_q^m$ such that $\mathbf{Az} = 0$ over $\mathcal{R}_q$ and $0 < \|\mathbf{z}\| \leq \beta$. An algorithm Adv is said to solve M-SIS$_{n,m,\beta}$ with advantage ϵ if*

$$\Pr\left[\mathbf{Az} = 0 \ \wedge \ 0 < \|\mathbf{z}\| \leq \beta : \mathbf{A} \leftarrow \mathcal{R}_q^{n \times m}, \ \mathbf{z} \leftarrow \mathsf{Adv}(\mathbf{A})\right] \geq \epsilon.$$

Definition 2 ($\mathbf{M\text{-}LWE}_{n,m,\chi}$). *Let $n, m > 0$ and let χ be an error distribution over $\mathcal{R}$. The M-LWE problem is to distinguish between the distributions $(\mathbf{A}, \mathbf{As}+\mathbf{e})$ and $(\mathbf{A}, \mathbf{b})$, where $\mathbf{A} \leftarrow \mathcal{R}_q^{n \times m}$, $\mathbf{s} \leftarrow \chi^m$, $\mathbf{e} \leftarrow \chi^n$, and $\mathbf{b} \leftarrow \mathcal{R}_q^n$. We say that an adversary Adv solves M-LWE$_{n,m,\chi}$ with advantage ϵ if*

$$|\Pr[\mathsf{Adv}(\mathbf{A}, \mathbf{As} + \mathbf{e}) = 1] - \Pr[\mathsf{Adv}(\mathbf{A}, \mathbf{b}) = 1]| \geq \epsilon.$$

Algorithm 1 Rejection Sampling Steps

$\mathrm{Rej}_1(\mathbf{z}, \mathbf{v}, \mathcal{S})$

1: $u \leftarrow [0,1)$ **if** $u > \frac{1}{L} \cdot \exp\left(-\frac{2\langle z, v\rangle + \|v\|^2}{2s^2}\right)$ **then**
2: **return** 1 (Reject) **else**
3: **return** 0 (Accept)

$\mathrm{Rej}_2(\mathbf{z}, \mathbf{v}, \mathcal{S})$

1: $u \leftarrow [0,1)$ **if** $u > \frac{1}{L \exp\left(-\frac{\|\overrightarrow{v}\|^2}{2\mathcal{S}^2}\right) \cosh\left(\frac{<\overrightarrow{z}, \overrightarrow{v}>}{\mathcal{S}^2}\right)}$ **then**
2: **return** 1 (Reject) **else**
3: **return** 0 (Accept)

2.3 Rejection sampling

Rejection sampling is employed to ensure that the signer's response vector $\mathbf{z}$ leaks no information about the secret or randomness vector $\mathbf{r}$, where $\mathbf{z} = \mathbf{y} + c\mathbf{r}$ and $\mathbf{y}$ is a masking vector. Due to different requirements, we use three distinct rejection sampling procedures. Let $V \subset \mathcal{R}^\ell$ be a set of polynomials with norm at most T, and let $\rho : V \to [0,1]$ be a probability distribution. Fix a standard deviation $\mathcal{S} = \gamma T$.

- $\mathrm{Rej}_1(\mathbf{z}, \mathbf{v}, \mathcal{S})$ [30]: Define $L = \exp\left(\frac{12}{\gamma} + \frac{1}{2\gamma^2}\right)$. Sample $\mathbf{v} \leftarrow \rho$ and $\mathbf{y} \leftarrow D^\ell_{\mathcal{S}}$, and set $\mathbf{z} = \mathbf{y} + \mathbf{v}$. Run the rejection sampling procedure $b \leftarrow \mathrm{Rej}_1(\mathbf{z}, \mathbf{v}, \mathcal{S})$ as defined in the Algorithm 1. Then, The probability that $b = 0$ satisfies $\Pr[b = 0] \geq \frac{1-2^{-100}}{L}$. The joint distribution $(\mathbf{v}, \mathbf{z})$ is within statistical distance at most 2^{-100} from the product distribution $\rho \times D^\ell_{\mathcal{S}}$.
- $\mathrm{Rej}_2(\mathbf{z}, \mathbf{v}, \mathcal{S})$ [33]: It is a variant of the previous rejection sampling method. In this version, the signer computes the inner product $\langle \overrightarrow{v}, \overrightarrow{z} \rangle$ at first, and rejects the sample if it is negative. The rest of the procedure is identical to the previous method, except that $L = \exp\big(1/(2\gamma^2)\big)$. Sample $\mathbf{v} \leftarrow \rho$, $\varsigma \in \{0,1\}$, and $\mathbf{y} \leftarrow D^\ell_{\mathcal{S}}$, and set $\mathbf{z} = \mathbf{y} + (-1)^\varsigma \mathbf{v}$. Run the rejection sampling procedure $b \leftarrow \mathrm{Rej}_2(\mathbf{z}, \mathbf{v}, \mathcal{S})$ as defined below. The probability that $b = 0$ satisfies $\Pr[b = 0] \geq \frac{1}{L}$. Moreover, the joint distribution $(\mathbf{v}, \mathbf{z})$ is identical to the product distribution $\rho \times D^\ell_{\mathcal{S}}$.

3 Technical Overview

We first recall the main idea behind the blind signature scheme of Beullens et al. [10] and then explain how its underlying proof system can be exploited to obtain a blind adaptor signature.

The blind signature scheme of Beullens et al. [10] is built from a GPV/Falcon-style lattice signature [23,40]. The signer possesses a trapdoor for a public matrix $\mathbf{A}$, which allows it to sample a short preimage $\mathbf{s}$ for a target vector $\mathbf{c}$ satisfying $\mathbf{As} = \mathbf{c}$. To blind a message μ, the user samples a short vector $\mathbf{r}$ and computes

$$\mathbf{c} = \mathbf{Dr} + H(\mu, G(\mathbf{r})), \tag{1}$$

where $\mathbf{D}$ is a public random matrix and $G(\cdot), H(\cdot)$ are modeled as random oracles. The user sends ciphertexts encrypting μ and $\mathbf{r}$, together with $\mathbf{c}$ and a NIZK proof of knowledge of a short vector $\mathbf{r}$ and a message μ satisfying the above relation. After verifying the proof, the signer uses its trapdoor to generate a short vector $\mathbf{s}$ satisfying $\mathbf{As} = \mathbf{c}$ and returns $\mathbf{s}$ to the user.

The user then considers $\rho = G(\mathbf{r})$ and outputs the signature (ρ, π), where π is a proof of knowledge of short vectors $\mathbf{r}$ and $\mathbf{s}$ satisfying $\mathbf{As} = \mathbf{Dr} + H(\mu, \rho)$. Verification consists of checking the proof π for the public statement

$$\exists, \mathbf{r}, \mathbf{s}; \mathbf{As} = \mathbf{Dr} + H(\mu, \rho), \quad |\mathbf{r}| \leq \beta_r, \quad |\mathbf{s}| \leq \beta_s.$$

Our starting point is to show that the ABDLOP proof system [31] underlying π admits a dichotomic structure in the sense of Gerhart et al. [24]. Informally, a proof system is dichotomic if its output can be decomposed into two components (σ_1, σ_2), where σ_1 depends only on a hidden image of the prover's randomness, while σ_2 depends on the underlying randomness itself through a suitable homomorphic relation. Gerhart et al. [24] showed that such a structure is sufficient for constructing adaptor signatures.

More concretely, let $\mathbf{u}$ denote the randomness used during proof generation and let $\mathrm{OWF}(\mathbf{u}) = \mathbf{Au}$ be the associated homomorphic one-way function. The ABDLOP proof can be viewed as consisting of a first component that depends only on the image $\mathbf{Au}$ and a second component that depends on the underlying randomness $\mathbf{u}$. Furthermore, verification depends on the second component only through its image under the one-way function. Consequently, the proof satisfies the structural requirements of the dichotomic framework.

This observation enables us to incorporate adaptor functionality directly into the proof system. Let $\mathbf{M} = \mathbf{A}'\mathbf{m}$ be an adaptor statement with $\rho' = G(\mathbf{M})$, where $\mathbf{m}$ is the corresponding witness and $\mathbf{A}'$ is a polynomial random matrix. Instead of generating the proof with respect to the original relation (Eq. (1)), we modify the target equation by embedding the public statement ρ': $\mathbf{c} = \mathbf{Dr} + H(\mu, \rho, \rho')$. After receiving the response $\mathbf{s}$ from the signer, the user generates a proof for the modified relation

$$\exists, \mathbf{r}, \mathbf{s}; \mathbf{As} = \mathbf{Dr} + H(\mu, \rho, \rho'), \quad |\mathbf{r}| \leq \beta_r, \quad |\mathbf{s}| \leq \beta_s.$$

The resulting object constitutes a blind adaptor pre-signature. A party possessing the witness $\mathbf{m}$ can adapt the proof to obtain a valid signature, while knowledge of both the pre-signature and the final signature enables extraction of $\mathbf{m}$. Thus, the standard adaptability and extractability properties follow directly from the dichotomic structure of the ABDLOP proof system.

Importantly, the adaptor mechanism is incorporated entirely within the existing ABDLOP proof relation. In this sense, our construction adds fairness guarantees to the original scheme with slightly increasing its signature size.

4 Blind Adaptor Signature Syntax and Security

BAS combines the privacy guarantees of blind signatures with the conditional-execution functionality of adaptor signatures. Informally, a user obtains an adap-

tor pre-signature on a hidden message through a blind signing protocol, ensuring that the signer learns neither the message nor its associated statement. The resulting pre-signature can subsequently be transformed into a valid signature by an adaptor possessing a witness for a public statement. This simultaneously provides blindness and conditional fairness.

Let Rel be an efficiently computable relation and define the associated language $L_{\text{Rel}} = \{M \mid \exists m \text{ such that } (M, m) \in \text{Rel}\}$. A blind adaptor signature scheme consists of seven polynomial-time algorithms BAS.KGen, BAS.BpSign, BAS.pSign, BAS.pVrfy, BAS.Adapt, BAS.Extract, and BAS.Vrfy corresponding to key generation, blind pre-signing, unblinding the blinded pre-signature, pre-verification, adaptation, extraction, and verification respectively.

- Key Generation. $(sk, pk) \leftarrow \text{BAS.KGen}(1^\lambda)$. A probabilistic algorithm that takes the security parameter 1^λ as input and outputs a signing key sk and a public key pk.
- Blind Pre-signing. $\tilde{\sigma}' \leftarrow \text{BAS.BpSign}(sk, \mu')$. A probabilistic algorithm that takes as input the signing key sk, a blinded message μ and a statement $M \in L_{\text{Rel}}$ denoted as μ'. It outputs a blinded adaptor pre-signature $\tilde{\sigma}'$.
- Pre-signing. $\tilde{\sigma} \leftarrow \text{BAS.pSign}(\tilde{\sigma}', \mu, M)$. A probabilistic algorithm that removes the blinding from the blinded adaptor pre-signature $\tilde{\sigma}'$ and outputs a valid adaptor pre-signature $\tilde{\sigma}$ on message μ with respect to statement M.
- Pre-verification. $b \leftarrow \text{BAS.pVrfy}(pk, \mu, \tilde{\sigma}, M)$. A deterministic algorithm that takes as input the public key pk, a message μ, a pre-signature $\tilde{\sigma}$, and a statement M. It outputs a bit $b \in \{0, 1\}$, where $b = 1$ indicates that $\tilde{\sigma}$ is a valid pre-signature on μ with respect to M.
- Adaptation. $\sigma \leftarrow \text{BAS.Adapt}(pk, \tilde{\sigma}, m)$. A probabilistic algorithm that takes as input the public key pk, a pre-signature $\tilde{\sigma}$, and a witness m satisfying $(M, m) \in$ Rel. It outputs a signature σ.
- Extraction. $m \leftarrow \text{BAS.Extract}(pk, \tilde{\sigma}, \sigma, M)$. A deterministic algorithm that takes as input the public key pk, a pre-signature $\tilde{\sigma}$, a valid signature σ, and a statement M. It outputs either a witness m satisfying $(M, m) \in$ Rel or the failure symbol $\perp$.
- Verification. $b \leftarrow \text{BAS.Vrfy}(pk, \mu, \sigma)$. A deterministic algorithm that takes as input the public key pk, a message μ, and a signature σ. It outputs a bit $b \in \{0, 1\}$, where $b = 1$ indicates that σ is a valid signature on message μ under the public key pk, and $b = 0$ otherwise.

The security requirements of blind signatures and adaptor signatures are reviewed in Appendix A.1 and Appendix A.2, respectively. A BAS scheme should simultaneously inherit the privacy guarantees of blind signatures and the fairness guarantees of adaptor signatures. Formally, a BAS scheme is secure if it satisfies the following properties: *strong blindness* which implies that the signer cannot associate a valid message/signature pair with the execution of the blind signing protocol from which it originated, *one-more unforgeability* which ensures that no polynomial-time adversary can produce more valid signatures than the number of valid pre-signatures it received from the signer, *extractability* which implies that given a valid pre-signature and its corresponding adapted signature, a valid

witness can be efficiently extracted, *unique extractability* which ensures that each valid pre-signature can be adapted into at most one valid signature, and *pre-verification soundness* which states that only valid statements can produce pre-signatures that are verified by the pre-verification algorithm.

5 ABDLOP Commit-and-Prove Protocol and Its Dichotomic Structure

In this section, we show how the ABDLOP commitment scheme and its associated zero-knowledge proof system can be used to obtain a dichotomic structure that enables adaptor functionality. We first briefly recall the commit-and-prove protocol of the ABDLOP construction [31], and then describe its dichotomic structure.

5.1 ABDLOP Commitment Scheme

Let $(\mathbf{t}_A, \mathbf{t}_B)$ be a commitment to a secret pair $(\mathbf{s}_1, \mathbf{m}) \in \mathcal{R}^{m_1 \times \ell}$ using randomness $\mathbf{s}_2 \leftarrow S_\beta^{m_2}$, defined as

$$\begin{aligned} \mathbf{t}_A &:= \mathbf{A}_1 \mathbf{s}_1 + \mathbf{A}_2 \mathbf{s}_2 \\ \mathbf{t}_B &:= \mathbf{B}\mathbf{s}_2 + \mathbf{m}, \end{aligned}$$

where $\mathbf{A}_1 \in \mathcal{R}_q^{n \times m_1}$, $\mathbf{A}_2 \in \mathcal{R}_q^{n \times m_2}$, and $\mathbf{B} \in \mathcal{R}_q^{\ell \times m_2}$ are uniformly random matrices. Assume the prover wants to prove knowledge of the structured witness

$$\hat{\mathbf{s}} := \begin{bmatrix} (\sigma^i(\mathbf{s}_1))_{i \in [k]} \\ (\sigma^i(\mathbf{m}))_{i \in [k]} \end{bmatrix} \in \mathcal{R}_q^{k(m_1+\ell)},$$

satisfying a quadratic relation $f(\hat{\mathbf{s}}) = 0$, where f is defined as

$$f(\hat{\mathbf{s}}) := \hat{\mathbf{s}}^\mathsf{T} \mathbf{R}_2 \hat{\mathbf{s}} + \mathbf{r}_1^\mathsf{T} \hat{\mathbf{s}} + r_0, \tag{2}$$

for known $r_0 \in \mathcal{R}_q$, $\mathbf{r}_1 \in \mathcal{R}_q^{k(m_1+\ell)}$, and $\mathbf{R}_2 \in \mathcal{R}_q^{k(m_1+\ell) \times k(m_1+\ell)}$. The prover first samples masking vectors $\mathbf{y}_1 \leftarrow D_{S_1}^{m_1}$ and $\mathbf{y}_2 \leftarrow D_{S_2}^{m_2}$, and computes $\mathbf{w} = \mathbf{A}_1 \mathbf{y}_1 + \mathbf{A}_2 \mathbf{y}_2$. It then computes $g_1 = \hat{\mathbf{s}}^\mathsf{T} \mathbf{R}_2 \mathbf{y} + \mathbf{y}^\mathsf{T} \mathbf{R}_2 \hat{\mathbf{s}} + \mathbf{r}_1^\mathsf{T} \mathbf{y}$, where

$$\mathbf{y} := \begin{bmatrix} (\sigma^i(\mathbf{y}_1))_{i \in [k]} \\ -(\sigma^i(\mathbf{B}\mathbf{y}_2))_{i \in [k]} \end{bmatrix} \in \mathcal{R}_q^{k(m_1+\ell)}.$$

Next, the prover computes $t = \mathbf{b}^\mathsf{T} \mathbf{s}_2 + g_1$, where $\mathbf{b} \in \mathcal{R}^{m_2}$ is a public vector, and uses it to mask g_1. It also computes $v = \mathbf{y}^\mathsf{T} \mathbf{R}_2 \mathbf{y} + \mathbf{b}^\mathsf{T} \mathbf{y}_2$ and sends $(\mathbf{w}, t, v)$ to the verifier.

Upon receiving a challenge $c \leftarrow \mathcal{C}$ (or deriving it via the Fiat–Shamir transform [21]), the prover computes $\mathbf{z}_i = c\mathbf{s}_i + \mathbf{y}_i$ for $i \in \{1, 2\}$ and applies rejection sampling. The output $(\mathbf{z}_1, \mathbf{z}_2)$ is accepted only if it passes the rejection sampling

procedure. For $\mathbf{z}_1$, rejection sampling is performed via $\mathrm{Rej}_1(\mathbf{z}_1, c\mathbf{s}_1, \mathcal{S}_1)$, and for $\mathbf{z}_2$ via $\mathrm{Rej}_2(\mathbf{z}_2, c\mathbf{s}_2, \mathcal{S}_2)$.

Given $(\mathbf{w}, t, v, c, \mathbf{z}_1, \mathbf{z}_2)$, the verifier first checks that $\mathbf{z}_1$ and $\mathbf{z}_2$ satisfy the required norm bounds. It then reconstructs $\mathbf{z}_m = c\mathbf{t}_B - \mathbf{B}\mathbf{z}_2$ and defines

$$\mathbf{z} := \begin{bmatrix} (\sigma^i(\mathbf{z}_1))_{i\in[k]} \\ (\sigma^i(\mathbf{z}_m))_{i\in[k]} \end{bmatrix} = c\hat{\mathbf{s}} + \mathbf{y}. \tag{3}$$

Finally, the verifier checks the following equations:

$$\mathbf{A}_1\mathbf{z}_1 + \mathbf{A}_2\mathbf{z}_2 \stackrel{?}{=} \mathbf{w} + c\mathbf{t}_A, \tag{4}$$

$$\mathbf{z}^\top\mathbf{R}_2\mathbf{z} + c\mathbf{r}_1^\top\mathbf{z} + c^2 r_0 - ct + \mathbf{b}^\top\mathbf{z}_2 \stackrel{?}{=} v. \tag{5}$$

This framework enables proving a wide class of relations over the committed values $\mathbf{s}_1$ and $\mathbf{m}$. In particular, it supports quadratic relations over $\mathcal{R}_q$ and $\mathbb{Z}_q$, norm-bounded relations with respect to the infinity and Euclidean norms, and approximate shortness constraints. The specific forms of $\mathbf{R}_2$, $\mathbf{r}_1$, and r_0 depend on the relation being encoded.

5.2 Dichotomic Structure of ABDLOP

A dichotomic signature scheme is a signature scheme in which each signature can be decomposed as $\sigma = (\sigma_1, \sigma_2)$. The first component is computed from a hidden image of the signing randomness, while the second component is computed from the original randomness and satisfies a suitable homomorphic relation. More precisely, let the signing algorithm use secret key sk, message m, and randomness $\mathbf{u}$. The randomness is first mapped through a homomorphic one-way function $\mathrm{OWF} : \mathcal{D}_R \to \mathcal{D}_{R'}$. In our lattice setting, this function is instantiated as $\mathrm{OWF}(\mathbf{u}) = \mathbf{A}\mathbf{u}$, where $\mathbf{u}$ is a polynomial vector.

The signing algorithm can then be written as $\Sigma.\mathrm{Sign}(\mathrm{sk}, m; \mathbf{u}) = (\sigma_1, \sigma_2)$, where $\sigma_1 = \Sigma_1(\mathrm{sk}, m; \mathbf{A}\mathbf{u})$ and $\sigma_2 = \Sigma_2(\mathrm{sk}, m, \sigma_1; \mathbf{u})$. Thus, the first component depends on the hidden randomness $\mathbf{A}\mathbf{u}$, while the second component depends on the original randomness $\mathbf{u}$.

The verification algorithm is compatible with this decomposition. In particular, verification depends on the second component only through its one-way image. That is, there exists a deterministic algorithm Vrfy' such that

$$\Sigma.\mathrm{Vrfy}(\mathrm{vk}, m, (\sigma_1, \sigma_2)) = \mathrm{Vrfy}'(\mathrm{vk}, m, \sigma_1, \mathbf{A}\sigma_2).$$

This property is essential for adaptor signatures, because it allows the verifier to check consistency with the image of an adapted second component without knowing the witness itself.

The second component must also be homomorphic with respect to the signing randomness. In our construction, for every admissible vector $\mathbf{m}$, we have $\Sigma_2(\mathrm{sk}, m, \sigma_1; \mathbf{u}) + \mathbf{m} = \Sigma_2(\mathrm{sk}, m, \sigma_1; \mathbf{u} + \mathbf{m})$. Hence, adding $\mathbf{m}$ to the second component has the same effect as adding $\mathbf{m}$ directly to the signing randomness.

This structure enables the adaptor mechanism. Let $M = \mathbf{Am}$ be the adaptor statement, where $\mathbf{m}$ is the corresponding witness. Instead of computing the first component from $\mathbf{Au}$, the signer computes it from the shifted hidden randomness $\mathbf{Au} + \mathbf{M} = \mathbf{A}(\mathbf{u} + \mathbf{m})$. The resulting pre-signature therefore corresponds to the hidden randomness of $\mathbf{u} + \mathbf{m}$, although the witness $\mathbf{m}$ is not yet revealed. Later, any party knowing $\mathbf{m}$ can complete the pre-signature by setting $\sigma_2 = \widetilde{\sigma}_2 + \mathbf{m}$. By the homomorphic property of the second component, the resulting pair (σ_1, σ_2) is a valid ordinary signature. Moreover, given both the pre-signature $\widetilde{\sigma}$ and the final signature σ, the witness can be extracted as $\mathbf{m} = \sigma_2 - \widetilde{\sigma}_2$.

The ABDLOP commitment admits to a dichotomic structure because its two commitment components play different roles. The first component, $\mathbf{t}_A = \mathbf{A}_1\mathbf{s}_1 + \mathbf{A}_2\mathbf{s}_2$, is tied to the short witness $\mathbf{s}_1$ and is verified through the linear relation involving $\mathbf{z}_1$ and $\mathbf{z}_2$. The second component, $\mathbf{t}_B = \mathbf{Bs}_2 + \mathbf{m}$, contains an additive witness parameter $\mathbf{m}$. This additive structure enables adaptor functionality by allowing the witness component $\mathbf{m}$ to be omitted during pre-signing and inserted later during adaptation.

Let $\mathbf{m}$ be the adaptor witness and let $\mathbf{M}$ be the corresponding public statement satisfying $\mathbf{M} = \mathbf{A}'\mathbf{m}$, for a public matrix $\mathbf{A}' \in \mathcal{R}_q^{n \times m}$. The adaptor reveals $\mathbf{M}$ but keeps $\mathbf{m}$ secret. Consequently, the prover constructs the truncated commitment

$$\mathbf{t}_A = \mathbf{A}_1\mathbf{s}_1 + \mathbf{A}_2\mathbf{s}_2, \quad \text{and} \quad \mathbf{t}'_B = \mathbf{Bs}_2.$$

The prover must nevertheless generate a valid ABDLOP proof with respect to the truncated commitment. Since the adaptor relation is linear, the witness $\mathbf{m}$ appears only in the linear part of Eq. (2). More precisely, for the relation $\mathbf{M} = \mathbf{A}'\mathbf{m}$, let the witness be partitioned as $\hat{\mathbf{s}} = \begin{bmatrix} \mathbf{s}_1 \\ \mathbf{m} \end{bmatrix}$, where $\hat{\mathbf{s}}_1$ contains all witness components other than the adaptor witness. Then the quadratic relation in Eq. (2) can be expressed as

$$\mathbf{R}_2 := \begin{bmatrix} \mathbf{R}_{\mathbf{s}_1} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}, \qquad \mathbf{r}_1^{\mathsf{T}} := \left[\mathbf{r}_{1_{\mathbf{s}_1}} \ \textstyle\sum_{i=1}^{n} \mathbf{a}_i \right], \qquad r_0 := r_{0_{\mathbf{s}_1}} - \sum_{i=1}^{n} M_i,$$

where $\mathbf{a}_i$ denotes the i-th row of $\mathbf{A}'$ and M_i denotes the i-th component of $\mathbf{M}$. Therefore, the computation of $g_1 = \hat{\mathbf{s}}^{\mathsf{T}}\mathbf{R}_2\mathbf{y} + \mathbf{y}^{\mathsf{T}}\mathbf{R}_2\hat{\mathbf{s}} + \mathbf{r}_1^{\mathsf{T}}\mathbf{y}$ does not require knowledge of $\mathbf{m}$. Indeed,

$$\hat{\mathbf{s}}^{\mathsf{T}}\mathbf{R}_2\mathbf{y} = \left[\mathbf{s}_1^{\mathsf{T}} \ \mathbf{m}^{\mathsf{T}} \right] \begin{bmatrix} \mathbf{R}_{\mathbf{s}_1} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \mathbf{y} = \mathbf{s}_1^{\mathsf{T}}\mathbf{R}_{\mathbf{s}_1}\mathbf{y},$$

and similarly $\mathbf{y}^{\mathsf{T}}\mathbf{R}_2\hat{\mathbf{s}} = \mathbf{y}^{\mathsf{T}}\mathbf{R}_{\mathbf{s}_1}\mathbf{s}_1$. Hence, all quadratic terms involving $\mathbf{m}$ vanish. As a result, the prover can compute g_1, and therefore $(\mathbf{w}, t, v)$, using only the known witness components. The first ABDLOP verification equation remains unchanged, since it depends only on $\mathbf{t}_A$. In particular,

$$\mathbf{A}_1\mathbf{z}_1 + \mathbf{A}_2\mathbf{z}_2 = \mathbf{A}_1\mathbf{y}_1 + \mathbf{A}_2\mathbf{y}_2 + c(\mathbf{A}_1\mathbf{s}_1 + \mathbf{A}_2\mathbf{s}_2) = \mathbf{w} + c\mathbf{t}_A.$$

The difference appears in the reconstruction of the $\mathbf{m}$-side response. Using the truncated commitment, the verifier computes $\widehat{\mathbf{z}}_m = c\mathbf{t}'_B - \mathbf{B}\mathbf{z}_2$. Substituting $\mathbf{t}'_B = \mathbf{B}\mathbf{s}_2$ and $\mathbf{z}_2 = \mathbf{y}_2 + c\mathbf{s}_2$ yields

$$\widehat{\mathbf{z}}_m = c\mathbf{B}\mathbf{s}_2 - \mathbf{B}(\mathbf{y}_2 + c\mathbf{s}_2) = -\mathbf{B}\mathbf{y}_2.$$

The corresponding response vector is

$$\widehat{\mathbf{z}} = \begin{bmatrix} (\sigma^i(\mathbf{z}_1))_{i\in[k]} \\ (\sigma^i(\widehat{\mathbf{z}}_m))_{i\in[k]} \end{bmatrix}.$$

Since the witness component $\mathbf{m}$ is omitted, its contribution is replaced by the public statement $\mathbf{M}$. The pre-verifier then checks $\mathbf{A}_1\mathbf{z}_1 + \mathbf{A}_2\mathbf{z}_2 \stackrel{?}{=} \mathbf{w} + c\mathbf{t}_A$ and

$$\widehat{\mathbf{z}}^\mathsf{T}\mathbf{R}_2\widehat{\mathbf{z}} + c\mathbf{r}_1^\mathsf{T}\widehat{\mathbf{z}} + c^2 r_0 - ct + \mathbf{b}^\mathsf{T}\mathbf{z}_2 + c^2\sum_i^n M_i \stackrel{?}{=} v. \tag{6}$$

The additional term $c^2\sum_i^n M_i$ corresponds to the contribution that would have been generated by the missing witness component through the linear relation $\mathbf{M} = \mathbf{A}'\mathbf{m}$.

During adaptation, the adaptor inserts the witness by computing $\mathbf{t}_B = \mathbf{t}'_B + \mathbf{m}$. The verifier then reconstructs

$$\mathbf{z}_m = c\mathbf{t}_B - \mathbf{B}\mathbf{z}_2 = c(\mathbf{B}\mathbf{s}_2 + \mathbf{m}) - \mathbf{B}(\mathbf{y}_2 + c\mathbf{s}_2) = c\mathbf{m} - \mathbf{B}\mathbf{y}_2,$$

and forms $\mathbf{z}$ as Eq. 3. At this point the witness is fully restored, and the signature is verified using the original ABDLOP equations

$$\begin{aligned} \mathbf{A}_1\mathbf{z}_1 + \mathbf{A}_2\mathbf{z}_2 &\stackrel{?}{=} \mathbf{w} + c\mathbf{t}_A, \\ \mathbf{z}^\mathsf{T}\mathbf{R}_2\mathbf{z} + c\mathbf{r}_1^\mathsf{T}\mathbf{z} + c^2 r_0 - ct + \mathbf{b}^\mathsf{T}\mathbf{z}_2 &\stackrel{?}{=} v. \end{aligned}$$

Finally, given both $\mathbf{t}_B$ and $\mathbf{t}'_B$, the adaptor witness can be recovered as $\mathbf{m} = \mathbf{t}_B - \mathbf{t}'_B$. The extracted witness is accepted only if it satisfies $\mathbf{M} \stackrel{?}{=} \mathbf{A}'\mathbf{m}$. Analogous to the dichotomic structure, $\mathbf{t}_A$ and other parts represent the first component, σ_1, while $\mathbf{t}_B$ represents the second component, σ_2.

5.3 Binding property of the commitment

The adaptor component $\mathbf{t}_B$ does not require a hiding property. Indeed, both $\mathbf{t}'_B = \mathbf{B}\mathbf{s}_2$ and $\mathbf{t}_B = \mathbf{B}\mathbf{s}_2 + \mathbf{m}$ are public values. Hence, anyone can recover the adaptor witness as $\mathbf{m} = \mathbf{t}_B - \mathbf{t}'_B$. This is intentional, since in adaptor signatures the witness is revealed after successful adaptation. Note that witness extraction is done publicly since pre-signature verification is public. We now discuss binding. Consider the commitment component $\mathbf{t}_B = \mathbf{B}\mathbf{s}_2 + \mathbf{m}$. A valid opening of $\mathbf{t}_B$ consists of a pair $(\mathbf{s}_2, \mathbf{m})$ satisfying the above equation. Suppose there exist two valid openings $(\mathbf{s}_2, \mathbf{m})$ and $(\mathbf{s}'_2, \mathbf{m}')$ such that both produce the same commitment

value $\mathbf{t}_B$. Then we have $\mathbf{B}\mathbf{s}_2 + \mathbf{m} = \mathbf{B}\mathbf{s}_2' + \mathbf{m}'$. Rearranging yields $\mathbf{m} - \mathbf{m}' = \mathbf{B}(\mathbf{s}_2' - \mathbf{s}_2)$, or equivalently,

$$\begin{bmatrix}\mathbf{B}\ \mathbf{I}\end{bmatrix}\begin{bmatrix}\mathbf{s}_2 - \mathbf{s}_2' \\ \mathbf{m}' - \mathbf{m}\end{bmatrix} = 0,$$

where $\mathbf{I}$ is the identity matrix. In addition, equality of the first commitment component

$$\mathbf{t}_A = \mathbf{A}_1\mathbf{s}_1 + \mathbf{A}_2\mathbf{s}_2 = \mathbf{A}_1\mathbf{s}_1' + \mathbf{A}_2\mathbf{s}_2'$$

implies $\mathbf{A}_1(\mathbf{s}_1 - \mathbf{s}_1') + \mathbf{A}_2(\mathbf{s}_2 - \mathbf{s}_2') = 0$, or equivalently,

$$\begin{bmatrix}\mathbf{A}_1\ \mathbf{A}_2\end{bmatrix}\begin{bmatrix}\mathbf{s}_1 - \mathbf{s}_1' \\ \mathbf{s}_2 - \mathbf{s}_2'\end{bmatrix} = 0.$$

We assume that when the adaptor reveals $\mathbf{M}$, it proves the smallness of $\mathbf{m}$. Since all witness components are required to be short, a non-zero difference yields a short vector in the kernel of $\begin{bmatrix}\mathbf{A}_1\ \mathbf{A}_2\end{bmatrix}$ and $\begin{bmatrix}\mathbf{B}\ \mathbf{I}\end{bmatrix}$, which contradicts the hardness of the M-SIS problem. Therefore, under the M-SIS assumption, the opening $(\mathbf{s}_1, \mathbf{s}_2, \mathbf{m})$ is unique, and in particular the adaptor witness $\mathbf{m}$ is uniquely determined by the commitment.

6 Lattice-based Blind Adaptor Signature (L-BAS) Scheme

In this section, we present a lattice-based instantiation of the BAS construction defined in section 4. Specifically using the dichotomic property of the ABDLOP commitment scheme, we build on the blind signature scheme in [10] and define the following BAS algorithms.

Key Generation. $(sk, pk) \leftarrow \text{BAS.KGen}(1^\lambda)$

Let $\mathcal{R}_q = \mathbb{Z}_q[X]/(X^d + 1)$ be the polynomial ring used by the scheme. The algorithm samples two short secret polynomials $f, g \leftarrow D^d_{\sigma_0} \subset \mathcal{R}_q$ and derives the public matrix $\mathbf{A} = (fg^{-1}, 1) \in \mathcal{R}_q^{1\times 2}$. The polynomials f and g are then used to construct a trapdoor basis $\mathbf{T} \in \mathcal{R}_q^{2\times 2}$ for the lattice $\mathcal{L}_q^\perp(\mathbf{A})$. This trapdoor is kept secret and enables the signer to sample short preimages during the signing protocol. The algorithm further samples a random matrix $\mathbf{A}' \leftarrow \mathcal{R}_q^{1\times 2}$ associated with the statement component.

The setup also samples an independent uniformly random public matrix $\mathbf{D} \in \mathcal{R}_q^{1\times 2}$, which is used as the public blinding matrix. In addition, a secret seed $\kappa \leftarrow \{0,1\}^\lambda$ is generated to provide deterministic randomness for the preimage-sampling procedure. Furthermore, the algorithm initializes the auxiliary public material required by the protocol, namely the public key for the public-key encryption mechanism and the public parameters of the ABDLOP proof system. These auxiliary components are treated as part of the public setup. Thus, the public key and the secret signing key are defined as $pk =$

$(\mathbf{A}', \mathbf{A}, \mathbf{D}, pk_{\text{PKE}}, pk_{\text{ABDLOP}})$ and $sk = (\mathbf{T}, \kappa)$.

Blind Pre-signing. $\tilde{\sigma}' \leftarrow \text{BAS.BpSign}(sk, \mu')$

An interactive algorithm run by the signer, user, and adaptor. Where the input μ' contains the blinded message, statement, and other parameters from both user and adaptor, and sk is the signer's secret key. The adaptor, user, and signer take the following steps in order to generate a blind pre-signature $\tilde{\sigma}'$.

1. The adaptor samples a bounded witness $\mathbf{m} \leftarrow \text{Bin}_1^{2d}$ and computes the corresponding adaptor statement $\mathbf{M} = \mathbf{A}'\mathbf{m}$. To prove that $\mathbf{M}$ is generated from an admissible bounded witness, the adaptor computes an ABDLOP commitment $\text{com}_{\mathbf{m}} = \text{ABDLOP.Com}(\mathbf{m}; \boldsymbol{s}_2')$. It then generates a zero-knowledge proof of knowledge
$$\pi_{\mathbf{M}} \leftarrow \text{NIZKP}\left\{(\mathbf{m}) : \mathbf{M} = \mathbf{A}'\mathbf{m} \ \wedge \ \mathbf{m} \in \text{Bin}_1^{2d}\right\}.$$
The adaptor sends $(\mathbf{M}, \text{com}_{\mathbf{m}}, \pi_{\mathbf{M}})$ to the user.
2. The user verifies $\pi_{\mathbf{M}}$. If the proof is invalid, the user aborts. Otherwise, the user accepts $\mathbf{M}$ as a valid adaptor statement and continues the pre-signing protocol.
3. The user samples a random polynomial vector $\mathbf{r} \leftarrow \text{Bin}_1^{2d}$ and computes $\mathbf{c} = \mathbf{D}\mathbf{r} + H(\mu, \rho, \rho')$, where $\rho = G(\mathbf{r})$ and $\rho' = G(\mathbf{M})$. Here, μ denotes the message, and $H(\cdot)$ and $G(\cdot)$ are modeled as random oracles which map onto $\mathcal{R}_q$. The user also computes a ciphertext $ct = PKE.\text{Enc}(\mathbf{r}, \mu; pk; r)$ encrypting $\mathbf{r}$ and μ under the public key pk. The ciphertext serves as encryption to the sky since the signer does not possess the decryption key and therefore cannot recover the encrypted values. However, it is required in proving the one-more unforgeability as it reduces the soundness error in the security proof [10].
4. The user sends $\mu' = \{\mathbf{c}, \text{Com}_{\mathbf{c}}, \pi_{\mathbf{c}}\}$ to the signer. $\pi_{\mathbf{c}}$ proves that $\mathbf{c}$ is a well-formed encryption and it also proves the smallness of $\mathbf{r}$.
$$\pi_{\mathbf{c}} \leftarrow \text{NIZKP}\left\{\begin{array}{l}(\mathbf{r}, \mu, r) : \mathbf{c} = \mathbf{D}\mathbf{r} + H(\mu, \rho, \rho') \\ \qquad\qquad \wedge \ ct = Enc(\mathbf{r}, \mu; pk, r) \ \wedge \ \mathbf{r} \in \text{Bin}_1^{2d}\end{array}\right\}.$$
5. After verifying $\pi_{\mathbf{c}}$, by using the trapdoor signing key, the signer computes a short vector $\mathbf{s}$ such that
$$\mathbf{A}\mathbf{s} = \mathbf{D}\mathbf{r} + H(\mu, \rho, \rho'), \quad \|\mathbf{s}\| \leq \beta_s. \tag{7}$$
The signer returns $\tilde{\sigma}' = \mathbf{s}$ as a blind pre-signature to the user. Note that the NIZK proofs of $\pi_{\mathbf{M}}$ and $\pi_{\mathbf{c}}$ are efficiently created using the linear-size zero-knowledge proof in Fig. 10 of [31, Figure 10] which proves smallness and correctness[2] of committed parameter in ABDLOP commitment.

[2] It is worth mentioning that this proof can establish more general properties; here, we present only the part required for our construction.

Pre-signing. $\tilde{\sigma} \leftarrow \text{BAS.pSign}(\tilde{\sigma}', \mu, \mathbf{M})$

Using the statement $\mathbf{M}$ and the message μ, the user takes the following steps to generate pre-signature $\tilde{\sigma}$ from the blind one $\tilde{\sigma}'$ which it received from the signer.

1. To generate pre-signature, the user proves the knowledge of short vectors $\mathbf{s}$ and $\mathbf{r}$ satisfying Eq. (7) and $\mathbf{M} = \mathbf{A}'\mathbf{m}$. Then, the user samples ABDLOP commitment randomness $\mathbf{s}_2 \leftarrow S_\beta^{m_2}$ and computes

$$\mathbf{t}_A = \mathbf{A}_1 \begin{bmatrix} \mathbf{s} \\ \mathbf{r} \end{bmatrix} + \mathbf{A}_2 \mathbf{s}_2, \qquad \mathbf{t}'_B = \mathbf{B}\mathbf{s}_2.$$

2. After computing $\mathbf{t}_A$ and $\mathbf{t}'_B$, the user samples masking vectors $\mathbf{y}_1$ and $\mathbf{y}_2$, and computes $\mathbf{w} = \mathbf{A}_1\mathbf{y}_1 + \mathbf{A}_2\mathbf{y}_2$. Define [3]

$$\hat{\mathbf{s}} = \begin{bmatrix} \mathbf{s} \\ \mathbf{r} \\ \mathbf{m} \end{bmatrix}, \qquad \mathbf{y} = \begin{bmatrix} \mathbf{y}_1 \\ -\mathbf{B}\mathbf{y}_2 \end{bmatrix}.$$

Then the user computes $g_1 = \hat{\mathbf{s}}^\mathsf{T}\mathbf{R}_2\mathbf{y} + \mathbf{y}^\mathsf{T}\mathbf{R}_2\hat{\mathbf{s}} + \mathbf{r}_1^\mathsf{T}\mathbf{y}$, and hides this value by setting $t = \mathbf{b}^\mathsf{T}\mathbf{s}_2 + g_1$. Additionally, the user computes $v = \mathbf{y}^\mathsf{T}\mathbf{R}_2\mathbf{y} + \mathbf{b}^\mathsf{T}\mathbf{y}_2$. The values $(\mathbf{w}, t, v)$ are then used to derive the challenge

$$c = H(\mu, \rho, \rho', \mathbf{t}_A, \mathbf{w}, t, v).$$

Finally, the user computes $\mathbf{z}_1 = \mathbf{y}_1 + c\begin{bmatrix} \mathbf{s} \\ \mathbf{r} \end{bmatrix}, \mathbf{z}_2 = \mathbf{y}_2 + c\mathbf{s}_2$. After applying the required rejection-sampling procedure, the user outputs the pre-signature $\tilde{\sigma} = (\mu, \rho, \mathbf{t}_A, \mathbf{t}'_B, \mathbf{w}, t, v, \mathbf{z}_1, \mathbf{z}_2)$.

Pre-verification. $b \leftarrow \text{BAS.pVrfy}(pk, \mu, \tilde{\sigma}, \mathbf{M})$

To verify the pre-signature $\tilde{\sigma}$ with respect to the adaptor statement $\mathbf{M}$ and $\rho' = G(\mathbf{M})$, the verifier first checks the norm bounds on $\mathbf{z}_1$ and $\mathbf{z}_2$. If either check fails, the verifier rejects. Otherwise, it reconstructs $\hat{\mathbf{z}} = \begin{bmatrix} \mathbf{z}_1 \\ c\mathbf{t}'_B - \mathbf{B}\mathbf{z}_2 \end{bmatrix}$. The verifier returns $b = 1$ if the following equations hold:

$$\mathbf{A}_1\mathbf{z}_1 + \mathbf{A}_2\mathbf{z}_2 - c\mathbf{t}_A \stackrel{?}{=} \mathbf{w}, \quad \text{and} \quad \widehat{\mathbf{z}}^\mathsf{T}\mathbf{R}_2\widehat{\mathbf{z}} + c\mathbf{r}_1^\mathsf{T}\widehat{\mathbf{z}} + c^2 r_0 - ct + \mathbf{b}^\mathsf{T}\mathbf{z}_2 + c^2\mathbf{M} \stackrel{?}{=} v.$$

In above equation, we write $\mathbf{M}$ rather than $\sum_{i=1}^{n} M_i$. Under our assumption on $\mathbf{A}'$, the dimension of $\mathbf{M}$ is one, and therefore $\mathbf{M}$ is equivalent to its sole component. For simplicity, we use $\mathbf{M}$.

This algorithm verifies the NIZK proof of $\mathbf{s}$. Specifically, in the pre-signing phase, the user generates an ABDLOP commitment to $\mathbf{s}$ along with a corresponding NIZK proof. In the pre-verification phase, the algorithm verifies the validity of the ABDLOP commitment and ensures that the required conditions are satisfied, including the norm bounds on $\mathbf{s}$ and $\mathbf{r}$.

[3] For simplicity, we remove the automorphism parameters $\sigma(\mathbf{s}), \sigma(\mathbf{r})$ and $\sigma(\mathbf{m})$.

Adaptation. $\sigma \leftarrow \text{BAS.Adapt}(pk, \tilde{\sigma}, \mathbf{m})$

Given a valid pre-signature $\tilde{\sigma}$ and the witness $\mathbf{m}$ corresponding to the adaptor statement $\mathbf{M} = \mathbf{A}'\mathbf{m}$, the adaptor computes $\mathbf{t}_B = \mathbf{t}'_B + \mathbf{m}$. The final signature is $\sigma = (\mu, \rho, \mathbf{t}_A, \mathbf{t}_B, \mathbf{t}'_B, \mathbf{w}, t, v, \mathbf{z}_1, \mathbf{z}_2)$. Anyone can verify the final signature by verification process.

Extraction. $m \leftarrow \text{BAS.Extract}(pk, \tilde{\sigma}, \sigma, \mathbf{M})$

Given a valid pre-signature $\tilde{\sigma}$ and the corresponding final signature σ, anyone can extract the adaptor witness by computing $\mathbf{m} = \mathbf{t}_B - \mathbf{t}'_B$. The extracted witness is then verified by checking $\mathbf{M} \stackrel{?}{=} \mathbf{A}'\mathbf{m}$, $\|\mathbf{m}\| \leq \mathcal{S}_1\sqrt{4d}$, where $\mathcal{S}_1$ denotes the standard deviation of Bin_1^{2d}, and $\rho' \stackrel{?}{=} G(\mathbf{A}'\mathbf{m})$. If both checks hold, then $\mathbf{m}$ is a valid witness for the adaptor statement $\mathbf{M}$.

Verification. $b \leftarrow \text{BAS.Vrfy}(pk, \mu, \sigma)$

The verifier can then recover the adaptor witness as $\mathbf{m} = \mathbf{t}_B - \mathbf{t}'_B$. After asserting that $\|\mathbf{m}\| \leq \mathcal{S}_1\sqrt{4d}$, it computes $\rho' = G(\mathbf{A}'\mathbf{m}) = G(\mathbf{A}'(\mathbf{t}_B - \mathbf{t}'_B))$. The value ρ' is subsequently used in the challenge computation and signature verification. If norm bound of $\mathbf{z}_1$ and $\mathbf{z}_2$ are verified, the verifier checks the validity of following equations

$$\begin{aligned} \mathbf{A}_1\mathbf{z}_1 + \mathbf{A}_2\mathbf{z}_2 &\stackrel{?}{=} \mathbf{w} + c\mathbf{t}_A, \\ \mathbf{z}^\mathsf{T}\mathbf{R}_2\mathbf{z} + c\mathbf{r}_1^\mathsf{T}\mathbf{z} + c^2 r_0 - ct + \mathbf{b}^\mathsf{T}\mathbf{z}_2 &\stackrel{?}{=} v, \\ \widehat{\mathbf{z}}^\mathsf{T}\mathbf{R}_2\widehat{\mathbf{z}} + c\mathbf{r}_1^\mathsf{T}\widehat{\mathbf{z}} + c^2 r_0 - ct + \mathbf{b}^\mathsf{T}\mathbf{z}_2 + c^2\mathbf{M} &\stackrel{?}{=} v. \end{aligned}$$

If three equations hold, it returns $b = 1$, otherwise $b = 0$. Consequently, any modification of the witness changes ρ' and invalidates the verification equation, thereby binding the adapted signature to the extracted witness. Note that pre-signature verification in adaptor-signature schemes is public. Therefore, witness extraction is inherently public as well, since anyone possessing a valid pre-signature and its corresponding adapted signature can recover the witness.

6.1 Security Analysis

In what follows, we prove that L-BAS satisfies extractability, unique extractability, and computational pre-verification soundness. The blindness and one-more unforgeability properties follow directly from the security of the underlying blind signature scheme [10].

Theorem 1 (Extractability) *Assuming the one-more unforgeability of the underlying blind signature scheme, the computational binding property of the ABDLOP commitment scheme, and the hardness of* $\text{M-SIS}_{1,2,\beta}$ *problem with* $\beta = 2\mathcal{S}_1\sqrt{4d}$*, the proposed L-BAS scheme satisfies witness extractability in the random-oracle model.*

Proof. Let $\mathcal{A}$ be a probabilistic polynomial-time adversary that wins the extractability experiment with non-negligible probability. We show that such an adversary can be used either to forge a signature for the underlying blind signature scheme Σ, or to violate binding property of ABDLOP commitment, or to solve the corresponding M-SIS problem.

Let $\mathcal{A}$ output a message μ^*, an adaptor statement $\mathbf{M}^*$, a valid pre-signature $\tilde{\sigma}^*$ and a valid final signature σ^*. Let $\tilde{c}^*$ be the challenge associated with the pre-signature and let c^* be the challenge associated with the final signature.

First, suppose that no relevant pre-signature exists for μ^* and $\mathbf{M}^*$. Then σ^* is a valid blind signature on μ^* that was not obtained through the appropriate signing oracle. Therefore, $\mathcal{A}$ directly yields a one-more forgery against the underlying blind signature scheme Σ. By the one-more unforgeability of Σ, this case occurs only with negligible probability.

It remains to consider the case where a relevant pre-signature exists. We distinguish two cases depending on whether the challenges are equal or not.

Case 1: $\tilde{c}^* = c^*$**.** In this case, the pre-signature and the final signature have the same Fiat–Shamir challenge. Hence, the adaptor component can be extracted directly from the difference between the corresponding commitment values. The extractor computes $\mathbf{a} = \mathbf{t}_B - \mathbf{t}'_B$. By the correctness of the pre-signature and final-signature verification equations, this extracted value satisfies $\mathbf{A}'\mathbf{a} = \mathbf{M}$. Therefore, $\mathbf{a}$ is a valid witness for the adaptor statement, provided that it also satisfies the required norm bound.

Now suppose that the adversary claims a different bounded witness $\mathbf{m}$ for the same statement, namely $\mathbf{A}'\mathbf{m} = \mathbf{M}$. Then we obtain

$$\mathbf{A}'(\mathbf{m} - \mathbf{a}) = \mathbf{A}'\mathbf{m} - \mathbf{A}'\mathbf{a} = \mathbf{M} - \mathbf{M} = \mathbf{0} \pmod{q}.$$

If $\mathbf{m} \neq \mathbf{a}$, then $\mathbf{m} - \mathbf{a}$ is a nonzero short vector in the kernel of the public matrix $\mathbf{A}'$. This gives a solution to the corresponding M-SIS problem. Consequently, except with negligible probability under the hardness of M-SIS$_{1,2,\beta}$, where $\beta = 2\mathcal{S}_1\sqrt{4d}$ and $\mathcal{S}_1$ denotes the standard deviation of Bin_1^{2d}, the extracted value $\mathbf{a}$ necessarily coincides with the unique valid witness. Thus, in the case $\tilde{c}^* = c^*$, the extractor successfully recovers a valid adaptor witness.

Case 2: $\tilde{c}^* \neq c^*$**.** Assume that the adversary outputs a valid final signature whose challenge c^* differs from the challenge $\tilde{c}^*$ associated with the corresponding pre-signature. By the forking lemma, we can rewind the adversary while keeping the same random-oracle input prefix and obtain two accepting transcripts

$$(c^*, \mathbf{z}_1^*, \mathbf{z}_2^*, \ldots) \quad \text{and} \quad (c^{*'}, \mathbf{z}_1^{*'}, \mathbf{z}_2^{*'}, \ldots),$$

with $c^* \neq c^{*'}$, for the same statement.

Since both transcripts are accepting and correspond to the same random-oracle input prefix, the reconstructed first-round commitments must be equal. In particular, $\mathbf{A}_1\mathbf{z}_1^* + \mathbf{A}_2\mathbf{z}_2^* - c^*\mathbf{t}_A = \mathbf{A}_1\mathbf{z}_1^{*'} + \mathbf{A}_2\mathbf{z}_2^{*'} - c^{*'}\mathbf{t}_A$. Hence,

$$\mathbf{A}_1(\mathbf{z}_1^* - \mathbf{z}_1^{*'}) + \mathbf{A}_2(\mathbf{z}_2^* - \mathbf{z}_2^{*'}) = (c^* - c^{*'})\mathbf{t}_A \pmod{q}.$$

Since $c^* - c^{*'}$ is invertible in $\mathcal{R}_q$ [31, subsection 2.7], we define

$$\widehat{\mathbf{s}}_1 = (c^* - c^{*'})^{-1}(\mathbf{z}_1^* - \mathbf{z}_1^{*'}), \qquad \widehat{\mathbf{s}}_2 = (c^* - c^{*'})^{-1}(\mathbf{z}_2^* - \mathbf{z}_2^{*'}).$$

It follows that $\mathbf{t}_A = \mathbf{A}_1\widehat{\mathbf{s}}_1 + \mathbf{A}_2\widehat{\mathbf{s}}_2 \pmod q$. Furthermore, the corresponding message component of the ABDLOP opening is uniquely determined by $\widehat{\mathbf{m}} = \mathbf{t}_B - \mathbf{B}\widehat{\mathbf{s}}_2 \pmod q$. Therefore, the two accepting transcripts yield a valid opening $(\widehat{\mathbf{s}}_1, \widehat{\mathbf{s}}_2, \widehat{\mathbf{m}})$ of the same ABDLOP commitment.

If this opening is distinct from the opening contained in the pre-signature, then the adversary has produced two different valid openings of the same ABDLOP commitment. This contradicts the computational binding property of the ABDLOP commitment scheme. Consequently, under the binding security of ABDLOP, this case can occur only with negligible probability.

Theorem 2 (Unique Extractability) *Assume that* $\mathrm{M-SIS}_{1,2,\beta}$ *is hard where* $\beta = 2\mathcal{S}_1\sqrt{4d}$. *Then no efficient adversary can produce two distinct admissible witnesses for the same adaptor statement. Consequently, the L-BAS scheme satisfies unique extractability.*

Proof. Suppose that an efficient adversary produces two distinct valid adaptations of the same pre-signature, yielding extracted witnesses $\mathbf{m}$ and $\mathbf{m}'$ for the same statement $\mathbf{M}$. Since both witnesses are accepted, we have $\mathbf{A}'\mathbf{m} = \mathbf{M}$ and $\mathbf{A}'\mathbf{m}' = \mathbf{M}$, with $\|\mathbf{m}\|, \|\mathbf{m}'\| \leq \mathcal{S}_1\sqrt{4d}$. Distinctness of the adaptations implies $\mathbf{m} \neq \mathbf{m}'$. Hence, for $\boldsymbol{\delta} = \mathbf{m} - \mathbf{m}'$, we obtain

$$\boldsymbol{\delta} \neq 0, \qquad \mathbf{A}'\boldsymbol{\delta} = 0, \qquad \|\boldsymbol{\delta}\| \leq 2\mathcal{S}_1\sqrt{4d}.$$

Thus, $\boldsymbol{\delta}$ is a nonzero short vector in the kernel of $\mathbf{A}'$, which solves the underlying $\mathrm{M-SIS}_{n,m,\beta}$ problem where $\beta = 2\mathcal{S}_1\sqrt{4d}$. This contradicts the assumed hardness of M − SIS problem. Therefore, except with negligible probability, the L-BAS scheme satisfies unique extractability.

Theorem 3 (Computational Pre-Verification Soundness) *Assuming the computational soundness of the ABDLOP proof system [31], the proposed L-BAS scheme satisfies computational pre-verification soundness with respect to the adaptor relation* L_{Rel}.

Proof. Assume that a probabilistic polynomial-time adversary $\mathcal{A}$ can generate a statement $\mathbf{M} \notin L_{\mathrm{Rel}}$, where $L_{\mathrm{Rel}} = \{\mathbf{M} : \exists, \mathbf{m} \text{ such that } \mathbf{M} = \mathbf{A}\mathbf{m} \wedge |\mathbf{m}| \leq \beta_m\}$, together with a proof $\pi_{\mathbf{m}}$ and a pre-signature $\tilde{\sigma}$ that passes pre-verification. We emphasize that the adversary generates both the statement $\mathbf{M}$ and its corresponding proof $\pi_{\mathbf{m}}$, in addition to the pre-signature $\tilde{\sigma}$. Successful pre-verification implies, in particular, that $\pi_{\mathbf{m}}$ is accepted by the NIZKP verifier. However, since $\mathbf{M} \notin L_{\mathrm{Rel}}$, therefore, the $\mathcal{A}$ produces an accepting NIZKP for a false statement, contradicting the computational soundness of the NIZKP system except with negligible probability. Hence, L-BAS satisfies computational pre-verification soundness.

Unlinkability. In a conventional adaptor signature scheme, unlinkability requires that adapted signatures be computationally indistinguishable from ordinary signatures.

In our 3-party construction, however, every valid signature is obtained through the adaptation procedure; hence, there is no separate distribution of ordinary (non-adapted) signatures against which adapted signatures need to be indistinguishable. Moreover, in the underlying blind signature scheme, the additional component $\mathbf{t}'_B$ may provide information that prevents unlinkability in the conventional sense. To address this issue, the user can itself act as an adaptor by generating an independent dummy statement and requesting a blind signature with respect to this statement. Consequently, signatures generated in such individual executions follow the same adaptation procedure as signatures associated with an actual adaptor, preventing an observer from distinguishing or linking a signature to the involvement of the actual adaptor based solely on its form.

In contrast, the first BAS scheme of Qin et al. [41] does not consider unlinkability as a security requirement. The rationale is that every valid signature is generated through the adaptation of the user's unblinded pre-signature issued by the signer. Consequently, all valid signatures originate from the same distribution, leaving no alternative signature distribution from which adapted signatures must be hidden.. In L-BAS, we follow the approach of Qin et al. [41] and do not include unlinkability as a core security requirement.

Discussion. One may argue that a malicious adaptor, after receiving a valid pre-signature, could simply withhold the corresponding adapted signature and thereby prevent the protocol from completing. We first note that such behavior also prevents the adaptor from obtaining the intended benefit of the exchange, since the service cannot be completed without releasing the adapted signature. Nevertheless, this abort behavior can be explicitly addressed at the protocol level by binding an expiration condition to the Fiat-Shamir challenge. In particular, when the user requests a signature, instead of computing $\mathbf{c} = \mathbf{Dr} + H(\mu, \rho, \rho')$, the user includes auxiliary information info, containing an expiration time and computes $\mathbf{c} = \mathbf{Dr} + H(\mu, \rho, \rho', \mathsf{info})$. Consequently, the pre-signature is cryptographically bound to the specified expiration condition. If the adaptor withholds the final signature beyond this time, the corresponding pre-signature can no longer be used for the intended exchange, and the user may invoke a refund mechanism of the underlying application protocol.
Moreover, if dispute resolution is required, the user can provide the relevant pre-signature transcript, including the randomness and auxiliary values necessary to recompute $\mathbf{c}$, thereby demonstrating that the pre-signature was generated for the claimed message and expiration condition. We emphasize that this mechanism does not cryptographically force the adaptor to complete the adaptation; rather, it prevents indefinite blocking by combining the adaptor-signature protocol with an application-level timeout and refund mechanism.

Parameter Selection: The proposed L-BAS scheme inherits the same parameter set as the lattice-based blind signature scheme of Beullens [10]. Indeed, our construction can be viewed as an adaptor extension of the original scheme, where the main modification is the introduction of the adaptor functionality through the ABDLOP commitment framework.

More specifically, unlike the original construction of [10], which does not contain a $\mathbf{t}_B$ term, our construction introduces $\mathbf{t}_B = \mathbf{B}\mathbf{s}_2 + \mathbf{m}$ to embed the message component $\mathbf{m}$ into the ABDLOP commitment. Since the witness is incorporated into the additive message slot of the commitment and does not affect the security analysis of the underlying blind signature, no changes to the original parameter selection are required.

The additional overhead arises during the statement-generation phase. Before requesting a pre-signature, the adaptor must convince the user that the public statement $\mathbf{M}$ is derived from a bounded witness $\mathbf{m}$. To this end, the adaptor provides an ABDLOP commit-and-prove protocol demonstrating knowledge of a witness $\mathbf{m}$ satisfying $\mathbf{M} = \mathbf{A}'\mathbf{m}$ and $\|\mathbf{m}\| \leq \beta_m$. According to the communication analysis of the ABDLOP framework, this proof requires approximately 14 KB of additional communication. Compared to the original blind signature scheme [10], our scheme introduces two additional components, namely $\mathbf{t}'_B, \mathbf{t}_B$. Under the parameterization of [10], the original signature size is approximately 22 KB, and the added component accounts for roughly 5.2 KB [4]. Therefore, the overall signature size becomes approximately 27.2 KB. A detailed discussion of the parameter settings and concrete instantiations is provided in [10].

7 Conclusion

In this paper, we formalized blind adaptor signatures (BAS) and introduced a lattice-based BAS (L-BAS) scheme that is suitable for privacy-preserving, fairness-oriented, and post-quantum secure environments. Our construction exploits the dichotomic structure of the ABDLOP commit-and-prove framework, which naturally facilitates the realization of blind adaptor signatures and is particularly well-suited for applications requiring strong fairness guarantees. Our construction introduces a small additional overhead: the signature size increases by approximately 5.2 KB, and the verification cost is also slightly increased due to the presence of the additional component. Nevertheless, these overheads remain moderate in practice.

References

1. Agrawal, S., Kirshanova, E., Stehlé, D., Yadav, A.: Practical, round-optimal lattice-based blind signatures. In: Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security. p. 39–53. CCS '22, Association for Computing Machinery, New York, NY, USA (2022)

[4] The component $\mathbf{t}'_B$ and $\mathbf{t}_B$ consist of 512 coefficients. For $q \approx 2^{41}$, each coefficient requires 41 bits, yielding a total size of $2 \times 512 \times 41 = 41{,}984$ bits ≈ 5.2 KB.

2. Alkeilani Alkadri, N., El Bansarkhani, R., Buchmann, J.: Blaze: Practical lattice-based blind signatures for privacy-preserving applications. In: Bonneau, J., Heninger, N. (eds.) Financial Cryptography and Data Security. pp. 484–502. Springer International Publishing, Cham (2020)
3. Alkeilani Alkadri, N., El Bansarkhani, R., Buchmann, J.: On lattice-based interactive protocols: An approach with less or no aborts. In: Liu, J.K., Cui, H. (eds.) Information Security and Privacy. pp. 41–61. Springer International Publishing, Cham (2020)
4. Aumayr, L., Ersoy, O., Erwig, A., Faust, S., Hostáková, K., Maffei, M., Moreno-Sanchez, P., Riahi, S.: Generalized channels from limited blockchain scripts and adaptor signatures. In: Tibouchi, M., Wang, H. (eds.) Advances in Cryptology – ASIACRYPT 2021. pp. 635–664. Springer International Publishing, Cham (2021)
5. Baldimtsi, F., Cheng, J., Goyal, R., Yadav, A.: Non-interactive blind signatures: Post-quantum and stronger security. In: Chung, K.M., Sasaki, Y. (eds.) Advances in Cryptology – ASIACRYPT 2024. pp. 70–104. Springer Nature Singapore, Singapore (2025)
6. Banaszczyk, W.: New bounds in some transference theorems in the geometry of numbers. Mathematische Annalen **296**(4), 625–635 (1993)
7. Banegas, G., Carrier, K., Chailloux, A., Couvreur, A., Debris-Alazard, T., Gaborit, P., Karpman, P., Loyer, J., Niederhagen, R., Sendrier, N., Smith, B., Tillich, J.P.: Wave. Technical report, National Institute of Standards and Technology (2023), https://csrc.nist.gov/Projects/pqc-dig-sig/round-1-additional-signatures
8. Beullens, W., Campos, F., Celi, S., Hess, B., Kannwischer, M.J.: MAYO. Technical report, National Institute of Standards and Technology (2024), https://csrc.nist.gov/Projects/pqc-dig-sig/round-2-additional-signatures
9. Beullens, W., Chen, M.S., Ding, J., Gong, B., Kannwischer, M.J., Patarin, J., Peng, B.Y., Schmidt, D., Shih, C.J., Tao, C., Yang, B.Y.: UOV—Unbalanced Oil and Vinegar. Technical report, National Institute of Standards and Technology (2024), https://csrc.nist.gov/Projects/pqc-dig-sig/round-2-additional-signatures
10. Beullens, W., Lyubashevsky, V., Nguyen, N.K., Seiler, G.: Lattice-based blind signatures: Short, efficient, and round-optimal. In: Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security. p. 16–29. CCS '23, Association for Computing Machinery, New York, NY, USA (2023)
11. Bootle, J., Lyubashevsky, V., Nguyen, N.K., Sorniotti, A.: A framework for practical anonymous credentials from lattices. In: Handschuh, H., Lysyanskaya, A. (eds.) Advances in Cryptology – CRYPTO 2023. pp. 384–417. Springer Nature Switzerland, Cham (2023)
12. Bouillaguet, C., Feneuil, T., Maire, J., Rivain, M., Sauvage, J., Vergnaud, D.: Blinding post-quantum hash-and-sign signatures. Cryptology ePrint Archive, Paper 2025/895 (2025), https://eprint.iacr.org/2025/895
13. Chase, M., Meiklejohn, S., Zaverucha, G.: Algebraic macs and keyed-verification anonymous credentials. In: Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security. p. 1205–1216. CCS '14, Association for Computing Machinery, New York, NY, USA (2014). https://doi.org/10.1145/2660267.2660328
14. Chaum, D.: Blind signatures for untraceable payments. In: Chaum, D., Rivest, R.L., Sherman, A.T. (eds.) Advances in Cryptology. pp. 199–203. Springer US, Boston, MA (1983)

15. Chaum, D., Moser, T.: ecash 2.0: Inalienably private and quantum-resistant to counterfeiting. https://chaum.com/publications/#iTK6j8 (2022), unpublished manuscript
16. Dai, W., Okamoto, T., Yamamoto, G.: Stronger security and generic constructions for adaptor signatures. In: Isobe, T., Sarkar, S. (eds.) Progress in Cryptology – INDOCRYPT 2022. pp. 52–77. Springer International Publishing, Cham (2022)
17. Decker, C., Wattenhofer, R.: A fast and scalable payment network with bitcoin duplex micropayment channels. In: Pelc, A., Schwarzmann, A.A. (eds.) Stabilization, Safety, and Security of Distributed Systems. pp. 3–18. Springer International Publishing, Cham (2015)
18. Eckey, L., Faust, S., Hostáková, K., Roos, S.: Splitting payments locally while routing interdimensionally. IACR Cryptol. ePrint Arch. **2020**, 555 (2020), https://api.semanticscholar.org/CorpusID:218653508
19. Erwig, A., Faust, S., Hostáková, K., Maitra, M., Riahi, S.: Two-party adaptor signatures from identification schemes. In: Garay, J.A. (ed.) Public-Key Cryptography – PKC 2021. pp. 451–480. Springer International Publishing, Cham (2021)
20. Esgin, M.F., Ersoy, O., Erkin, Z.: Post-quantum adaptor signatures and payment channel networks. In: Chen, L., Li, N., Liang, K., Schneider, S. (eds.) Computer Security – ESORICS 2020. pp. 378–397. Springer International Publishing, Cham (2020)
21. Fiat, A., Shamir, A.: How to prove yourself: Practical solutions to identification and signature problems. In: Odlyzko, A.M. (ed.) Advances in Cryptology — CRYPTO' 86. pp. 186–194. Springer Berlin Heidelberg, Berlin, Heidelberg (1987)
22. Fischlin, M.: Round-optimal composable blind signatures in the common reference string model. In: Dwork, C. (ed.) Advances in Cryptology - CRYPTO 2006. pp. 60–77. Springer Berlin Heidelberg, Berlin, Heidelberg (2006)
23. Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing. p. 197–206. STOC '08, Association for Computing Machinery, New York, NY, USA (2008). https://doi.org/10.1145/1374376.1374407, https://doi.org/10.1145/1374376.1374407
24. Gerhart, P., Schröder, D., Soni, P., Thyagarajan, S.A.: Foundations of adaptor signatures. In: Joye, M., Leander, G. (eds.) Advances in Cryptology – EUROCRYPT 2024. pp. 161–189. Springer Nature Switzerland, Cham (2024)
25. Glaeser, N., Maffei, M., Malavolta, G., Moreno-Sanchez, P., Tairi, E., Thyagarajan, S.A.K.: Foundations of coin mixing services. In: Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security. p. 1259–1273. CCS '22, Association for Computing Machinery, New York, NY, USA (2022). https://doi.org/10.1145/3548606.3560637
26. Hanzlik, L.: Non-interactive blind signatures for random messages. In: Hazay, C., Stam, M. (eds.) Advances in Cryptology – EUROCRYPT 2023. pp. 722–752. Springer Nature Switzerland, Cham (2023)
27. Hauck, E., Kiltz, E., Loss, J., Nguyen, N.K.: Lattice-based blind signatures, revisited. In: Micciancio, D., Ristenpart, T. (eds.) Advances in Cryptology – CRYPTO 2020. pp. 500–529. Springer International Publishing, Cham (2020)
28. Hisai, M., Yanai, N.: Blind adaptor signatures, revisited: Stronger security definitions and their construction toward practical applications. Cryptology ePrint Archive, Paper 2026/060 (2026), https://eprint.iacr.org/2026/060
29. Langlois, A., Stehlé, D.: Worst-case to average-case reductions for module lattices. Designs, Codes and Cryptography **75**, 565–599 (2015)

30. Lyubashevsky, V.: Lattice signatures without trapdoors. In: Pointcheval, D., Johansson, T. (eds.) Advances in Cryptology – EUROCRYPT 2012. pp. 738–755. Springer Berlin Heidelberg, Berlin, Heidelberg (2012)
31. Lyubashevsky, V., Nguyen, N.K., Plançon, M.: Lattice-based zero-knowledge proofs and applications: Shorter, simpler, and more general. p. 71–101. Springer-Verlag, Berlin, Heidelberg (2022)
32. Lyubashevsky, V., Nguyen, N.K., Plancon, M.: Efficient lattice-based blind signatures via gaussian one-time signatures. In: Hanaoka, G., Shikata, J., Watanabe, Y. (eds.) Public-Key Cryptography – PKC 2022. pp. 498–527. Springer International Publishing, Cham (2022)
33. Lyubashevsky, V., Nguyen, N.K., Seiler, G.: Shorter lattice-based zero-knowledge proofs via one-time commitments. In: Garay, J.A. (ed.) Public-Key Cryptography – PKC 2021. pp. 215–241. Springer International Publishing, Cham (2021)
34. Madathil, V., Thyagarajan, S.A.K., Vasilopoulos, D., Fournier, L., Malavolta, G., Moreno-Sanchez, P.: Cryptographic oracle-based conditional payments. In: 30th Annual Network and Distributed System Security Symposium, NDSS 2023, San Diego, California, USA, February 27 - March 3, 2023. The Internet Society (2023)
35. National Institute of Standards and Technology: Module-Lattice-Based Digital Signature Standard (ML-DSA). Tech. Rep. FIPS 204, NIST (2024), https://doi.org/10.6028/NIST.FIPS.204
36. National Institute of Standards and Technology: Module-Lattice-Based Key-Encapsulation Mechanism Standard (ML-KEM). Tech. Rep. FIPS 203, NIST (2024), https://doi.org/10.6028/NIST.FIPS.203
37. National Institute of Standards and Technology: Stateless Hash-Based Digital Signature Standard (SLH-DSA). Tech. Rep. FIPS 205, NIST (2024), https://doi.org/10.6028/NIST.FIPS.205
38. National Security Agency: Announcing the Commercial National Security Algorithm Suite 2.0. Tech. Rep. U/OO/194427-22, National Security Agency (2025), https://media.defense.gov/2025/May/30/2003728741/-1/-1/0/CSA-CNSA-2.0-ALGORITHMS.PDF
39. del Pino, R., Katsumata, S.: A new framework for more efficient round-optimal lattice-based (partially) blind signature via trapdoor sampling. In: Dodis, Y., Shrimpton, T. (eds.) Advances in Cryptology – CRYPTO 2022. pp. 306–336. Springer Nature Switzerland, Cham (2022)
40. Prest, T., Fouque, P.A., Hoffstein, J., Kirchner, P., Lyubashevsky, V., Pornin, T., Ricosset, T., Seiler, G., Whyte, W., Zhang, Z.: FALCON specification (round 1 submission). Tech. rep., NIST Post-Quantum Cryptography Project (2017), https://falcon-sign.info/
41. Qin, X., Pan, S., Mirzaei, A., Sui, Z., Ersoy, O., Sakzad, A., Esgin, M.F., Liu, J.K., Yu, J., Yuen, T.H.: Blindhub: Bitcoin-compatible privacy-preserving payment channel hubs supporting variable amounts. In: 2023 IEEE Symposium on Security and Privacy (SP). pp. 2462–2480 (2023)
42. Rückert, M.: Lattice-based blind signatures. In: Abe, M. (ed.) Advances in Cryptology - ASIACRYPT 2010. pp. 413–430. Springer Berlin Heidelberg, Berlin, Heidelberg (2010)
43. Schnorr, C.P.: Security of blind discrete log signatures against interactive attacks. In: Qing, S., Okamoto, T., Zhou, J. (eds.) Information and Communications Security. pp. 1–12. Springer Berlin Heidelberg, Berlin, Heidelberg (2001)
44. Tairi, E., Moreno-Sanchez, P., Maffei, M.: Post-quantum adaptor signature for privacy-preserving off-chain payments. In: Borisov, N., Diaz, C. (eds.) Financial

Cryptography and Data Security. pp. 131–150. Springer Berlin Heidelberg, Berlin, Heidelberg (2021)

A Additional Preliminaries

In this section, we review the additional preliminaries that are required for our construction.

A.1 Blind Signature Definition

A blind signature scheme typically consists of the following algorithms and protocols:

1. **Key Generation:** The signer generates a public/private key pair. The private key is used for signing, while the public key is used for verification.
2. **Signing:** The requester, who owns the message to be signed, first blinds the message using a blinding factor. This factor prevents the signer from learning any information about the content of the message during the signing process. The signer then signs the blinded message, thereby producing a signature on the blinded version of the message. After receiving the blinded signature, the requester removes the blinding factor and obtains a valid signature on the original message.
3. **Verification:** Anyone can later verify the validity of the signature using the signer's public key, ensuring that the signature is legitimate and corresponds to the original message.

We require blind signature schemes to satisfy three main security properties: correctness, blindness, and one-more unforgeability. Correctness requires that honestly generated signatures verify with probability close to one. Blindness requires that the signer cannot link a valid signature to the specific signing interaction in which it was generated. One-more unforgeability requires that if a user is allowed to execute the signing protocol Q times, then the user cannot obtain more than Q valid signatures. The formal security requirements are defined as follows.

Definition 3 (Correctness). *A blind signature scheme* $(\mathrm{KeyGen}, \mathrm{Sign}, \mathrm{Verify})$ *is correct if there exists a negligible function* $\mathrm{negl}(\lambda)$ *such that, for any message* μ*, we have:*

$$\Pr\left[\mathrm{Verify}(\mathrm{pk}, \mu, \sigma) = 1 \;\middle|\; \begin{array}{l}(\mathrm{sk}, \mathrm{pk}) \leftarrow \mathrm{KeyGen}(1^\lambda) \\ \sigma \leftarrow \mathrm{Sign}(\mathrm{sk}, \mathrm{pk}, \mu)\end{array}\right] \geq 1 - \mathrm{negl}(\lambda).$$

Definition 4 (Blindness). *A blind signature scheme is* blind *if, for every PPT adversarial signer* $\mathcal{A}$*, the following experiment reveals no information about the association between the signing sessions and the resulting signatures, except with negligible probability.*

The adversary $\mathcal{A}$ chooses two messages m_0, m_1 and a public key pk*. The challenger samples a uniformly random bit $b \leftarrow \{0,1\}$ and invokes two honest users on messages μ_b and μ_{1-b}, respectively. The adversary interacts with both users as the signer and may arbitrarily schedule and interleave the two signing sessions. At the end of the experiment, the resulting signatures σ^b, σ^{1-b} are given to $\mathcal{A}$, which outputs a bit b' indicating its guess for b. Formally,*

$$\left| \Pr\left[b' = b \;\middle|\; \begin{array}{l} (\mu_0, \mu_1, \mathrm{pk}) \leftarrow \mathcal{A}(1^\lambda), \qquad b \xleftarrow{\$} \{0,1\} \\ (\sigma^b, \sigma^{1-b}) \leftarrow \langle \mathcal{A}, \mathrm{User}(\mathrm{pk}, \mu_b), \mathrm{User}(\mathrm{pk}, \mu_{1-b}) \rangle \\ \textit{if } \sigma^b = \perp \;\; \textit{or } \sigma^{1-b} = \perp, \;\; \textit{then } (\sigma^b, \sigma^{1-b}) \leftarrow (\perp, \perp) \\ b' \leftarrow \mathcal{A}(\sigma^b, \sigma^{1-b}) \end{array} \right] - \frac{1}{2} \right| \leq \mathrm{negl}(\lambda).$$

A blind signature scheme is said to satisfy perfect blindness *if the above advantage is exactly zero, i.e.,*

$$\left| \Pr\left[\mathcal{A}(\sigma^b, \sigma^{1-b}) = b' \right] - \frac{1}{2} \right| = 0.$$

Equivalently, the signer cannot distinguish which completed signature corresponds to which signing session, and its success probability is exactly that of random guessing.

Definition 5 (One-More Unforgeability). *A blind signature scheme is one-more unforgeable if, for every polynomial-time adversary $\mathcal{A}$ that makes at most Q signing queries, where Q is a polynomially bounded function of λ, there exists a negligible function* $\mathrm{negl}(\lambda)$ *such that:*

$$\Pr\left[\begin{array}{l} \forall\, 1 \leq i, j \leq Q+1 : \mu_i \neq \mu_j \\ \wedge\; \forall\, 1 \leq i \leq Q+1 : \\ \mathrm{Verify}(\mathrm{pk}, \mu_i, \sigma_i) = 1 \end{array} \;\middle|\; \begin{array}{l} (\mathrm{sk}, \mathrm{pk}) \leftarrow \mathrm{KeyGen}(1^\lambda), \\ \{(\mu_i, \sigma_i)\}_{i \in [Q+1]} \leftarrow \mathcal{A}^{\mathrm{Sign}(\mathrm{sk}, \cdot)}(\mathrm{pk}) \end{array} \right] < \mathrm{negl}(\lambda).$$

A.2 Adaptor Signature Definition

An adaptor signature scheme is expected to satisfy several security properties, including extractability, unique extractability, unlinkability, and pre-verification soundness [24]. These properties capture distinct aspects of correctness and security for pre-signatures and their corresponding adapted signatures. Although all four properties are generally desirable, some applications may require only a subset of them. We note that related definitions appear in the literature, for example in [4,16,19], with minor differences in formulation. Formal definition of these security requirements are as follow.

Definition 6 (Extractability). *Extractability guarantees that every valid signature output by an efficient adversary, after interacting with signing and pre-signing oracles, is either not a fresh signature or can be justified through the extraction of a valid witness from a previously issued pre-signature for the same message. Equivalently, an adversary should not be able to produce a fresh valid signature that cannot be associated, via the extraction procedure, with any pre-signature generated during the experiment.*

$\underline{\mathrm{Exp}^{\mathrm{ext}}_{\mathcal{A},AS}(1^\lambda)}$
1 : $(sk, pk) \leftarrow \mathrm{AS.KGen}(1^\lambda)$; $S \leftarrow \emptyset$; $T \leftarrow \emptyset$
2 : $(\mu^*, \sigma^*) \leftarrow \mathcal{A}^{\mathrm{NewY,Sign,pSign}}(pk)$
3 : if $\Sigma.\mathrm{Vrfy}(pk, \mu^*, \sigma^*) \neq 1$ then return 0
4 : if $\mu^* \in S$ then return 0
5 : for every $(M, \tilde{\sigma}) \in T[\mu^*]$ do
6 : $w \leftarrow \mathrm{AS.Extract}(pk, \tilde{\sigma}, \sigma^*, M)$
7 : if $(M, w) \in \mathrm{Rel}$ then return 0
8 : return 1.

$\underline{\mathrm{NewM}(1^\lambda)}$
1 : $(M, m) \leftarrow \mathrm{Rel.RGen}(1^\lambda)$
2 : return M

$\underline{\mathrm{Sign}(\mu)}$
1 : $\sigma \leftarrow \Sigma.\mathrm{Sign}(sk, \mu)$
2 : $S \leftarrow S \cup \{\mu\}$
3 : return σ

$\underline{\mathrm{pSign}(\mu, M)}$
1 : $\tilde{\sigma} \leftarrow \mathrm{AS.pSign}(sk, \mu, M)$
2 : $T[\mu] \leftarrow T[\mu] \cup \{(M, \tilde{\sigma})\}$
3 : return $\tilde{\sigma}$

Fig. 2. Extractability experiment and its associated oracles.

Formally, let S be the set of messages submitted to the signing oracle, and let T be a table containing all pre-signatures returned by the pre-signing oracle, organized according to their corresponding messages. Consider the experiment $\mathrm{Exp}^{\mathrm{ext}}_{\mathcal{A},AS}(1^\lambda)$ in Figure 2.

The adaptor signature scheme AS satisfies extractability if, for every probabilistic polynomial-time adversary $\mathcal{A}$, there exists a negligible function ν such that $\Pr\left[\mathrm{Exp}^{\mathrm{ext}}_{\mathcal{A},AS}(1^\lambda) = 1\right] \leq \nu(\lambda)$. Hence, except with negligible probability, no efficient adversary can output a fresh valid signature that is not extractable with respect to all relevant pre-signatures obtained during the experiment.

Definition 7 (Unique Extractability). *Unique extractability requires that a valid pre-signature have at most one valid completion that is meaningful under the extraction algorithm. More specifically, an efficient adversary should not be able to produce a message, a statement, a valid pre-signature, and two different valid signatures for the same message such that the extraction algorithm outputs valid witnesses from both signatures with respect to the same statement.*

In contrast to the extractability experiment, the adversary in this experiment is not given access to any oracle. It receives only the public key and attempts to construct a tuple that violates the uniqueness of the extraction process. Consider Unique Extractability experiment $\mathrm{Exp}^{\mathrm{uext}}_{\mathcal{A},AS}(1^\lambda)$ in Figure 3.

The adaptor signature scheme AS satisfies unique extractability if, for every probabilistic polynomial-time adversary $\mathcal{A}$, there exists a negligible function ν such that $\Pr\left[\mathrm{Exp}^{\mathrm{uext}}_{\mathcal{A},AS}(1^\lambda) = 1\right] \leq \nu(\lambda)$.

Equivalently, with all but negligible probability, no valid pre-signature can be completed into two distinct valid signatures such that both signatures lead to valid witnesses through the extraction algorithm.

$\underline{\mathrm{Exp}^{\mathrm{uext}}_{\mathcal{A},AS}(1^\lambda)}$
1 : $(sk, pk) \leftarrow AS.\mathrm{KGen}(1^\lambda)$
2 : $(\mu, M, \tilde{\sigma}, \sigma_0, \sigma_1) \leftarrow \mathcal{A}(pk)$
3 : if $AS.\mathrm{pVrfy}(pk, \mu, \tilde{\sigma}, M) \neq 1$ then return 0
4 : if $\sigma_0 = \sigma_1$ then return 0
5 : if $\Sigma.\mathrm{Vrfy}(pk, \mu, \sigma_0) \neq 1$ then return 0
6 : if $\Sigma.\mathrm{Vrfy}(pk, \mu, \sigma_1) \neq 1$ then return 0
7 : $w_0 \leftarrow AS.\mathrm{Extract}(M, \tilde{\sigma}, \sigma_0)$
8 : $w_1 \leftarrow AS.\mathrm{Extract}(M, \tilde{\sigma}, \sigma_1)$
9 : if $(M, w_0) \in \mathrm{Rel}$ and $(M, w_1) \in \mathrm{Rel}$ then return 1
10 : return 0.

Fig. 3. Unique Extractability experiment

$\underline{\mathrm{Exp}^{\mathrm{unlink}}_{\mathcal{A},AS}(1^\lambda, b)}$
1 : $(sk, pk) \leftarrow AS.\mathrm{KGen}(1^\lambda)$
2 : $b' \leftarrow \mathcal{A}^{\mathrm{Chall}_b, \mathrm{Sign}, \mathrm{pSign}}(pk)$
3 : return b'

$\underline{\mathrm{Chall}_b(\mu, (M, m))}$
1 : if $(M, m) \notin \mathrm{Rel}$ then return $\perp$
2 : $\tilde{\sigma} \leftarrow AS.\mathrm{pSign}(sk, \mu, M)$
3 : $\sigma_0 \leftarrow AS.\mathrm{Adapt}(pk, \tilde{\sigma}, m)$
4 : $\sigma_1 \leftarrow \Sigma.\mathrm{Sign}(sk, \mu)$
5 : return σ_b

$\underline{\mathrm{Sign}(\mu)}$
1 : $\sigma \leftarrow \Sigma.\mathrm{Sign}(sk, \mu)$
2 : return σ

$\underline{\mathrm{pSign}(\mu, M)}$
1 : $\tilde{\sigma} \leftarrow AS.\mathrm{pSign}(sk, \mu, M)$
2 : return $\tilde{\sigma}$

Fig. 4. Unlinkability experiment and its associated oracles.

Definition 8 (Unlinkability). *Unlinkability requires that a signature obtained by adapting a pre-signature be computationally indistinguishable from a freshly generated ordinary signature on the same message. This property is particularly important in applications such as atomic swaps, where revealing whether a signature was derived from a pre-signature may compromise privacy.*

The experiment is parameterized by a hidden bit $b \in \{0, 1\}$. If $b = 0$, the challenge oracle returns an adapted signature; if $b = 1$, it returns a freshly generated ordinary signature. The adversary is given access to the challenge, signing, and pre-signing oracles and attempts to distinguish which type of signature is returned by the challenge oracle. Consider the experiment $\mathrm{Exp}^{\mathrm{unlink}}_{\mathcal{A},AS}(1^\lambda, b)$ as illustrated in Figure 4. The challenge oracle takes as input a message μ and a statement-witness pair (M, m). It first checks that $M, m)$ satisfies the relation Rel. If the check fails, the oracle returns $\perp$. Otherwise, it generates both an adapted signature and a fresh ordinary signature on the same message, and re-

$\underline{\mathrm{Exp}^{\mathrm{pvs}}_{\mathcal{A},AS}(1^\lambda, M)}$
1 : $(sk, pk) \leftarrow AS.\mathrm{KGen}(1^\lambda)$
2 : $(\mu, \tilde{\sigma}) \leftarrow \mathcal{A}(sk, M)$
3 : if $AS.\mathrm{pVrfy}(pk, \mu, \tilde{\sigma}, M) = 1$ then return 1
4 : return 0.

Fig. 5. Pre-Verification soundness experiment

turns one of them according to the hidden bit b. The adaptor signature scheme AS satisfies unlinkability if, for every probabilistic polynomial-time adversary $\mathcal{A}$, there exists a negligible function ν such that

$$\left|\Pr\left[\mathrm{Exp}^{\mathrm{unlink}}_{\mathcal{A},AS}(1^\lambda, 0) = 1\right] - \Pr\left[\mathrm{Exp}^{\mathrm{unlink}}_{\mathcal{A},AS}(1^\lambda, 1) = 1\right]\right| \leq \nu(\lambda).$$

Equivalently, even when the adversary chooses the witness used for adaptation, it cannot distinguish, except with negligible advantage, whether the resulting valid signature was obtained by adapting a pre-signature or by running the ordinary signing algorithm directly.

Definition 9 (Pre-Verification Soundness). *Pre-verification soundness requires that the pre-verification algorithm be sound with respect to the relation* Rel. *In particular, a pre-signature generated with respect to a statement $M \notin L_{\mathrm{Rel}}$ should not be accepted by the pre-verification algorithm. This property prevents a malicious pre-signer from convincing a verifier to accept a pre-signature associated with a false statement, namely a statement for which no valid witness exists.*

Computational pre-verification soundness. Computational pre-verification soundness requires that no efficient adversary can produce an accepting pre-signature for a statement outside the language L_{Rel}, even when the adversary is given the secret signing key. More formally, let $M \notin L_{\mathrm{Rel}}$ be a statement whose length is polynomially bounded in the security parameter. Figure 5 shows the experiment $\mathrm{Exp}^{\mathrm{pvs}}_{\mathcal{A},AS}(1^\lambda, M)$. The adaptor signature scheme AS satisfies computational pre-verification soundness if, for every probabilistic polynomial-time adversary $\mathcal{A}$ and every polynomially bounded statement $M \notin L_{\mathrm{Rel}}$, there exists a negligible function ν such that

$$\Pr\left[\mathrm{Exp}^{\mathrm{pvs}}_{\mathcal{A},AS}(1^\lambda, M) = 1\right] \leq \nu(\lambda).$$

Equivalently, even an adversary that knows the secret signing key cannot generate, except with negligible probability, a pre-signature that is accepted with respect to a statement outside the relation language.

Statistical pre-verification soundness. Statistical pre-verification soundness is a stronger notion. It requires that accepting pre-signatures for false statements do not exist, regardless of the computational power of the adversary. The scheme AS satisfies statistical pre-verification soundness if, for every security parameter $\lambda \in \mathbb{N}$, every polynomially bounded statement $M \notin L_{\mathrm{Rel}}$, every key pair (sk, pk) in the support of $AS.\mathrm{KGen}(1^\lambda)$, every message μ, and every pre-signature $\tilde{\sigma}$, it holds that

$$AS.\mathrm{pVrfy}(pk, \mu, \tilde{\sigma}, M) \neq 1.$$

Equivalently, for any statement outside the language L_{Rel}, no message and no pre-signature can make the pre-verification algorithm accept.

A.3 Probability Distributions

The proposed scheme employs two probability distributions: the discrete Gaussian distribution for rejection sampling, and the binomial distribution.

Definition 10 (Discrete Gaussian Distribution). *The discrete Gaussian distribution over $\mathcal{R}^\ell$ centered at $\boldsymbol{v} \in \mathcal{R}^\ell$ with standard deviation $\mathcal{S} > 0$ is defined as:*

$$D^\ell_{\boldsymbol{v},\mathcal{S}}(\boldsymbol{z}) = \frac{e^{-\|\boldsymbol{z}-\boldsymbol{v}\|^2/(2\mathcal{S}^2)}}{\sum\limits_{\boldsymbol{z'}\in\mathcal{R}^\ell} e^{-\|\boldsymbol{z'}\|^2/(2\mathcal{S}^2)}}.$$

When the distribution is centered at the origin $\mathbf{0} \in \mathcal{R}^\ell$, we write $D^\ell_{\mathcal{S}} = D^\ell_{\mathbf{0},\mathcal{S}}$.

Lemma 1 (Tail Bound). *Let $\boldsymbol{z} \leftarrow D^m_{\mathcal{S}}$. Then:*

$$\Pr\left[\, \|\boldsymbol{z}\| > t \cdot \mathcal{S}\sqrt{md} \,\right] < \left(t \cdot e^{\frac{1-t^2}{2}}\right)^{md}.$$

For a Gaussian vector $\mathbf{z} \in \mathcal{R}^m$, the probability that its length exceeds t times the typical length $\mathcal{S}\sqrt{m}$ decreases exponentially with m, so unusually long vectors are extremely unlikely [6, Lemma 1.5(i)].

Definition 11 (Binomial Distribution). *The binomial distribution with parameter $\varkappa > 0$, denoted as $\mathrm{Bin}_\varkappa$, is defined as: $\mathrm{Bin}_\varkappa = \sum_{i=1}^{\varkappa}(a_i - b_i)$, where each $a_i, b_i \leftarrow \{0,1\}$ uniformly at random. The variance of $\mathrm{Bin}_\varkappa$ is $\varkappa/2$, and it satisfies the additive property: $\mathrm{Bin}_{\varkappa_1} \pm \mathrm{Bin}_{\varkappa_2} = \mathrm{Bin}_{\varkappa_1+\varkappa_2}$.*