BABE: Verifying Proofs on Bitcoin Made 1000x Cheaper

Sanjam Garg[1], Dimitris Kolonelos[1], Mikhail Sergeevitch[3], Srivatsan Sridhar[*4,1], and David Tse[4,5]

[1]University of California, Berkeley
[3]Babylon Labs
[4]Byzantine Research
[5]Stanford University

Abstract

Endowing Bitcoin with the ability to verify succinct proofs has been a longstanding problem with important applications such as scaling Bitcoin and allowing the Bitcoin asset to be used in other blockchains trustlessly. It is a challenging problem due to the lack of expressiveness in the Bitcoin scripting language and the small Bitcoin block space. BitVM2 [LAA+25] is the state-of-the-art verification protocol for Bitcoin used in several mainnets and testnets [Bit25a, Cit25, BOB25a], but it suffers from very high on-chain Bitcoin transaction fees in the unhappy path (over $\$14,000$ in a recent experiment [LAA+25]). Recent research BitVM3 dramatically reduces this on-chain cost by using a garbled SNARK verifier circuit to shift most of the verification off-chain [Rub24, Lin24], but each garbled circuit is 42 GiBytes in size, so the off-chain storage and setup costs are huge. This paper introduces BABE, a new proof verification protocol on Bitcoin, which preserves BitVM3's savings of on-chain costs but reduces its off-chain storage and setup costs by three orders-of-magnitude. BABE uses a witness encryption scheme for linear pairing relations [GKPW24] to verify Groth16 proofs. Since Groth16 verification involves non-linear pairings, this witness encryption scheme is augmented with a secure two-party computation protocol implemented using a very efficient garbled circuit for scalar multiplication on elliptic curves. The design of this garbled circuit builds on the recent work of Argo MAC [EL26], a garbling primitive that efficiently computes homomorphic MACs on such curves.

[*]Work done while at Byzantine Research, currently at University of California, Berkeley.

Contents

1 Introduction **4**
- 1.1 Motivating Example . . . 4
- 1.2 Verifying Proofs on Bitcoin . . . 4
  - 1.2.1 BitVM and BitVM2 . . . 4
  - 1.2.2 BitVM3 . . . 5
- 1.3 New Verification Protocol: BABE . . . 6
  - 1.3.1 Witness Encryption as a Starting Point . . . 6
  - 1.3.2 Witness encryption for linear pairings . . . 7
  - 1.3.3 Garbled Circuit for Scalar Multiplication . . . 9
- 1.4 Other Applications . . . 10
- 1.5 Paper Outline . . . 10

2 Preliminaries **10**
- 2.1 Basic Notation . . . 10
- 2.2 Bilinear Groups . . . 11
  - 2.2.1 Generic Bilinear Group Model . . . 11
- 2.3 Succinct Non-Interactive Arguments of Knowledge (SNARKs) . . . 12
  - 2.3.1 The Groth16 SNARK . . . 12
- 2.4 Garbling Schemes . . . 13
- 2.5 The Bitcoin Ledger . . . 14

3 The BitVM-core Primitive **18**

4 Witness Encryption for Linear Pairing Relation **20**

5 Garbled Circuit for BN254 Scalar Multiplication **23**
- 5.1 Overview . . . 23
- 5.2 Elliptic Curve Addition and Notation . . . 26
- 5.3 Decomposable Randomized Encodings: Definitions and Preliminaries . . . 28
- 5.4 Decomposable Randomized Encodings Constructions . . . 29
- 5.5 DRE for the Scalar Multiplication . . . 30
- 5.6 Completing the Garbled Circuit . . . 32

6 BABE Protocol **34**
- 6.1 Honest Setup Protocol . . . 34
  - 6.1.1 Setup Phase . . . 34
  - 6.1.2 Proving Phase . . . 35
  - 6.1.3 Transaction Graph . . . 38
- 6.2 Verifying Setup Correctness . . . 39

7 Security Proof **40**
- 7.1 Security Proof Assuming Honest Setup . . . 40
  - 7.1.1 Proof of Honest-Setup u-Robustness . . . 40
  - 7.1.2 Proof of Honest-Setup Knowledge Soundness . . . 42

8 Extensions and Optimizations **46**
8.1 Multiple Verifiers and Provers . . . 46
8.2 Optimistic Path . . . 46

9 Evaluation **46**
9.1 Honest Setup . . . 46
9.2 Cut-and-Choose Setup Verification . . . 48
9.3 Soldering (zk-SNARK-soldering) . . . 49
9.4 Verifiable Shamir Secret Sharing . . . 50

A Honest-Setup BABE Protocol Details **57**
A.1 Transactions . . . 57
A.1.1 Notation . . . 58
A.1.2 Example: The Assert Transaction . . . 58

B Protocol for Malicious Security **60**

1 Introduction

1.1 Motivating Example

Bob has 1 BTC and would like to use it as collateral to borrow \$50K worth of stablecoins for real-world use. Currently Bob's only options are to go to a centralized entity such as Tether to directly borrow, or to a centralized entity like Coinbase which mints Bob a wrapped asset cbBTC that he can use in a smart contract lending protocol on Ethereum or other chains. However, in both cases the centralized entity would take custody of Bob's BTC, violating the core ethos of Bitcoin: *trustlessness*.[1] If Bob had 1 ETH, on the other hand, he would be able to participate directly in a trustless lending protocol like Aave on Ethereum without going through a centralized entity. Unfortunately, the Bitcoin chain does not support smart contract protocols such as Aave. This begs the question:
How do we allow the \$1.4 *trillion worth of BTC to participate trustlessly in smart contracts like Aave on chains like Ethereum?*

Consider what such a protocol would look like:

1. Borrower Bob deposits his BTC using a Bitcoin transaction;
2. A lending position is created on Ethereum with the BTC as collateral and \$50K in stablecoin is lent from a lender Larry to Bob;
3. If Bob returns the loan, then the BTC collateral can be withdrawn by Bob;
4. If Larry liquidates the loan because BTC price has dropped below a threshold, then the BTC collateral can be withdrawn by Larry.

To achieve step 2 trustlessly, Bob has to prove a certain state of the Bitcoin chain to Ethereum. This proof can be verified by smart contracts which already exist on Ethereum [Cat25]. To achieve steps 3 and 4 trustlessly, each party has to prove a certain state of the Ethereum chain to Bitcoin so that he can withdraw. Due to the rudimentary nature of the Bitcoin scripting language, however, even a succinct proof verifier takes about 900 Mbytes of Bitcoin block space, equivalent to 225 Bitcoin blocks [LAA+25]. Moreover, this verifier has to be submitted on-chain as a Bitcoin transaction every time a proof needs to be verified, thus rendering it totally impractical (the transaction fees will be millions of dollars).

Note that the lending example is only one application. The ability to verify proofs is important for any application in which the withdrawal of the deposited BTC depends on the state of another chain, such as light-client based Bitcoin bridges to rollups and other chains [LAA+25].

1.2 Verifying Proofs on Bitcoin

1.2.1 BitVM and BitVM2

BitVM [Lin23, AAL+24] initiated a line of work to reduce the on-chain footprint of proof verification using optimistic methods. It is a two-party protocol between the Prover and the Verifier, with the Verifier helping Bitcoin verify the Prover's proof by challenging the Prover if its proof is invalid. This protocol ends if either the Prover wins, i.e. convinces Bitcoin its proof is correct, or the Verifier wins, i.e. convinces Bitcoin the Prover's proof is incorrect. This framework can be applied to our motivating example: when Bob tries to withdraw, he is the Prover, with Larry being the Verifier. Bob wins if he succeeds in withdrawing, and Larry wins if he succeeds in stopping Bob from withdrawing. (We formalize the BitVM setting in Sec. 3 and use it to prove security in this paper.)

[1] In Satoshi Nakamoto's initial post introducing Bitcoin to the world, he described it as: "It's completely decentralized, without relying on central servers or trusted parties, because everything is based on crypto proofs rather than trust." [Nak09]

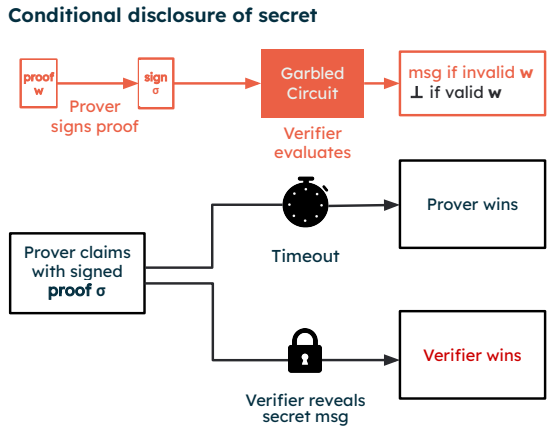

(a) Verification on Bitcoin using a SNARK verifier garbled circuit [WAA+26]. The *Prover encrypts* a secret msg which the *Verifier decrypts* iff the Prover signs an *invalid* SNARK proof w. Decryption is interactive: the Prover posts a Lamport signature σ on his SNARK proof w to Bitcoin. Upon seeing this, the Verifier uses σ as the input labels to evaluate the garbled circuit, which implements the SNARK verifier. If w is an invalid SNARK proof, the evaluation reveals the secret msg, which the Verifier posts on Bitcoin to open the hash lock and win; otherwise, after a timelock, the Prover wins.

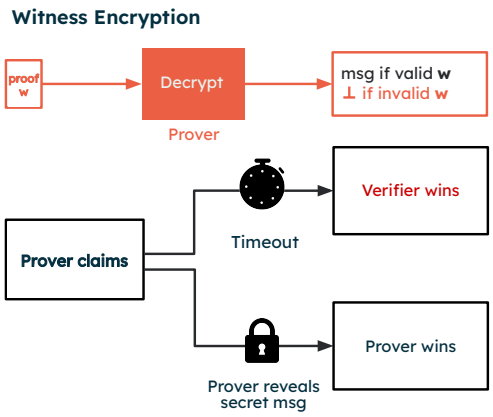

(b) Verification on Bitcoin using witness encryption (if it were practical). The roles and decryption condition are reversed: the *Verifier encrypts* a secret msg and the *Prover decrypts* when he knows a *valid* SNARK proof w. Decryption is non-interactive: the Prover computes the SNARK proof w, then runs the decryption algorithm of the witness encryption scheme. If w is an valid SNARK proof, decryption reveals the secret msg, which the Prover posts on Bitcoin to open the hash lock and win; otherwise, after a timelock, the Verifier wins.

Figure 1: Verification on Bitcoin using a SNARK verifier garbled circuit (a) versus witness encryption (b). Both approaches are solutions to the conditional disclosure of secrets problem [GIKM00].

The state-of-the-art protocol in that line of work is BitVM2 [LAA+25]; this protocol follows the framework of naysayer proofs [SGB24] where the Prover posts the intermediate states of the proof verification computation trace on-chain and the Verifier can challenge a specific state transition by posting the correct verification trace only for that state transition. Even though this drastically reduces the on-chain costs as compared to posting the entire verification trace, the on-chain cost is still substantial as the intermediate states have to be signed by Lamport signatures [Lam79], verifiable on Bitcoin using hashlocks. (Lamport signatures are very large, requiring 256 bits per bit of each intermediate state.) In a mainnet experiment on Bitcoin [LAA+25], the on-chain footprint of posting these dispute transactions is over 5.4 Mbytes, costing over $\$15,000$ of transaction fees. This high challenge cost necessitates large amounts of capital locked as a bond to pay for the challenge. Moreover, the Bitcoin transactions involved in the challenge are non-standard[2] and require special services for submission to Bitcoin, thus increasing fee rates and centralization risks.

1.2.2 BitVM3

To reduce the on-chain costs of BitVM2, [Rub24] suggested using garbled circuits (GC) to shift the proof verification off-chain. Follow-up works include [Lin24, Che25, Eag25] under the general umbrella of BitVM3. Garbled circuits, originally invented by Yao [Yao82] for private two-party computation, are used in BitVM3 for *authenticated* computation. In the two-party setting, it works as follows (Fig. 1a). At setup,

[2]Non-standard transactions are valid as per Bitcoin consensus rules but most miners do not consider them for inclusion in the blocks they mine because, for example, the transactions are too large.

the Prover generates a secret msg and shares a garbled circuit encrypting msg with the Verifier. The garbled circuit implements a SNARK verifier and can be used to decrypt the secret if an invalid proof is input into the garbled circuit. The Prover commits to the hashed value $H(\mathsf{msg})$ to create a hash lock such that the Verifier can open the hash lock and stop the Prover's claim if the Verifier can learn msg. The Verifier stores the garbled circuit. At the proving phase, the Prover posts a Lamport signed proof on Bitcoin, and the Verifier feeds the signed proof as input into the garbled circuit, with the Lamport signatures as input labels.[3] So if the proof is invalid, the Verifier will be able to stop the Prover by decrypting the scret and use it to open the hashlock. Otherwise, after a timeout, the Prover will win. This is basically a hash-time lock contract (HTLC) [Wik21] used in many Bitcoin protocols including atomic swaps [Wik20, Her18], Lightning [PD16], and Bitcoin staking [DLT+24]. HTLC had earlier been generalized to verification of signatures by discreet log contracts [Dry17], and one can view BitVM3 as a final generalization to verification of arbitrary computations using the SNARK verifier garbled circuit.

The on-chain cost of BitVM3 compared to BitVM2 is reduced from $\$14,000$ to less than $\$40$ because only the signed proof has to be put on-chain instead of the signed intermediate states, and a hashlock script of negligible size replaces the verification traces. The challenge for BitVM3 is the significant *off-chain* costs. A Boolean garbled circuit for the Groth16 verifier [Gro16] was implemented [Bit25b, WAA+26], and even with free-XOR [KS08] and privacy-free half-gate [ZRE15] optimizations, the garbled circuit size is still 42 Gibytes (total of 10 B gates and 3 B non-free gates). Garbling time per circuit is 6 minutes on a single core. Moreover, proving the correctness of the garbled circuit with a zero-knowledge proof is infeasible and with a standard cut-and-choose method [LP07] requires garbling many instances of the circuits, translating to a total setup time of hours.

The large size of the garbled circuit for the Groth16 verifier primarily stems from the expensive pairing operations on the BN254 curve [BN06]. By replacing the Groth16 by a designated-verifier SNARK and by replacing the BN254 curve with a curve on the binary field, [Eag25] shows that the garbled circuit size can be reduced to 12 Mbytes. Unfortunately, the security of binary curves is not widely accepted by the community and so this approach is not currently pursued in practice. Arithmetic circuits were proposed recently [FBFL25] for garbling the Groth16 verifier, but the size is still expected to be 100 million ciphertexts (a few GBs).

1.3 New Verification Protocol: BABE

In this paper, we introduce a new verification protocol, BABE, which verifies Groth16 proofs (on the standard BN254 curve) on Bitcoin. BABE improves the off-chain costs of the existing Groth16 verifier garbled circuit [Bit25b, WAA+26] in BitVM3 by a factor of more than 1000 while preserving its low on-chain costs relative to BitVM2 (Fig. 2).

Instead of using a garbled circuit for the Groth16 verifier, our approach uses witness encryption [GGSW13] as a starting point, but sidesteps the difficulty of general-purpose witness encryption by using interactivity. In simple terms, a witness encryption (WE) scheme for a relation R allows one to encrypt a secret under an NP statement $\mathbb{x} \in \mathcal{R}$ such that anyone that knows the corresponding witness $\mathbb{w}$ can decrypt.[4]

1.3.1 Witness Encryption as a Starting Point

Let $\mathcal{R}_G$ be the relation for Groth16 verification, i.e. $\mathbb{x}$ is the statement (public inputs), $\mathbb{w}$ is the Groth16 proof (witness of this relation), and $(\mathbb{x}, \mathbb{w}) \in \mathcal{R}_G$ is true iff the proof $\mathbb{w}$ is valid for the statement $\mathbb{x}$. If we had a solution for WE on $\mathcal{R}_G$, then we could use that to solve the Prover-Verifier problem above (Fig. 1b).

[3] In an uncanny coincidence, Yao's input labels correspond exactly to a Lamport signature and vice-versa.

[4] Technically, this describes a stronger notion of witness encryption called "extractable witness encryption" [GKP+13]. For the sake of this overview we omit this delicate difference; we elaborate in the technical sections.

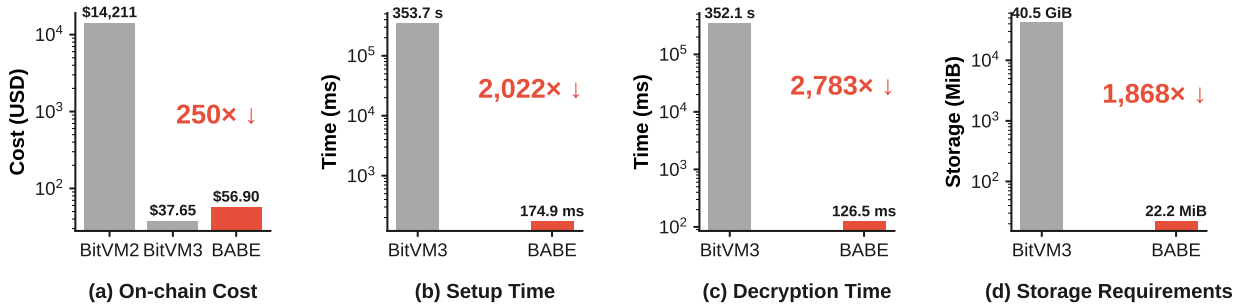

CPU: AMD Ryzen 7 7840U; timing: single-threaded. BTC price: $95,500 (Jan 3, 2026). Fee rate: 2.2sats/vB for BitVM3, BABE, 11sats/vB for BitVM2 (premium rate for non-standard transactions). Details: Sec. 9.

Figure 2: On-chain and off-chain costs of BitVM2 [LAA+25], BitVM3 (Boolean garbled circuit for Groth16 verification) [Bit25b, WAA+26], and BABE. On-chain costs are derived from experiments on Bitcoin mainnet. These metrics are in the honest-setup setting where the Prover and Verifier are assumed to be honest during setup (but can be malicious afterwards). Sec. 9 evaluates multiple methods for achieving malicious security.

In this hypothetical solution, the Verifier plays the role of the WE-encryptor. At setup, it encrypts a secret msg using the statement $\mathbb{x}$ to generate ciphertext ct. It creates a hash lock by committing to the hashed value $H(\mathsf{msg})$ of the secret msg. The Prover plays the role of the WE-decryptor and stores the ciphertext ct. At the proving phase, the Prover obtains the proof $\mathbb{w}$ of the application-specific relation, and runs $\mathsf{Dec}(\mathsf{ct}, \mathbb{w})$ to decrypt the secret msg. This is used to open the hash lock to win.

1.3.2 Witness encryption for linear pairings

Even before the advent of BitVM, witness encryption had already been proposed as a technique to perform trustless verification on Bitcoin (in the context of building Bitcoin bridges) [Hio22]. Unfortunately, witness encryption for all relations is a notoriously challenging cryptographic primitive, and to date it lacks an efficient realization for general relations.

Specifically, the Groth16 relation $\mathcal{R}_G$ involves verifying an equation containing pairing terms:

$$e(\pi_1, \pi_2) + e(\pi_3, x_2) = x_3 \tag{1}$$

Here, $\mathbb{w} = (\pi_1, \pi_2, \pi_3)$ is the proof, $\mathbb{x} = (x_2, x_3)$ is the statement. The issue is the pairing between the two witness elements ($e(\pi_1, \pi_2)$): There is no known efficient construction of WE for this type of relation.

However, we do know [BC16, GKPW24] an efficient WE scheme for relations that are pairing equations *linear* in the witness $\mathbb{w}$. Consider for example, the linear pairing relation $e(x_1, w_1) + e(w_2, x_2) = x_3$. Then we can obtain a WE scheme as follows:[5]

- **Relation**: $\mathcal{R}_{\text{linear}} = \left\{(\mathbb{x} = (x_1, x_2, x_3); \mathbb{w} = (w_1, w_2)) : e(x_1, w_1) + e(w_2, x_2) = x_3\right\}$
- **Encryption** (for $\mathbb{x}$): $\mathsf{ct} = (rx_1, rx_2, rx_3 + \mathsf{msg})$
- **Decryption** (using $\mathbb{w}$): $\mathsf{msg} = \mathsf{ct}_3 - e(w_2, \mathsf{ct}_2) - e(\mathsf{ct}_1, w_1)$

where r is an additional private uniformly random field element generated by the WE-encryptor. This simple yet powerful observation has been leveraged to build efficient witness encryption schemes for

[5]Formally, for the security proof we need the "masking" term of the ciphertext to be passed through a random oracle, i.e. $\mathsf{RO}(rx_3)$. For simplicity, we omit this technicality from this overview and refer to Sec. 4.

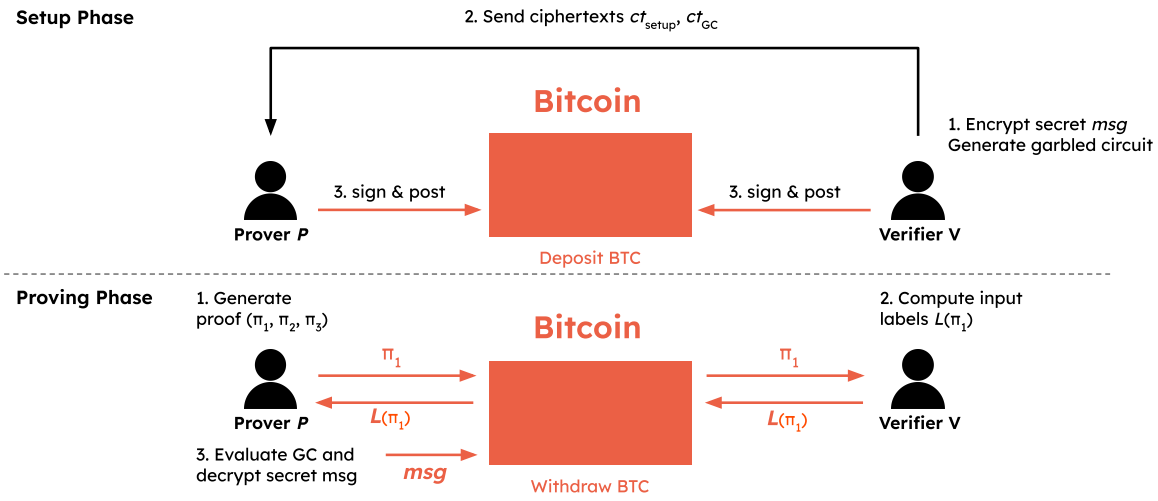

Figure 3: BABE: Verification of arbitrary computation on Bitcoin using a combination of linear witness encryption and an interactive protocol that allows the Prover to compute $r\pi_1$ without knowing r. The protocol is not a witness encryption scheme for general relations, but a solution to conditional disclosure of secrets (CDS) [GIKM00] in which decryption is interactive.

specific relations across a wide range of cryptographic settings, both implicitly (e.g. [BL20, FKdP23, CFK24, FHAS24]) and explicitly following [GKPW24] (e.g. [CGPW25, AFP25, BFOQ25, MLLP25]), and has been recently formalized in [GHK+25]. To adapt this WE scheme for Groth16 verification, we treat one of the proof elements, π_1, as a public input x_1 of the WE relation so that the Groth16 relation (Eq. (1)) turns into:

$$e(x_1, \pi_2) + e(\pi_3, x_2) = x_3, \tag{2}$$

a linear pairing relation in the witness (π_2, π_3) and public inputs (x_1, x_2, x_3). Following the aforementioned framework for linear pairing equations, the ciphertext is given by:

$$\mathsf{ct} = (rx_1, rx_2, rx_3 + \mathsf{msg}). \tag{3}$$

But in reality, $x_1 := \pi_1$ is not available during setup, which means the ciphertext (Eq. (3)) cannot be entirely computed during setup; the computation must be deferred to the proving phase. In particular, the Prover must be able to compute $r\pi_1$ from π_1 in the proving phase without knowing the private randomness r of the Verifier.

To achieve this, the Prover and the Verifier will set up a maliciously secure two-party computation protocol (Fig. 3). This two-party computation takes as the Verifier's input the private randomness r at setup and as the Prover's input the proof element π_1 in the proving phase, and either the Prover learns $r\pi_1$ without learning r, or the protocol aborts and the Prover wins.

Concretely, during setup, the Verifier creates a garbled circuit ct_{GC} that hard codes the randomness r and shares it with the Prover. This garbled circuit computes $r\pi_1$ given input labels for π_1 with r private from the Prover. During the proving phase, when the Prover has the Groth16 proof, the Prover reveals π_1 on-chain and obtains the input label $L(\pi_1)$ from the Verifier. If the Verifier responds, then the Prover can compute $r\pi_1$ by inputting the labels into the garbled circuit. If the Verifier does not respond within a timeout or does not send the correct labels for the Prover's π_1, the protocol halts and the Prover wins.

Once the Prover has learned $r\pi_1$, together with $\mathsf{ct}_{setup} := (rx_2, rx_3 + \mathsf{msg})$ which it obtained from the Verifier at setup, it has the entire ciphertext (Eq. (3)). Now the Prover can use the rest of the proof (π_2, π_3) to decrypt the secret msg as in linear witness encryption and wins, provided that the proof is valid.

1.3.3 Garbled Circuit for Scalar Multiplication

In BitVM3, the garbled circuit (GC) allows authentic evaluation of a Groth16 verifier to reveal a secret. In BABE, the garbled circuit allows evaluation of a scalar multiplication $r\pi_1$ with privacy on r. So the garbled circuits play totally different roles in the two protocols. But more importantly, computing a single scalar multiplication is much simpler than computing a Groth16 verifier, which involves multiple pairings. More specifically, our goal is to design an efficient garbled circuit that

- has a private input $r \in \mathbb{F}_p$ from the Verifier (fixed at setup phase),
- takes as input labels a Lamport signature on the Prover's public input $\pi_1 = (x, y) \in \mathbb{F}_p^2$ (Bitcoin friendly).
- Outputs $r\pi_1$.

A natural first approach is to leverage Yao's garbling [Yao82] for the scalar multiplication $r\pi_1$. While this approach offers a substantial $\approx 5\times$ improvement over the Groth16 verifier garbled circuit (BitVM3), the resulting solution would still be on the order of Gbytes. Essentially, the cost comes from the fact that one has to encode large ($\mathbb{F}_p$, where $\log(p) \approx 254$) field operations into binary circuits,[6] which introduces unnecessary overhead.

In this direction, a recent innovation [EL26] advocates an approach to keep most of the operations of the garbled circuit at the level of the relevant elliptic curve groups. Specifically, the paper introduces a new garbling primitive that computes vector homomorphic MACs (HMACs) on the BN254 elliptic curve. Such homomorphic MACs allow free additions of BN254 group elements, analogous to free-XORs for bits. The main result is an efficient garbling method to compute the components of the vector HMAC of a group element G, each of which effectively involves adding a fixed group element to G. The construction directly works over $\mathbb{F}_p$ and builds on the Ishai-Wee partial garbling for branching programs [IW14], which directly treats $\mathbb{F}_p$-elements without decomposing them to binary. Their main observation is that for BN254, a group addition to a fixed group element can be expressed as a low-degree polynomial f, such that $\pi_1 + \phi = f_\phi(x, y)$ where x, y are the x and y coordinates of π_1, (and thus a branching program) and garbled with [IW14]. Notably, the Ishai-Wee garbled circuit is information-theoretic and free, meaning that the garbled circuit has zero size, and the only cost comes from the encoded input (i.e. the labels). Crucially, the Ishai-Wee labels are not Bitcoin-friendly (i.e. Lamport signatures) as is the case for Yao. Therefore, the proposal suggests a pre-processing garbled circuit to compute the Ishai-Wee labels from the Lamport signature on π_1.

Building on this work, we develop a highly efficient garbling scheme for BN254 scalar multiplications. First, using the technique of decomposable randomized encodings [IK00, IK02, AIK04, Ish13] for linear operations, we decompose the BN254 scalar multiplication problem into the problem of computing a 254-dimensional vector HMAC. Second, to garble each component, we again rely on DRE for linear operations rather than the arguably more complex branching programs technique of Ishai-Wee used in [EL26]. Our core observation is that we can 'linearize' the BN254 group addition function $f_\phi(x, y)$ to $\widetilde{f}_\phi(x, y, x^2, y^2, xy)$, by simply giving low-degree monomials as inputs. Now our $\widetilde{f}_\phi$ is a linear function. Then the operation boils down to an inner-product, amenable to the [Ish13] DRE. The computation of the low-degree monomials is now deferred to the Yao garbled circuit, which turns out to be of low cost.

Comparing to [EL26], our technique is arguably simpler, allowing for a full security proof derived directly from the well-established decomposable randomized encodings for linear relations [FKN94, Ish13] (and a composition theorem within). Our final GC construction (after some optimizations) is about 22 MiB.

[6] Recall that Yao's garbling works on binary circuits.

1.4 Other Applications

In addition to Bitcoin, several other blockchains (e.g. Cardano) do not natively support verification of general-purpose proofs. In our system BABE, the bulk of the protocol runs off-chain, and is therefore largely blockchain-agnostic. The on-chain component requires only very basic functionality including hashlocks, timelocks, and signature verification, which most blockchains already support. As a result, BABE or a system built on top of it, could be deployed across a broad class of chains, unlocking a wide range of applications beyond Bitcoin.

From a separate perspective, this work speaks to a broader usability bottleneck in cryptography: *non-black-box techniques* are often considered impractical because they require "opening up" cryptographic algorithms and expressing them as circuits, which can incur substantial overhead. A canonical example comes from zkSNARKs, where proving statements about cryptographic computations requires unrolling those computations into circuits, something that was historically seen as prohibitive. Over time, however, modern SNARK constructions and implementations have shown that these efficiency barriers can be overcome and that doing so can have far-reaching practical impact [BCTV14, OWWB20, BCMS20, BMM+21, GMN22, KST22, CFH+22, GGW24, OKMZ25, WOS+25, GGKS25].

Secure computation (e.g., via garbled circuits) faces a closely related challenge today. Despite significant progress, general-purpose garbling still performs poorly on many cryptographic workloads and other non-black-box computations, which remain expensive in practice. In this work, we focus on a highly specialized cryptographic task of real-world concern and demonstrate practical performance. Beyond the immediate application, this suggests a broader takeaway: with the right specialization, non-black-box uses of cryptography can be made practical for secure computation tasks, potentially enabling new deployments in a variety of settings.

1.5 Paper Outline

The rest of the paper is organized as follows. Sec. 2 describes the basic notations and cryptographic concepts used in this paper. Sec. 3 formalizes the security model for proof verification on Bitcoin. Sec. 4 presents the witness encryption protocol for linear pairing relations that we use. Sec. 5 describes our efficient garbled circuit construction for scalar multiplication. Sec. 6 combines these components into the full BABE protocol. Sec. 8 discusses several optimizations. Sec. 9 presents experimental results comparing BABE to prior work. Mathematical notation, preliminary background, and security proofs are in the appendix.

2 Preliminaries

2.1 Basic Notation

We use λ for the security parameter and $\mathsf{negl}(\lambda)$ for a negligible function, i.e., a function that is less than $1/f(\lambda)$ for *any* polynomial f. We also define κ to be a statistical security parameter. The security parameters are implicitly taken as input to every algorithm and, for brevity, we omit explicitly writing it. For events A and B, we let $\Pr[A] \approx_\epsilon \Pr[B]$ denote that $|\Pr[A] - \Pr[B]| \leqslant \epsilon$. Row vectors will be written in small bold font, e.g. $\boldsymbol{x} = (x_1, \ldots, x_n)$ and matrices in capital bold, e.g. $\boldsymbol{A} = (a_{ij})_{i,j}$. We use the operator '$\times$' for the matrix multiplication and the operator $\otimes$ for the tensor product. $x \leftarrow\!\!\$\ \mathcal{X}$ is used to imply that x is being uniformly sampled from a finite set $\mathcal{X}$. "PPT" stands for Probabilistic Polynomial-Time. Every algorithm (including the adversaries and simulators) is stateful. For an execution of a stateful algorithm $\mathcal{A}$, we write $\mathsf{View}_{\mathcal{A}}$ for all inputs received by $\mathcal{A}$ throughout the execution, together with $\mathcal{A}$'s random coins.

2.2 Bilinear Groups

A Bilinear Group $\mathcal{BG}$, generated as $(\mathbb{F}_q, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, g_1, g_2, e) \leftarrow \mathcal{BG}(1^\lambda)$, is specified by a field $\mathbb{F}_q$ of prime order $q = 2^{\Theta(\lambda)}$, three groups $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$ (the first two we call "source groups" and the third "target group"), a bilinear map $e : \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_T$ that we call "pairing" and random generators g_1, g_2 for $\mathbb{G}_1, \mathbb{G}_2$ respectively. We use the *implicit notation*, i.e., $[x]_s := xg_s$ for $s \in \{1, 2, T\}$. Also, we denote the group operation additively, $[x]_s + [y]_s = [x + y]_s$, for $s \in \{1, 2, T\}$. The pairing has the property that $e([x]_1, [y]_2) = xye([1]_1, [1]_2)$. In our constructions, we will omit writing explicitly the Bilinear Group in the algorithms' inputs, even though it is implicitly taken as input.

2.2.1 Generic Bilinear Group Model

The Generic Group Model (GGM) [Sho97, Mau05] is an idealized model, that formalizes a 'generic' adversary. That is, the adversary does not have access to the concrete representation of the group elements but can only use generic group operations (addition, inverse element, scalar multiplication). This model captures the 'algebraic' attacks that an adversary can perform.

In this work, we use the Maurer's GGM [Mau05], which is extended to Bilinear Groups by [BBG05]. There, the adversary makes oracle queries for each generic group operation they wish to perform and receives a handle for the resulting group element, instead of the actual element itself. All group elements resulting from the adversary's queries are recorded—together with their handles—in three lists $\boldsymbol{L}_1, \boldsymbol{L}_2, \boldsymbol{L}_T$ for $\mathbb{G}_1$, $\mathbb{G}_2$ and $\mathbb{G}_T$ respectively.

A standard GGM technique in security proofs is the 'symbolic' equivalence. We call 'symbolic' experiment (and symbolic group representation, respectively) the model where polynomials instead of group elements are stored and polynomial operations instead of group operations are performed. The formal variables of the polynomials are the initial elements that the adversary received. For example, a generic adversary to the discrete logarithm problem is initially receiving $[1], [x]$; thus, the formal variables are $1, X$, and then can perform any generic group operation which is going to be symbolically performed with the corresponding polynomials in $\mathbb{Z}_q[1, X]$.

Master Theorem. We recall the 'Master Theorem' [BBG05, Boy08]. The Master Theorem is used to determine the probability loss between the symbolic group representation and the actual generic group representation, where the formal variables are instantiated.

Theorem 1 (Master theorem [BBG05, Boy08])**.** *Let $\boldsymbol{L}_1 \in \mathbb{Z}_q[X_1, \ldots, X_n]^{\nu_1}$, $\boldsymbol{L}_2 \in \mathbb{Z}_q[X_1, \ldots, X_n]^{\nu_2}$, $\boldsymbol{L}_T \in \mathbb{Z}_q[X_1, \ldots, X_n]^{\nu_T}$ be three lists of n-variate polynomials over $\mathbb{Z}_q$ of maximum degree $d_{\boldsymbol{L}_1}, d_{\boldsymbol{L}_2}, d_{\boldsymbol{L}_T}$, respectively. Let $f \in \mathbb{Z}_q[X_1, \ldots, X_n]$ be an n-variate polynomial of degree d_f and denote $d = \max\{d_{\boldsymbol{L}_1} + d_{\boldsymbol{L}_2}, d_{\boldsymbol{L}_T}, d_f\}$, $\nu = \nu_1 + \nu_2 + \nu_T$. If f is* independent *of $(\boldsymbol{L}_1, \boldsymbol{L}_2, \boldsymbol{L}_T)$, then for any generic adversary $\mathcal{A}$ that makes at most Q group oracle queries:*

$$\left| \Pr\left[\mathcal{A}\begin{pmatrix} q, \mathsf{h}_1[\boldsymbol{L}_1(\boldsymbol{x})], \\ \mathsf{h}_2[\boldsymbol{L}_2(\boldsymbol{x})], \mathsf{h}_T[f(\boldsymbol{x})] \end{pmatrix} = 1 \right] - \Pr\left[\mathcal{A}\begin{pmatrix} q, \mathsf{h}_1[\boldsymbol{L}_1(\boldsymbol{x})], \\ \mathsf{h}_2[\boldsymbol{L}_2(\boldsymbol{x})], \mathsf{h}_T[r] \end{pmatrix} = 1 \right] \right| \leqslant \frac{(Q + \nu + 2)^2 d}{2q},$$

where $\mathsf{h}_1, \mathsf{h}_2, \mathsf{h}_T$ denote the corresponding handles, and the probabilities are taken over the choices of $\boldsymbol{x} \leftarrow\!\!\$\ (\mathbb{Z}_q)^n$ and $r \leftarrow\!\!\$\ \mathbb{Z}_q$.

Here f-dependence on $\boldsymbol{L} = (\boldsymbol{L}_1, \boldsymbol{L}_2, \boldsymbol{L}_T)$ means that the polynomial f is in the span of the polynomials in the list $\boldsymbol{C}(\boldsymbol{L}) := \{\boldsymbol{L}_1 \otimes \boldsymbol{L}_2\} \cup \boldsymbol{L}_T$ (intuitively $\{\boldsymbol{L}_1 \otimes \boldsymbol{L}_2\}$ are all the elements in $\mathbb{G}_T$ that can be computed using pairings). Naturally, the opposite case is called f-independence of $\boldsymbol{L}$.

2.3 Succinct Non-Interactive Arguments of Knowledge (SNARKs)

We recall the definition of SNARKs. In this work we do not focus on the zero-knowledge property, therefore we omit it.

Definition 1 (SNARKs). *A SNARK for a family of relations $\mathcal{R}_{\mathsf{fam}}$ is a tuple of three algorithms* $(\mathtt{Gen}, \mathtt{Prove}, \mathtt{Verify})$*:*

- $\mathtt{Gen}(\mathcal{R}) \to (\mathsf{crs})$*: On input a relation $\mathcal{R} \in R_{\mathsf{fam}}$ generates a common reference string* crs*.*
- $\mathtt{Prove}(\mathsf{crs}, \boldsymbol{x}, \boldsymbol{w}) \to \pi$*: On input the common reference string* crs*, a statement $\boldsymbol{x}$ and the corresponding witness $\boldsymbol{w}$ computes a proof π.*
- $\mathtt{Verify}(\mathsf{crs}, \boldsymbol{x}, \pi) \to 0/1$*: On input the common reference string* crs*, a statement $\boldsymbol{x}$ and a proof π the verification algorithm outputs either 1 for accept or 0 for reject.*

It is further required that the following properties hold.
(Perfect) Correctness. *For every relation $\mathcal{R} \in \mathcal{R}_{\mathsf{fam}}$, and every statement-witness pair $(\boldsymbol{x}, \boldsymbol{w}) \in \mathcal{R}$:*

$$\Pr\left[\ \mathtt{Verify}(\mathsf{crs}, \boldsymbol{x}, \pi) = 1 \ \ : \ \begin{array}{c} \mathsf{crs} \leftarrow \mathtt{Gen}(\mathcal{R}) \\ \pi \leftarrow \mathtt{Prove}(\mathsf{crs}, \boldsymbol{x}, \boldsymbol{w}) \end{array}\right] = 1$$

Knowledge Soundness. *For every PPT adversarial Prover $\mathcal{P}^*$, there exists a PPT extractor $\mathcal{E}_{\mathcal{P}^*}$ such that for every benign auxiliary input* $\mathsf{aux} \in \{0,1\}^{\mathsf{poly}(\lambda)}$*, and every relation $\mathcal{R} \in \mathcal{R}_{\mathsf{fam}}$:*

$$\Pr\left[\begin{array}{l} \mathtt{Verify}(\mathsf{crs}, \boldsymbol{x}, \pi) = 1 \\ \wedge (\boldsymbol{x}, \boldsymbol{w}) \notin \mathcal{R} \end{array} \ : \ \begin{array}{c} \mathsf{crs} \leftarrow \mathtt{Gen}(\mathcal{R}) \\ (\boldsymbol{x}, \pi) \leftarrow \mathcal{P}^*(\mathsf{crs}, \mathsf{aux}) \\ \boldsymbol{w} \leftarrow \mathcal{E}_{\mathcal{P}^*}(\mathsf{View}_{\mathcal{P}^*}) \end{array}\right] = \mathsf{negl}(\lambda)$$

Succinctness. *There exists a universal polynomial $p(\cdot)$ such that, for every security parameter $\lambda \in \mathbb{N}$, every relation $\mathcal{R} \in \mathcal{R}_{\mathsf{fam}}$, and every statement-witness pair $(\boldsymbol{x}, \boldsymbol{w})$:*

- *An honestly generated proof π has size $p(\lambda + \log |\boldsymbol{w}|)$.*
- *The Verifier algorithm $\mathtt{Verify}(\mathsf{crs}, \boldsymbol{x}, \pi)$ runs in time $p(\lambda + |\boldsymbol{x}| + \log |\boldsymbol{w}|)$.*

2.3.1 The Groth16 SNARK

We recall the Groth16 proof system [Gro16], excluding the zero-knowledge property.

Rank-1 constraint satisfiability (R1CS). Groth16 works for relations encoded with the rank-1 constraint satisfiability (R1CS). Formally, an R1CS relation is of the form:

$$\mathcal{R} = \left\{(\boldsymbol{x}; \boldsymbol{w}) : (\boldsymbol{A} \times \boldsymbol{z}^T) \circ (\boldsymbol{B} \times \boldsymbol{z}^T) = \boldsymbol{C} \times \boldsymbol{z}^T \wedge \boldsymbol{z} = (\boldsymbol{x} \| \boldsymbol{w}) \wedge x_1 = 1\right\}$$

where the relation is characterized by the fixed matrices $\boldsymbol{A}, \boldsymbol{B}, \boldsymbol{C} \in \mathbb{Z}_q^{n \times m}$, the statement $\boldsymbol{x} \in \mathbb{Z}_q^{\ell}$ is an ℓ-sized vector,[7] the witness $\boldsymbol{w} \in \mathbb{Z}_q^{m-\ell}$ is an $(m-\ell)$-sized vector, and $\boldsymbol{z} \in \mathbb{Z}_q^m$ is called the 'extended witness', consisting of the witness and the statement. Here '$\circ$' is the Hadamard product. R1CS generalizes arithmetic circuits.

[7] In the R1CS encoding, the first entry of $\boldsymbol{x}$ is equal to 1 to rule out the degenerate case of the $\boldsymbol{z} = \boldsymbol{0}$ solution.

The Groth16 construction. For the proof system each column of the R1CS matrices $\boldsymbol{A}, \boldsymbol{B}, \boldsymbol{C}$ is interpolated into polynomials as: $a_i(X) = \sum_{j=1}^{n} a_{i,j} L_j(X), b_i(X) = \sum_{j=1}^{n} b_{i,j} L_j(X), c_i(X) = \sum_{j=1}^{n} c_{i,j} L_j(X)$ for each $i \in [m]$, where $L_j(x)$ the corresponding Lagrange polynomial. Then the relation boils down to the following polynomial relation:

$$\left(\sum_{i=1}^{m} z_i a_i(X)\right)\left(\sum_{i=1}^{m} z_i b_i(X)\right) - \sum_{i=1}^{m} z_i c_i(X) = q(X)V(X)$$

where $V(X) = \prod_{i=1}^{n}(X - \omega^i)$ is the vanishing polynomial. This polynomial relation is essentially equivalent to the R1CS satisfiability (we refer to [GGPR13, PHGR13, Gro16] for more details).

The actual Groth16 SNARK is described below.

- $\mathtt{Gen}(\mathcal{R}) \to \mathsf{crs}$: Samples uniformly $\tau, \alpha, \beta, \gamma, \delta \leftarrow\!\!\$\ \mathbb{Z}_q$ and outputs:

$$\mathsf{crs} = \left\{ \{[\alpha]_1, [\beta]_2, [\gamma]_2, [\delta]_1, [\delta]_2, \left\{ \left[\tau^i\right]_1, \left[\tau^i\right]_2 \right\}_{i=0}^{n-1}, \left\{ \left[\frac{V(\tau)\tau^i}{\delta}\right]_1 \right\}_{i=0}^{n-2}, \right.$$
$$\left\{[a_i(\tau)]_1, [b_i(\tau)]_1, [b_i(\tau)]_2\right\}_{i=1}^{m}, \left\{ \left[\frac{\beta a_i(\tau) + \alpha b_i(\tau) + c_i(\tau)}{\gamma}\right]_1 \right\}_{i=1}^{\ell},$$
$$\left. \left\{ \left[\frac{\beta a_i(\tau) + \alpha b_i(\tau) + c_i(\tau)}{\delta}\right]_1 \right\}_{i=\ell+1}^{m} \right\}$$

- $\mathtt{Prove}(\mathsf{crs}, \boldsymbol{x}, \boldsymbol{w}) \to \pi$: Sets $\boldsymbol{z} = (\boldsymbol{x} \| \boldsymbol{w})$. Computes the quotient polynomial $q(X) = \sum_{i=0}^{n-2} q_i X^i = \left\lfloor \frac{\left(\sum_{i=1}^{m} z_i a_i(X)\right)\left(\sum_{i=1}^{m} z_i b_i(X)\right) - \sum_{i=1}^{m} z_i c_i(X)}{V(X)} \right\rfloor$ and then outputs $\pi = (\pi_1, \pi_2, \pi_3)$, where:

$$\pi_1 = [\alpha]_1 + \sum_{i=1}^{m} z_i [a_i(\tau)]_1$$
$$\pi_2 = [\beta]_2 + \sum_{i=1}^{m} z_i [b_i(\tau)]_2$$
$$\pi_3 = \sum_{i=\ell+1}^{m} z_i \left[\frac{\beta a_i(\tau) + \alpha b_i(\tau) + c_i(\tau)}{\delta}\right]_1 + \sum_{i=0}^{n-2} q_i \left[\frac{V(\tau)\tau^i}{\delta}\right]_1$$

- $\mathtt{Verify}(\mathsf{crs}, \boldsymbol{x}, \pi) \to 0/1$: Outputs 1 iff:

$$e(\pi_1, \pi_2) = e\left([\alpha]_1, [\beta]_2\right) + e\left(\sum_{i=1}^{\ell} x_i \left[\frac{\beta a_i(\tau) + \alpha b_i(\tau) + c_i(\tau)}{\gamma}\right]_1, [\gamma]_2\right) + e(\pi_3, [\delta]_2) \quad (4)$$

Theorem 2 (Groth16 Security [Gro16]). *The above protocol is a SNARK, satisfying Perfect Correctness and Knowledge Soundness in the Generic Bilinear Group Model.*

2.4 Garbling Schemes

Here we recall the definition of garbling schemes for binary circuits introduced by Yao in 1982 [Yao82] and further formalized in [LP09, BHR12b].

Definition 2. *A garbling scheme for a family of binary circuits $\mathcal{C}_{\mathsf{fam}}$ is a tuple of PPT algorithms* $\mathsf{GC} = (\mathsf{Garble}, \mathsf{Encode}, \mathsf{Eval})$ *such that:*

- $\mathsf{Garble}(C) \to (\mathsf{ct}_{\mathsf{GC}}, \mathsf{ek})$*: Takes as input a circuit* $C : \{0,1\}^n \to \{0,1\}^k \in \mathcal{C}_{\mathsf{fam}}$ *and outputs a garbled circuit description* $\mathsf{ct}_{\mathsf{GC}}$ *and an encoding key* ek*.*
- $\mathsf{Encode}(\mathsf{ek}, x) \to L_x$*: Takes as input an encoding key* ek *and a circuit input* $x \in \{0,1\}^n$ *and outputs the encoded input* L_x*. We also refer to* L_x *as the 'input labels' (or sometimes just 'labels').*
- $\mathsf{Eval}(\mathsf{ct}_{\mathsf{GC}}, L_x)$*: On input the garbled circuit* $\mathsf{ct}_{\mathsf{GC}}$ *and the labels of the input* L_x*, outputs the circuit output* $y \in \{0,1\}^k$

Furthermore, it should satisfy the following properties.

Correctness. *We require that for any binary circuit* $C \in \mathcal{C}_{\mathsf{fam}}$ *and any input* $x \in \{0,1\}^{|\mathsf{inp}(C)|}$*,*

$$\Pr\left[\; C(x) = \mathsf{Eval}(\mathsf{ct}_{\mathsf{GC}}, L_x) \;\; : \;\; \begin{array}{c} (\mathsf{ct}_{\mathsf{GC}}, \mathsf{ek}) \leftarrow \mathsf{Garble}(1^\lambda, C) \\ L_x \leftarrow \mathsf{Encode}(\mathsf{ek}, x) \end{array} \right] = 1$$

Adaptive Privacy. *We require that there exists a PPT simulator* $\mathsf{Sim} = (\mathsf{Sim}_1, \mathsf{Sim}_2)$ *such that for any PPT adversary* $\mathcal{A}$*, for every circuit* $C \in \mathcal{C}_{\mathsf{fam}}$ *and for every auxiliary information* $\mathsf{aux} \in \{0,1\}^{\mathsf{poly}(\lambda)}$*:*

$$\left| \Pr\left[\mathcal{A}(\mathsf{ct}_{\mathsf{GC}}, L_x, \mathsf{aux}) = 1 \;\; : \;\; \begin{array}{r} (\mathsf{ct}_{\mathsf{GC}}, \mathsf{ek}) \leftarrow \mathsf{Garble}(C) \\ x \leftarrow \mathcal{A}(\mathsf{ct}_{\mathsf{GC}}, \mathsf{aux}) \\ L_x \leftarrow \mathsf{Encode}(\mathsf{ek}, x) \end{array} \right] \right.$$
$$\left. - \Pr\left[\mathcal{A}(\mathsf{ct}_{\mathsf{GC}}, L_x, \mathsf{aux}) = 1 \;\; : \;\; \begin{array}{r} (\mathsf{ct}_{\mathsf{GC}}, \mathsf{st}) \leftarrow \mathsf{Sim}_1(\mathsf{topo}(C)) \\ x \leftarrow \mathcal{A}(\mathsf{ct}_{\mathsf{GC}}, \mathsf{aux}) \\ L_x \leftarrow \mathsf{Sim}_2(\mathsf{st}, x, C(x)) \end{array} \right] \right| \leqslant \mathsf{negl}(\lambda)$$

The garbling scheme is called privacy-free *if it satisfies only correctness.*

In Yao's garbling scheme the encoding key is $\{L^b_{x,i}\}^n_{i=1,b=0,1}$ and the labels of x are $L_x = (L^{x_1}_{x,1}, \ldots, L^{x_n}_{x,n})$. Notice that fortuitously the algorithm coincide with a Lamport signature [Lam79] on x. Yao GC can be simply proven adaptively secure in the random oracle model [BHR12a]. We will make use of the well-established free-XOR [KS08] and half-gate [ZRE15] optimizations.

2.5 The Bitcoin Ledger

Ledger Model. We model Bitcoin as a distributed ledger functionality $\mathcal{F}_{\mathsf{BTC}}$ maintaining an append-only sequence of confirmed transactions (a *ledger*). We assume a global clock which counts time in rounds $r \in \mathbb{N}$ and $r \leqslant \mathsf{poly}(\lambda)$. Any party can submit a transaction tx to the ledger at any round r by calling $\mathcal{F}_{\mathsf{BTC}}.\mathsf{WRITE}(\mathsf{tx})$. For a party P and round $r \in \mathbb{N}$, the notation $\mathcal{L}^r_P$ denotes party P's local view of the Bitcoin ledger at round r.

The height of a ledger $\mathcal{L}$ is denoted $h(\mathcal{L})$. A ledger of height h is a sequence of blocks $\mathcal{L}[0], \ldots, \mathcal{L}[h-1]$. and $\mathcal{L}[:h]$ denotes the sequence of all blocks in $\mathcal{L}$ less than height h. $\mathcal{L}||B$ denotes the ledger formed by appending the block B to the ledger $\mathcal{L}$. Each block $\mathcal{L}[i]$ is a sequence of transactions. By flattening the blocks, the ledger itself can be viewed as a sequence of transactions. We denote by $\mathcal{L}_1 \preceq \mathcal{L}_2$ that $\mathcal{L}_1$ is a prefix of $\mathcal{L}_2$.

We assume the Bitcoin ledger satisfies the two fundamental properties safety and liveness. Informally, safety ensures that the ledgers of honest parties are consistent with each other, and liveness ensures that new valid transactions are eventually included in the ledgers of honest parties. Safety is defined as in [GKL15].

Definition 3 (Ledger Safety). *A ledger functionality $\mathcal{F}_{\mathsf{BTC}}$ is* safe *if with probability $1 - \mathsf{negl}(\lambda)$: (i) for any honest party P and rounds $r_1 \leq r_2$, $\mathcal{L}_P^{r_1} \preceq \mathcal{L}_P^{r_2}$ (self-consistency); and (ii) for any honest parties P_1, P_2 and any round r, $\mathcal{L}_{P_1}^r \preceq \mathcal{L}_{P_2}^r$ or $\mathcal{L}_{P_2}^r \preceq \mathcal{L}_{P_1}^r$ (view-consistency)*

To precisely define liveness, we first define the validity of a Bitcoin transaction, beginning with a model for Bitcoin transactions. We discuss only the features of Bitcoin transactions that are relevant to BABE.

UTXO Model and Taproot Scripts. The ledger's state is represented as a set of unspent transaction outputs (UTXOs). Each UTXO is a pair $\mathsf{out} = (a, \mathsf{lockScript})$ where $a \in \mathbb{R}_{\geqslant 0}$ is the amount of coins (in BTC) in that UTXO, and $\mathsf{lockScript}$ is a program (the *locking script*) that determines under which conditions the UTXO can be spent.

In BABE, we widely use Taproot Trees [WNT20], or Taptrees, which make a UTXO spendable by satisfying one among multiple locking scripts. We will represent the locking script of such a UTXO as $\mathsf{lockScript} = \langle \mathsf{leaf}_0, \ldots, \mathsf{leaf}_{k-1} \rangle$.

Transactions. A transaction is a triple $\mathsf{tx} = (\mathsf{inputs}, \mathsf{tx_witnesses}, \mathsf{outputs})$ where:

- $\mathsf{inputs} = [\mathsf{in}_1, \ldots, \mathsf{in}_n]$, where each input references an output of a previous transaction, indexed as $\mathsf{in} = (\mathsf{PrevTx}, \mathsf{outIndex}, \mathsf{leaf})$ where PrevTx is the previous transaction, $\mathsf{outIndex}$ is the index of the output in the previous transaction, and leaf is the leaf of the Taproot tree to be satisfied.
- $\mathsf{outputs} = [\mathsf{out}_1, \ldots, \mathsf{out}_m]$ where each output is a pair $(a, \mathsf{lockScript})$ as above;
- $\mathsf{tx_witnesses} = [w_1, \ldots, w_n]$, where w_i is a transaction witness (e.g., signature, data, etc.) intended to satisfy the corresponding input's leaf script.

The pair $\overline{\mathsf{tx}} = (\mathsf{inputs}, \mathsf{outputs})$ is called the *transaction skeleton.*

Cryptographic Primitives. We assume the existence of the following cryptographic primitives that are used by the ledger functionality:

- *Signature Scheme:* Sig_{BTC} with algorithms $\mathsf{Sig}_{BTC}.\mathtt{Gen}(1^\lambda) \to (\mathsf{sk}, \mathsf{pk})$, $\mathsf{Sig}_{BTC}.\mathsf{Sign}(\mathsf{sk}, \overline{\mathsf{tx}}) \to \sigma$, and $\mathsf{Sig}_{BTC}.\mathtt{Verify}(\mathsf{pk}, \overline{\mathsf{tx}}, \sigma) \to 0/1$ which is EUF-CMA secure.
- *Hash Function:* $\mathsf{Hash}_{BTC}(m) \to h$ modeled as a random oracle.
- *Lamport One-Time Signature Scheme [Lam79]:* LampSig with algorithms $\mathsf{LampSig}.\mathtt{Gen}(1^\lambda, \ell) \to (\mathsf{lsk}, \mathsf{lpk})$, $\mathsf{LampSig}.\mathsf{Sign}(\mathsf{lsk}, m) \to \mu$, and $\mathsf{LampSig}.\mathtt{Verify}(\mathsf{lpk}, m, \mu) \to 0/1$ where ℓ is the number of bits in the message m. A Lamport signature scheme constructed using Hash_{BTC} can be efficiently verified in Bitcoin script.

$$\mathsf{lsk} := \begin{pmatrix} L_0^0 & \ldots & L_{\ell-1}^0 \\ L_0^1 & \ldots & L_{\ell-1}^1 \end{pmatrix} \leftarrow\!\!\$ \left(\{0,1\}^\lambda\right)^{2\times\ell} \tag{5}$$

$$\mathsf{lpk} := \begin{pmatrix} \mathsf{Hash}_{BTC}(L_0^0) & \ldots & \mathsf{Hash}_{BTC}(L_{\ell-1}^0) \\ \mathsf{Hash}_{BTC}(L_0^1) & \ldots & \mathsf{Hash}_{BTC}(L_{\ell-1}^1) \end{pmatrix} \tag{6}$$

$$\mu := \begin{pmatrix} L_0^{m_0} & \ldots & L_{\ell-1}^{m_{\ell-1}} \end{pmatrix} \tag{7}$$

$$\mathsf{LampSig}.\mathtt{Verify}(\mathsf{lpk}, m, \mu) = 1 \iff \forall i \in \{0, \ldots, \ell-1\}, \mathsf{Hash}_{BTC}(\mu_i) = \mathsf{lpk}_i^{m_i} \tag{8}$$

Locking Scripts. The locking scripts of Bitcoin that are used by BABE are described below:

- *Signature Check:* $\mathsf{CheckSig}(\mathsf{pk})$ — Requires w to contain a digital signature σ such that $\mathsf{Sig}_{BTC}.\mathtt{Verify}(\mathsf{pk}, \overline{\mathsf{tx}}, \sigma) = 1$.
- *Relative Timelock:* $\mathsf{RelTimelock}(\tau)$ — Requires that at least τ blocks have elapsed since the output referenced by the input was created. That is, if a transaction with an input of the form $\mathsf{in} = (\mathsf{PrevTx}, \mathsf{outIndex}, \langle \mathsf{RelTimelock}(\tau) \rangle)$ is included in a block at height h, then PrevTx must have been included in a block at height at most $h - \tau$.
- *Hash Lock:* $\mathsf{HashLock}(h)$ — Requires that the transaction witness contains a hash pre-image of h, i.e. w such that $\mathsf{Hash}_{BTC}(w) = h$.
- *Logical Operations:* The above scripts can be combined using logical operator $\wedge$ (AND) and $\vee$ (OR) to form more complex locking conditions.

Using the above scripts and logical operator $\wedge$ (AND) and $\vee$ (OR), we can define more complex locking scripts, for example:

- *Lamport Signature:* $\mathsf{CheckLampSig}(\mathsf{lpk})$ — Requires that the transaction witness contains a Lamport signature for the public key lpk on some message. In BABE, we will use this script to ensure that the Prover, while posting this transaction, makes a binding commitment to some proof.

$$\mathsf{CheckLampSig}(\mathsf{lpk}) \;:=\; \bigwedge_{i=1}^{\ell} \left(\mathsf{HashLock}(\mathsf{lpk}_i^0) \vee \mathsf{HashLock}(\mathsf{lpk}_i^1)\right) \tag{9}$$

- *Check Lamport Signatures from Multiple Parties on Same Message:* $\mathsf{CheckLampSigsMatch}(\mathsf{lpk}_A, \mathsf{lpk}_B)$ — Requires that the transaction witness contains a Lamport signature for each public key $\mathsf{lpk}_A, \mathsf{lpk}_B$, both on the same message. In BABE, we will use this script to ensure that the Verifier (who holds only lsk_B) commits to the same proof that the Prover (who holds lsk_A) committed to. This is guaranteed because the Verifier, who does not know lsk_A, cannot create a valid Lamport signature under lsk_A for any message other than the one committed to by the Prover. The Prover can then use the Verifier's commitment to evaluate the garbled circuit.

$$\begin{aligned}\mathsf{CheckLampSigsMatch}(\mathsf{lpk}_A, \mathsf{lpk}_B, \mathsf{msg}) \;:=\; \bigwedge_{i=1}^{\ell} \Big[&\left(\mathsf{HashLock}((\mathsf{lpk}_A)_i^0) \wedge \mathsf{HashLock}((\mathsf{lpk}_B)_i^0)\right) \\ &\vee \left(\mathsf{HashLock}((\mathsf{lpk}_A)_i^1) \wedge \mathsf{HashLock}((\mathsf{lpk}_B)_i^1)\right)\Big]\end{aligned} \tag{10}$$

Definition 4 (Transaction Validity)**.** *A transaction* $\mathsf{tx} = (\mathsf{inputs}, \mathsf{tx_witnesses}, \mathsf{outputs})$ *is* valid *with respect to a ledger* $\mathcal{L}$*, denoted* $\mathsf{Valid}_{\mathcal{L}}(\mathsf{tx})$*, if:*

1. All Inputs Unspent: *For each* $\mathsf{in}_i = (\mathsf{PrevTx}, \mathsf{outIndex}, \mathsf{leaf})$ *in* inputs*, the transaction* PrevTx *exists in* $\mathcal{L}$*.*
2. All Locking Scripts Satisfied: *Each input's specified leaf script is one of the leaves of the taptree for that input. That is, for each* $\mathsf{in}_i = (\mathsf{PrevTx}, \mathsf{outIndex}, \mathsf{leaf})$*,* $\mathsf{leaf} \in \langle \mathsf{leaf}_0, \ldots, \mathsf{leaf}_{k-1} \rangle = \mathsf{PrevTx.outputs}[\mathsf{outIndex}].\mathsf{lockScript}$*. Moreover, the transaction witness* $\mathsf{tx_witnesses}[i]$ *must satisfy the script* leaf*.*

3. Value Preservation: *The total amount of coins in the outputs is less than or equal to the total amount of coins in the inputs.*

$$\sum_{j=1}^{m} \mathsf{outputs}[j].a \leq \sum_{i=1}^{n} \mathsf{in}_i.\mathsf{PrevTx}.\mathsf{outputs}[\mathsf{in}_i.\mathsf{outIndex}].a \tag{11}$$

A ledger $\mathcal{L}$ is valid, *denoted* $\mathsf{Valid}(\mathcal{L})$, *if for all transactions* tx *in* $\mathcal{L}$, $\mathsf{Valid}_{\mathcal{L}[:\mathsf{tx}]}(\mathsf{tx}) = 1$ *($\mathcal{L}[:\mathsf{tx}]$ is the ledger containing all transactions in $\mathcal{L}_P^r$ before* tx*).*

Definition 5 (Ledger Validity)**.** *If a transaction* tx *appears in the ledger view $\mathcal{L}_P^r$ of any party P at any round r, then* tx *is valid with respect to the state $\mathcal{L}_P^r[:\mathsf{tx}]$ (the ledger containing all transactions in $\mathcal{L}_P^r$ before* tx*).*

Finally, to define liveness, we note that only valid transactions may be included in the ledger. However, the adversary may delay the inclusion of a valid transaction. The adversary may also include his own transactions in the ledger which may cause the honest party's transaction to become invalid thereafter, e.g. the adversary's transaction may use the same input as the honest party's transaction. Thus, liveness is guaranteed when the adversary is (computationally) unable to exclude a transaction for too many blocks. To make this precise, we extend the definition of unambiguous transactions [GKL15] to *unstoppable transactions*. Informally, a transaction is unstoppable if inserting this transaction in any sequence of u valid blocks preserves their validity. For example, a transaction that requires the honest party's signature will be unstoppable since the adversary cannot forge the honest party's signature, hence the inputs for the transaction cannot be spent by the adversary before the honest party's transaction is included in the ledger. We assume that for large enough u, Bitcoin satisfies liveness for u-unstoppable transactions. In general, the adversary may have access to some state st, which may include certain signatures, ciphertexts, etc. that the adversary has access to.

Definition 6 (u-Unstoppable Transactions)**.** *A transaction* tx *is* u-unstoppable *with respect to a ledger $\mathcal{L}$ and adversarial state* st*, if for all* PPT *adversaries $\mathcal{A}$:*

$$\Pr\left[\mathsf{Valid}(\mathcal{L}||B_1||\ldots||B_i'||\ldots||B_u) = 1 \;:\; \begin{array}{l} B_1,\ldots,B_u,i,B_i' \leftarrow \mathcal{A}(\mathcal{L},\mathsf{st},\mathsf{tx}) \\ \mathsf{Valid}(\mathcal{L}||B_1||\ldots||B_u) = 1 \\ B_i' \setminus B_i = \{\mathsf{tx}\} \end{array}\right] \geqslant 1 - \mathsf{negl}(\lambda) \tag{12}$$

Here, $B_i' \setminus B_i = \{\mathsf{tx}\}$ means that the block B_i' contains the same sequence of transactions as B_i except that tx *is inserted at some location in the sequence.*

Definition 7 (Ledger Liveness)**.** *A ledger functionality $\mathcal{F}_{\mathsf{BTC}}$ is* u-live *($u \in \mathbb{N}$) if for all adversarial states* st*, with probability $1 - \mathsf{negl}(\lambda)$, for any party P calling $\mathcal{F}_{\mathsf{BTC}}.\mathsf{WRITE}(\mathsf{tx})$ at any round r such that* tx *is u-unstoppable with respect to $\mathcal{L}_P^r$ and* st*, for all honest parties H and rounds r' with $h(\mathcal{L}_H^{r'}) \geqslant h(\mathcal{L}_P^r) + u$,* $\mathsf{tx} \in \mathcal{L}_H^{r'}[:h(\mathcal{L}_P^r)+u]$.

While the above definition of liveness guarantees inclusion of a transaction within a certain number of blocks, we use the *chain growth* property [GKL15] to guarantee that the height of every party's ledger grows.

Theorem 3 (τ-Chain Growth)**.** *For all $s > \lambda$, for all honest parties P,* $\Pr\left[\forall\, r : h(\mathcal{L}_P^{r+s}) < h(\mathcal{L}_P^r) + \tau s\right] \leqslant \mathsf{negl}(\lambda)$.

3 The BitVM-core Primitive

We define the BitVM-core primitive (c.f. BitVM2-core [LAA+25], ChainVM [ADJK25]).[8] A BitVM-core protocol is an interactive protocol run between a Prover P and a Verifier V interacting with the Bitcoin ledger functionality $\mathcal{F}_{\mathsf{BTC}}$. The protocol consists of two phases: a setup phase and a proving phase. In the setup phase, the Prover and Verifier are given a common reference string crs and a statement $\boldsymbol{x}$ to be proven, and interact off-chain to agree on a set of Bitcoin transactions $\mathcal{S}$. These transactions are called the *success transactions* because they can be posted on Bitcoin in the proving phase if and only if he has a valid witness for the statement. During setup, the Prover and Verifier also obtain and store local state (e.g., transactions, ciphertexts, private keys) to be used in the proving phase. To defend against malicious behavior, both the Prover and Verifier validate information received from the other party and may abort the setup if they detect invalid information.

During the proving phase, both the Prover and Verifier interact with the Bitcoin ledger. The Prover, who now knows a witness for the statement to be proven, and the Verifier take turns posting transactions to the Bitcoin ledger. In the end, the Prover wins by posting a success transaction from $\mathcal{S}$ or the Verifier wins by preventing the Prover from ever posting a success transaction. Security means that the Prover must win if he has a valid witness (even if the Verifier is malicious) and the Verifier must win otherwise (even if the Prover is malicious).

This definition of BitVM-core abstracts away the deposit and withdrawal of Bitcoin, similar to ChainVM [ADJK25], and indicates the validity of the proof by the presence of a transaction from $\mathcal{S}$ on the ledger. Applications such as the trustless lending protocol described in Sec. 1 or wrapped asset bridges [LAA+25, WAA+26, ADJK25] may use a separate transaction to deposit Bitcoin and use this primitive to allow withdrawal if and only if the Prover has a valid witness.

Protocol Syntax. A BitVM-core protocol for a relation $\mathcal{R}$ consists of the following PPT algorithms:

- $\mathtt{Gen}(\mathcal{R}) \to \mathsf{crs}$: Takes as input the relation and outputs a common reference string crs.
- $P_{\mathtt{Setup}}(\mathsf{crs}, \boldsymbol{x})$: Run by the Prover during the setup phase, interacts with the Verifier and the ledger functionality $\mathcal{F}_{\mathsf{BTC}}$, and outputs the Prover's state st_P and a set of transactions $\mathcal{S}$. Alternatively, the Prover may output $\perp$ to indicate that the setup failed.
- $V_{\mathtt{Setup}}(\mathsf{crs}, \boldsymbol{x})$: Run by the Verifier during the setup phase, interacts with the Prover and the ledger functionality $\mathcal{F}_{\mathsf{BTC}}$, and outputs the Verifier's state st_V and a set of transactions $\mathcal{S}$. Alternatively, the Verifier may output $\perp$ to indicate that the setup failed.
- $P_{\mathtt{Prove}}(\mathsf{st}_P, \boldsymbol{w})$: Run by the Prover knowing a witness $\boldsymbol{w}$, interacts with the ledger functionality $\mathcal{F}_{\mathsf{BTC}}$, and outputs a bit $b \in \{0, 1\}$ ($b = 1$ indicates success of the proving phase).
- $V_{\mathtt{Prove}}(\mathsf{st}_V)$: Run by the Verifier during the proving phase, interacts with the ledger functionality $\mathcal{F}_{\mathsf{BTC}}$.

We denote by $\langle P_{\mathtt{Setup}}(\mathsf{crs}, \boldsymbol{x}), V_{\mathtt{Setup}}(\mathsf{crs}, \boldsymbol{x})\rangle$ the random variable representing the transcript of the setup phase run interactively by the Prover and Verifier. The outputs of the Prover and Verifier during setup are denoted respectively by

$$\mathsf{out}_P(\langle P_{\mathtt{Setup}}(\mathsf{crs}, \boldsymbol{x}), V_{\mathtt{Setup}}(\mathsf{crs}, \boldsymbol{x})\rangle) = (\mathcal{S}, \mathsf{st}_P) \text{ or } \perp \tag{13}$$

$$\mathsf{out}_V(\langle P_{\mathtt{Setup}}(\mathsf{crs}, \boldsymbol{x}), V_{\mathtt{Setup}}(\mathsf{crs}, \boldsymbol{x})\rangle) = (\mathcal{S}, \mathsf{st}_V) \text{ or } \perp. \tag{14}$$

[8] Compared to BitVM2-core [LAA+25] and ChainVM [ADJK25], we simplify the definition by considering a single Verifier and excluding BitVM2-core's optimistic path. We discuss extensions in Sec. 8.

and the union of both outputs is denoted by

$$\mathsf{out}\langle P_{\mathtt{Setup}}(\mathsf{crs}, \boldsymbol{x}), V_{\mathtt{Setup}}(\mathsf{crs}, \boldsymbol{x})\rangle = (\mathcal{S}, \mathsf{st}_P, \mathsf{st}_V) \tag{15}$$

if neither party aborted, and $\perp$ otherwise. Similarly, the outputs of each party during the proving phase are denoted respectively by $\mathsf{out}_P(\langle P_{\mathtt{Prove}}(\cdot), V_{\mathtt{Prove}}(\cdot)\rangle)$ and $\mathsf{out}_V(\langle P_{\mathtt{Prove}}(\cdot), V_{\mathtt{Prove}}(\cdot)\rangle)$ where $(\cdot)$ is replaced by the inputs of the corresponding algorithm. We use $\langle P_{\mathtt{Prove}}(\cdot), V_{\mathtt{Prove}}(\cdot)\rangle_r$ to denote that the proving phase began at round r whenever the starting round is relevant.

Robustness ensures that as long as the Prover is honest, accepted that the setup was performed correctly, and has a valid witness, a malicious Verifier cannot prevent the Prover from succeeding (i.e., the required success transaction appears on the ledger) except with negligible probability. Thus, the worst thing a malicious Verifier can do is to cause the Prover to abort the setup, in which case no Bitcoin is ever lost by any party.

Definition 8 (u-Robustness)**.** *For all NP relations $\mathcal{R}$, for all statements x and witness w such that $(x, w) \in \mathcal{R}$, for all* PPT *adversarial Verifiers V^*, and all rounds $r \in \mathbb{N}$, the following holds:*

$$\Pr\left[\begin{array}{rcl} & & \mathsf{crs} \leftarrow \mathtt{Gen}(\mathcal{R}) \\ \mathsf{out}_P \neq \perp & & \mathsf{out}_P \leftarrow \mathsf{out}_P(\langle P_{\mathtt{Setup}}(\mathsf{crs}, \boldsymbol{x}), V^*(\mathsf{crs}, \boldsymbol{x})\rangle) \\ \wedge \big((b \neq 1) \vee (\forall\, \mathsf{tx} \in \mathcal{S}, & : & \textit{If } \mathsf{out}_P \neq \perp : \\ \mathsf{tx} \notin \mathcal{L}_P^{r+u})\big) & & \quad (\mathcal{S}, \mathsf{st}_P) \leftarrow \mathsf{out}_P \\ & & \quad b \leftarrow \mathsf{out}_P(\langle P_{\mathtt{Prove}}(\mathsf{st}_P, \boldsymbol{w}), V^*(\mathsf{crs}, \boldsymbol{x})\rangle_r) \end{array}\right] \leqslant 2^{-\kappa} + \mathsf{negl}(\lambda) \tag{16}$$

Knowledge soundness ensures that if the Verifier accepts the setup, then a malicious Prover cannot successfully include the success transaction on Bitcoin unless he knows a valid witness for the chosen relation. As usual in knowledge soundness, the Prover's knowledge of the witness is defined by the existence of an extractor that can extract the witness using the Prover's inputs and its view View_{P^*} of its interactions with the Verifier and Bitcoin. Below, assigning the proving interaction to $\tau_{\mathtt{Prove}}$ denotes executing that interaction, including its ledger effects; its parties' outputs are not used.

Definition 9 (Knowledge Soundness)**.** *For all NP relations $\mathcal{R}$, all* PPT *adversarial Provers P^*, there exists a* PPT *extractor $\mathcal{E}_{P^*}$ such that for every benign auxiliary input $\mathsf{aux} \in \{0,1\}^{\mathsf{poly}(\lambda)}$:*

$$\Pr\left[\begin{array}{rcl} & & \mathsf{crs} \leftarrow \mathtt{Gen}(\mathcal{R}) \\ \mathsf{out}_V \neq \perp & & \boldsymbol{x} \leftarrow P^*(\mathsf{crs}, \mathsf{aux}) \\ \wedge (\exists\, r \in \mathbb{N}, \mathsf{tx} \in \mathcal{S} : & & \mathsf{out}_V \leftarrow \mathsf{out}_V(\langle P^*(\mathsf{crs}, \mathsf{aux}), V_{\mathtt{Setup}}(\mathsf{crs}, \boldsymbol{x})\rangle) \\ \mathsf{tx} \in \mathcal{L}_V^r) & : & \textit{If } \mathsf{out}_V \neq \perp : \\ \wedge\, (\boldsymbol{x}, \boldsymbol{w}) \notin \mathcal{R} & & \quad (\mathcal{S}, \mathsf{st}_V) \leftarrow \mathsf{out}_V \\ & & \quad \tau_{\mathtt{Prove}} \leftarrow \langle P^*(\mathsf{crs}, \mathsf{aux}), V_{\mathtt{Prove}}(\mathsf{st}_V)\rangle \\ & & \quad \boldsymbol{w} \leftarrow \mathcal{E}_{P^*}(\mathsf{View}_{P^*}) \end{array}\right] \leqslant 2^{-\kappa} + \mathsf{negl}(\lambda) \tag{17}$$

Remark 1. *Assuming the existence of strong cryptographic primitives, there are auxiliary information distributions for which not all SNARKs admit an extractor [BCPR14, BP15]. Following the SNARK literature, we therefore, formally assume that the adversary should have access only to "benign" auxiliary inputs. Notably, these results are highly theoretical and do not affect SNARKs' security in practice.*

Together, robustness and knowledge soundness make a BitVM-core protocol *trustless*, solving the problem stated in Sec. 1. It is trustless because once the Prover and Verifier have mutually accepted that setup was performed correctly, an honest Prover can withdraw Bitcoin even if the Verifier is malicious, and an honest Verifier can prevent any malicious Prover from withdrawing Bitcoin.

Since a trivial protocol where the setup always aborts satisfies the above two properties, we also require that if both parties are honest, then neither party aborts during setup.

Definition 10 (Setup Correctness). *For all NP relations $\mathcal{R}$, all statements $\boldsymbol{x}$,*

$$\Pr\left[\begin{array}{c}\mathsf{out}\langle P_{\mathtt{Setup}}(\mathsf{crs},\boldsymbol{x}),V_{\mathtt{Setup}}(\mathsf{crs},\boldsymbol{x})\rangle=\bot:\\ \mathsf{crs}\leftarrow\mathtt{Gen}(\mathcal{R})\end{array}\right]\leqslant\mathsf{negl}(\lambda)$$

4 Witness Encryption for Linear Pairing Relation

In this section we present a core building block of our protocol, namely a Witness Encryption (WE) for Groth16, under the intermediate assumption that one proof element, π_1, is known to the Encryptor. Looking ahead, even though this is an unnatural assumption, combined with a garbling scheme it will constitute the cryptographic core of our full BABE construction.

We recall the definition of Extractable Witness Encryption [GKP+13]. Informally, an extractable witness encryption scheme for a relation $\mathcal{R}$ allows one to encrypt a message under an NP statement $\mathbb{x}$ such that anyone can decrypt if they *know* the corresponding witness $\mathbb{w}$, such that $(\mathbb{x},\mathbb{w})\in\mathcal{R}$.[9]

Let any R1CS relation $\mathcal{R}$ and a Groth16 proof system for $\mathcal{R}$. Let $\mathsf{crs}\leftarrow\mathsf{Groth16.}\mathtt{Gen}(\mathcal{R})$. We define the following relation. The statement of this relation $\mathcal{R}'$ consists of the statement $\boldsymbol{x}$ of the R1CS relation $\mathcal{R}$, and a specific proof element π_1 of a Groth16 proof for $\boldsymbol{x}$.

$$\mathcal{R}'=\left\{\big((\boldsymbol{x},\pi_1);\boldsymbol{w}\big):\ (\boldsymbol{x},\boldsymbol{w})\in\mathcal{R}\wedge\exists\pi_2,\pi_3\text{ s.t. }\mathsf{Groth16.}\mathtt{Verify}(\mathsf{crs},\boldsymbol{x},(\pi_1,\pi_2,\pi_3))=1\right\}\tag{18}$$

Construction 1 (WE for Groth16 with known π_1). *Let $\mathcal{R}$ be an R1CS relation. Below is* $\mathsf{WE}=(\mathtt{Gen},\mathtt{Enc},\mathsf{Dec})$*, our witness encryption scheme for $\mathcal{R}'$:*
$\mathtt{Gen}(\mathcal{R})$*: Output* $\mathsf{crs}=\mathsf{Groth16.}\mathtt{Gen}(\mathcal{R})$
$\mathtt{Enc}(\mathsf{crs},(\boldsymbol{x},\pi_1),\mathsf{msg})$*: Samples* $r\leftarrow\$\ \mathbb{F}_q^*$ *and executes the two sub-algorithms:*

- $\mathtt{Enc}_{\mathsf{setup}}(\mathsf{crs},\boldsymbol{x},\mathsf{msg};r)$*: Set* $Y:=e\big([\alpha]_1,[\beta]_2\big)+e\big(X,[\gamma]_2\big)$ *and* $X:=\sum_{i=1}^{\ell}x_i\left[\frac{\beta a_i(\tau)+\alpha b_i(\tau)+c_i(\tau)}{\gamma}\right]_1$ *and output* $\mathsf{ct}_{\mathsf{setup}}=(r[\delta]_2,\mathsf{RO}(rY)+\mathsf{msg})$
- $\mathtt{Enc}_{\mathsf{prove}}(\mathsf{crs},\pi_1;r)$: *Output* $\mathsf{ct}_{\mathsf{prove}}=r\pi_1$
- *Outputs* $(\mathsf{ct}_{\mathsf{prove}},\mathsf{ct}_{\mathsf{setup}})$

$\mathsf{Dec}(\mathsf{crs},(\boldsymbol{x},\pi_1),\mathsf{ct},\boldsymbol{w})$*: Parse* $\mathsf{ct}:=(\mathsf{ct}_{\mathsf{prove}},\mathsf{ct}_{\mathsf{setup}})=(\mathsf{ct}_1,(\mathsf{ct}_2,\mathsf{ct}_3))$*. Compute* $\pi_1',\pi_2,\pi_3=\mathsf{Groth16.}\mathtt{Prove}(\mathsf{crs},\boldsymbol{x},\boldsymbol{w})$ *and output* $\mathsf{msg}=\mathsf{ct}_3-\mathsf{RO}(e(\mathsf{ct}_1,\pi_2)-e(\pi_3,\mathsf{ct}_2))$.

We use the deterministic (non-rerandomized) variant of $\mathsf{Groth16.}\mathtt{Prove}$, so Dec recomputes the same π_1; equivalently Dec may take (π_2,π_3) directly.

Lemma 1 (Correctness of WE for canonical Groth16 proofs). *The construction in Constr. 1 satisfies the following correctness property: for every R1CS relation $\mathcal{R}$, every $(\boldsymbol{x},\boldsymbol{w})\in\mathcal{R}$, and every message* msg*, the following experiment outputs* msg*:*

$$\begin{aligned}\mathsf{crs}&\leftarrow\mathsf{Groth16.}\mathtt{Gen}(\mathcal{R}),\\(\pi_1,\pi_2,\pi_3)&\leftarrow\mathsf{Groth16.}\mathtt{Prove}(\mathsf{crs},\boldsymbol{x},\boldsymbol{w}),\\\mathsf{ct}&\leftarrow\mathtt{Enc}(\mathsf{crs},(\boldsymbol{x},\pi_1),\mathsf{msg}),\\\mathsf{Dec}(\mathsf{crs},(\boldsymbol{x},\pi_1),\mathsf{ct},\boldsymbol{w})&=\mathsf{msg}.\end{aligned}$$

[9] In this paper, we denote the statement and witness used in a SNARK as $\boldsymbol{x},\boldsymbol{w}$, and those used in witness encryption as $\mathbb{x},\mathbb{w}$.

Proof. Write $\mathsf{ct} = (\mathsf{ct}_1, (\mathsf{ct}_2, \mathsf{ct}_3))$. By construction, $\mathsf{ct}_1 = r\pi_1$, $\mathsf{ct}_2 = r[\delta]_2$, and $\mathsf{ct}_3 = \mathsf{RO}(rY) + \mathsf{msg}$. Since we use deterministic Groth16 proving, decryption recomputes the same (π_1, π_2, π_3). Groth16 correctness and its verification equation give $Y = e(\pi_1, \pi_2) - e(\pi_3, [\delta]_2)$. Therefore, by bilinearity,

$$e(\mathsf{ct}_1, \pi_2) - e(\pi_3, \mathsf{ct}_2) = r\big(e(\pi_1, \pi_2) - e(\pi_3, [\delta]_2)\big) = rY.$$

Thus decryption subtracts $\mathsf{RO}(rY)$ from ct_3 and outputs msg. □

Security of WE. In the following we show that this is effectively a witness encryption scheme for the relation $\mathcal{R}$, in the sense that the Decryptor cannot learn any information about msg unless she knows a valid witness $\boldsymbol{w}$ for the R1CS relation $\mathcal{R}$. Witness Encryption was originally introduced by Garg et al. [GGSW13] with a weaker security property that, roughly, the scheme is secure if there exists no $\mathbb{w}$ such that $(\mathbb{x}, \mathbb{w}) \in \mathcal{R}$. In this work, we make use of the stronger version of extractable witness encryption, that can be secure even if $\mathbb{x} \in \mathcal{R}$ (but the adversary doesn't know the witness). Throughout the paper, we, nevertheless, sometimes abuse the terminology and call it 'WE', omitting the 'extractable'.

In fact, we prove a stronger notion of security for our scheme, where an adversary adaptively chooses one part of the statement, π_1, after receiving a (partial) ciphertext, $\mathsf{ct}_{\mathsf{setup}}$ corresponding to $\boldsymbol{x}$. The type of adaptivity is not captured by the standard WE definition [GGSW13]; it is rather idiosyncratic to our scheme, as our encryption algorithm is explicitly split into two phases (setup and prove). We define this type of adaptive extractable one-way security for our scheme directly in Def. 11 below. Looking ahead, we will need this type of upgraded security for our full BABE protocol.

We prove security in the generic bilinear group model (see Sec. 2.2.1) and random oracle model. Thus, the extractor's access to $\mathsf{View}_{\mathcal{A}}$ includes the adversary's queries and the corresponding responses from both oracles. The random oracle ensures that the extractor gets access to the group element that "randomizes" the message (here rY). Otherwise, from the point of view of the adversary the ciphertext perfectly hides the message msg. So, intuitively, the only way for the adversary to learn anything from the ciphertext is to query rY and learn the masking term. But then the extractor gets access to rY as well.

From there we use a standard generic group model argument to extract valid proof elements π_2, π_3. Finally, we use the extractor of Groth16 to extract the original witness $\boldsymbol{w}$.

Remark 2. *As discussed in Sec. 1.3.2, our construction falls in the framework of witness encryption for general linear pairing relations introduced in [BC16, GKPW24] and subsequently formalized in [GHK+25]. Despite that, the security of our scheme (even the non-adaptive version) cannot be directly inferred from these prior works: [BC16] insists on relations in the standard model (does not capture Groth16), [GKPW24] identifies the general paradigm for arbitrary relations in the GGM, but provides security proof only for their particular instantiation and [GHK+25] formalizes the idea under a slightly different abstraction (linearly verifiable SNARKs). Consequently, we provide a complete security proof of our WE scheme for our specific relation.*

Definition 11 (Adaptive Extractable One-Way Security of WE for $\mathcal{R}'$)**.** *Let $\mathcal{M} = \mathcal{M}_\lambda$ be the message space, and assume that $1/|\mathcal{M}| = \mathsf{negl}(\lambda)$. A witness encryption scheme $\mathsf{WE} = (\mathtt{Gen}, (\mathtt{Enc}_{\mathsf{setup}}, \mathtt{Enc}_{\mathsf{prove}}), \mathsf{Dec})$ for $\mathcal{R}'$ is adaptively extractably one-way secure if, for every stateful PPT adversary $\mathcal{A}$, there exists a PPT extractor $\mathcal{E}_{\mathcal{A}}$ such that, for every benign auxiliary input $\mathsf{aux} = (\mathsf{aux}_1, \mathsf{aux}_2, \mathsf{aux}_3) \in \{0,1\}^{\mathsf{poly}(\lambda)}$,*

$$\Pr\left[\begin{array}{c} \mathsf{msg} = \mathsf{msg}' \\ \wedge\, (\boldsymbol{x}, \boldsymbol{w}) \notin \mathcal{R} \end{array} : \begin{array}{r} \mathsf{crs} \leftarrow \mathtt{Gen}(\mathcal{R}) \\ \boldsymbol{x} \leftarrow \mathcal{A}(\mathsf{crs}, \mathsf{aux}_1) \\ \mathsf{msg} \leftarrow\!\$\, \mathcal{M}, \quad r \leftarrow\!\$\, \mathbb{F}_q^* \\ \mathsf{ct}_{\mathsf{setup}} \leftarrow \mathtt{Enc}_{\mathsf{setup}}(\mathsf{crs}, \boldsymbol{x}, \mathsf{msg}; r) \\ \pi_1 \leftarrow \mathcal{A}(\mathsf{ct}_{\mathsf{setup}}, \mathsf{aux}_2) \\ \mathsf{ct}_{\mathsf{prove}} \leftarrow \mathtt{Enc}_{\mathsf{prove}}(\mathsf{crs}, \pi_1; r) \\ \mathsf{msg}' \leftarrow \mathcal{A}(\mathsf{ct}_{\mathsf{prove}}, \mathsf{aux}_3) \\ \boldsymbol{w} \leftarrow \mathcal{E}_{\mathcal{A}}(\mathsf{View}_{\mathcal{A}}) \end{array}\right] \leqslant \mathsf{negl}(\lambda). \tag{19}$$

Lemma 2. *Our* WE *construction for $\mathcal{R}'$ satisfies the adaptive extractable one-way security of Def. 11 in the generic bilinear group and random oracle models.*

Proof. Let $\mathcal{A}$ be a PPT adversary in the adaptive extractable one-way security experiment of Def. 11. We will construct an extractor $\mathcal{E}_{\mathcal{A}}$ that, given $\mathsf{View}_{\mathcal{A}}$, outputs a valid witness $\boldsymbol{w}$, i.e. $\boldsymbol{w} \leftarrow \mathcal{E}_{\mathcal{A}}(\mathsf{View}_{\mathcal{A}})$, such that $\mathcal{R}(\boldsymbol{x}, \boldsymbol{w}) = 1$.

Let $W := (\mathsf{msg} = \mathsf{msg}')$ be the message-recovery event, let $F := ((\boldsymbol{x}, \boldsymbol{w}) \notin \mathcal{R})$ be the extraction-failure event, and let Q be the event that $\mathcal{A}$ queries the random oracle at rY. Splitting according to Q gives

$$\Pr[W \wedge F] = \Pr[W \wedge F \wedge Q] + \Pr[W \wedge F \wedge \neg Q] \tag{20}$$
$$\leqslant \Pr[F \wedge Q] + \Pr[W \wedge \neg Q]. \tag{21}$$

If Q does not occur, then the value $\mathsf{RO}(rY)$ is uniformly random and information-theoretically hides msg in $\mathsf{ct}_{\mathsf{setup}}$. Consequently, conditioned on $\neg Q$, the entire view $\mathsf{View}_{\mathcal{A}}$ is independent of msg. Since msg is uniform in $\mathcal{M}$, it follows that

$$\Pr[W \wedge \neg Q] \leqslant \frac{1}{|\mathcal{M}|} = \mathsf{negl}(\lambda). \tag{22}$$

Now in the case where rY was queried by $\mathcal{A}$ we argue that, unless with negligible probability, $\mathcal{E}_{\mathcal{A}}$ can use the generic group oracle queries in $\mathsf{View}_{\mathcal{A}}$ to extract a valid Groth16 proof $\pi = (\pi_1, \pi_2, \pi_3)$, i.e. $\mathsf{Groth16.Verify}(\mathsf{crs}, \boldsymbol{x}, \pi) = 1$. We start from the symbolic group representation (see Sec. 2.2.1). Let $\mathfrak{r}$ be the symbolic variable of r, then information-theoretically the monomial $Y\mathfrak{r}$ can only be obtained from ct_1 and ct_2 queries, because no other element that the adversary receives contains the variable $\mathfrak{r}$. In particular, π_1 is $\mathfrak{r}$-free: $\mathsf{ct}_{\mathsf{setup}}$ contains no $\mathbb{G}_1$ handle involving $\mathfrak{r}$, and generic group operations and pairings provide no map from $\mathbb{G}_2$ or $\mathbb{G}_T$ back to $\mathbb{G}_1$. Therefore:

$$(Y)\mathfrak{r} = (a_2\pi_1)\mathfrak{r} + (a_3)\mathfrak{r}\mathfrak{d} \Rightarrow Y = a_2\pi_1 + a_3\mathfrak{d}$$

for a_2, a_3 chosen by the adversary ($\mathfrak{d}$ is the formal variable for δ).

Switching to the actual generic group model:

$$Y = e(\pi_1, [a_2]_2) + e([a_3]_1, [\delta]_2)$$

except with probability $\mathsf{negl}_{\mathrm{GGM}}(\lambda)$, by the Master Theorem (see Thm. 1). Recall that

$$Y := e\big([\alpha]_1, [\beta]_2\big) + e\Big(\sum_{i=1}^{\ell} x_i \left[\frac{\beta a_i(\tau) + \alpha b_i(\tau) + c_i(\tau)}{\gamma}\right]_1, [\gamma]_2\Big)$$

so π_1, $\pi_2 = [a_2]_2, \pi_3 = -[a_3]_1$ are valid Groth16 proof elements that $\mathcal{E}_{\mathcal{A}}$ can compute from the generic group oracle queries in $\mathsf{View}_{\mathcal{A}}$.

Consider the derived Groth16 adversarial Prover $\mathcal{P}^*$ that follows this execution of $\mathcal{A}$ and outputs $(\boldsymbol{x}, \pi)$, where $\pi := (\pi_1, [a_2]_2, -[a_3]_1)$. Since $\mathcal{P}^*$ only adds this deterministic post-processing, its execution view is efficiently reconstructible from $\mathsf{View}_{\mathcal{A}}$. Therefore, $\mathcal{E}_{\mathcal{A}}$ can invoke the extractor $\mathcal{E}_{\mathcal{P}^*}$ guaranteed by Groth16 knowledge soundness with the same execution view and extract a valid witness $\boldsymbol{w}$, unless with a negligible probability $\mathsf{negl}_{\text{Groth16-KS}}(\lambda)$.

The preceding extraction argument shows that

$$\Pr[F \wedge Q] \leqslant \mathsf{negl}_{\mathrm{GGM}}(\lambda) + \mathsf{negl}_{\text{Groth16-KS}}(\lambda). \tag{23}$$

Combining Eqs. (21) to (23),

$$\Pr[W \wedge F] \leqslant \frac{1}{|\mathcal{M}|} + \mathsf{negl}_{\text{GGM}}(\lambda) + \mathsf{negl}_{\text{Groth16-KS}}(\lambda) \tag{24}$$

$$= \mathsf{negl}(\lambda). \tag{25}$$

□

5 Garbled Circuit for BN254 Scalar Multiplication

5.1 Overview

Goal. The objective of this section is to construct a compact garbled circuit for BN254 scalar multiplication with a *hard-coded secret* scalar. Concretely, the garbling algorithm fixes a secret $r \in \mathbb{F}_q^*$, and the Prover provides a public input point $\pi \in \mathbb{G}$ (authenticated on-chain via a Lamport signature). The garbled evaluation outputs the group element $f_r(\pi) = r\pi$.

Bitcoin constraint. BABE's proving phase is between mutually distrustful parties. Every proving-phase message must appear on-chain because an off-chain crash is indistinguishable from Byzantine deviation, whereas an on-chain transcript makes deviations detectable and the parties accountable. Thus the input encodings produced by Encode — labels the Verifier supplies after the Prover commits to π — must be revealed on-chain.

Bitcoin script constrains the *form* of these encodings. The Prover commits to π on-chain with a Lamport signature ($\mathsf{CheckLampSig}$, Eq. (9)); the Verifier then reveals the labels, and $\mathsf{CheckLampSigsMatch}$ (Eq. (10))uses Hash_{BTC} hashlocks, implemented by `OP_HASH160`, to ensure that they encode the same π. Since the script can check only these hashlock/Lamport relations, each label must be affine in one coordinate bit. Writing

$$(x_0(\pi), \ldots, x_{n-1}(\pi), y_0(\pi), \ldots, y_{n-1}(\pi)) \in \{0,1\}^{2n}$$

for the bit-decomposition of $(x(\pi), y(\pi))$, a label for any bit z has the form L_z^0 if $z = 0$ and L_z^1 if $z = 1$. We call this requirement the *Bitcoin constraint*.

Input representation. The coordinate-bit labels above supply the public point $\pi = (x(\pi), y(\pi)) \in \mathbb{F}_p^2$ to the garbled circuit. From π, the circuit derives a small collection of algebraic features

$$\boldsymbol{u}(\pi) = \big(1,\ x(\pi),\ y(\pi),\ x(\pi)^2,\ y(\pi)^2,\ x(\pi)y(\pi)\big) \in \mathbb{F}_p^6,$$

and we also work with its bit-decomposition $\overline{\boldsymbol{u}}(\pi) \in \{0,1\}^\ell$ where $\ell = 1 + 5n$ and $n = \lceil \log_2 p \rceil$. The linear map $\boldsymbol{G} \in \mathbb{F}_p^{6\times\ell}$ reconstructs $\boldsymbol{u}(\pi)$ from $\overline{\boldsymbol{u}}(\pi)$, i.e., $\boldsymbol{u}^T(\pi) = \boldsymbol{G} \times \overline{\boldsymbol{u}}^T(\pi)$. Details in Sec. 5.2.

Observation: group addition becomes an inner product. A central bottleneck for naively garbling $r\pi$ is that scalar multiplication entails many elliptic-curve additions/doublings, which are prohibitively expensive inside a Boolean garbled circuit. Our first observation, as made by [EL26] is that, for BN254, addition of an input point $\pi \in \mathbb{G}$ to a fixed point $\phi \in \mathbb{G}$ (both in affine coordinates) can be expressed as *linear functions* of the feature vector $\boldsymbol{u}(\cdot)$. More precisely, assuming $x(\phi) \neq x(\pi)$, Lem. 4 shows that Jacobian coordinates of $\pi + \phi$ can be written as

$$(X, Y, Z) = \boldsymbol{A}(\phi) \times \boldsymbol{u}(\pi)^T,$$

for a matrix $\boldsymbol{A}(\phi) \in \mathbb{F}_p^{3\times 6}$ that depends only on ϕ. Since the Prover sees π only in bit form, we immediately lift this to the bit-decomposed domain using $\boldsymbol{G}$, yielding an equivalent form

$$(X, Y, Z) \;=\; \boldsymbol{B}(\phi) \times \overline{\boldsymbol{u}}(\pi)^T \qquad \text{where} \qquad \boldsymbol{B}(\phi) = \boldsymbol{A}(\phi) \times \boldsymbol{G} \in \mathbb{F}_p^{3\times \ell}.$$

Finally, because we will compute $\delta\pi + \phi$ where δ is a bit, we use the bit-gated variant (Lem. 5), which yields a matrix form $\boldsymbol{D}(\delta, \phi) \times \overline{\boldsymbol{u}}(\pi)^T$ for $\delta\pi + \phi$ where the coefficient matrix $\boldsymbol{D}$ depends only on (δ, ϕ).

Decomposable Randomized Encodings (DREs) To garble the resulting linear function $\boldsymbol{D}(\delta, \phi) \times \overline{\boldsymbol{u}}(\pi)^T$, we use a decomposable randomized encoding (DRE) for linear functions [Ish13]. A DRE for a function $f : X_1 \times \ldots \times X_n \to Y$ is a randomized encoding

$$\hat{f}(x = (x_1, \ldots, x_\ell), \rho) = (\hat{f}_1(x_1, \rho), \ldots, \hat{f}_\ell(x_\ell, \rho))$$

with randomness ρ such that there exists a decoder $\mathsf{Dec}(\hat{f}(x, \rho)) = f(x)$ but $\hat{f}(x, \rho)$ does not reveal any additional information about x (Defs. 12 and 13). Unlike garbled circuits, decoding a linear DRE requires no ciphertexts.

High-level plan. At a high level, we linearize the field operations (using the feature vector $\overline{\boldsymbol{u}}$), use the DRE decoding to compute all linear field- and group-level operations (which are expensive in a boolean circuit), and use the boolean garbled circuit for only low-degree non-linear operations:

1. Step 1: Create a *DRE for weighted group sum* to compute the scalar multiplication $r\pi$.
2. Step 2: Create another DRE to compute the DRE used in Step 1.
3. Step 3: Create a boolean garbled circuit for (i) validating the input is a valid Elliptic curve point, (ii) deriving the feature vector $\overline{\boldsymbol{u}}$, and (iii) computing the DRE used in Step 2.

Step 1: DRE for the weighted group sum with public weights. To efficiently do scalar multiplication, we take $r = \sum_{i=0}^{n-1} r_i 2^i$, the binary expansion of the hard-coded scalar $r \in \mathbb{F}_q^*$. Using the DRE for weighted summation in abelian groups (Lem. 9), we obtain a DRE for the function

$$f(\, r_0\pi, \ldots, r_{n-1}\pi \,) = \sum_{i=0}^{n-1} 2^i (r_i \pi) = r\pi.$$

The DRE for this function is

$$\hat{f}(\, r_0\pi, \ldots, r_{n-1}\pi \,) = (r_0\pi + \rho_0, \ldots, r_{n-1}\pi + \rho_{n-1} \,)$$

where $\rho_0, \ldots, \rho_{n-1}$ are sampled uniformly from $\mathbb{G}_1 \setminus \{\mathcal{O}\}$ subject to $\sum_{i=0}^{n-1} 2^i \rho_i = 0$.[10] This DRE shifts this *weighted summation* (here $2^0, \ldots, 2^{n-1}$ are the weights) from the circuit to the Prover: the Prover simply computes the weighted sum of the elements of $\hat{f}$ to obtain $r\pi$.

[10] Note that in the actual construction, we also need to randomize the representation on Jacobian coordinates. We refer the reader to Remark 3 and Lem. 6 for more details.

Step 2: DRE for producing each $r_i\pi + \rho_i$ from $\overline{\boldsymbol{u}}(\pi)$. This step creates a DRE to compute $r_i\pi + \rho_i$ for random masks ρ_i used in Step 1. Here we use Lem. 5 to express $r_i\pi + \rho_i$ in Jacobian coordinates as a linear map of $\overline{\boldsymbol{u}}(\pi)$:

$$(X_i, Y_i, Z_i) \;=\; \boldsymbol{D}(r_i, \rho_i) \times \overline{\boldsymbol{u}}(\pi)^T,$$

for a matrix $\boldsymbol{D}(r_i, \rho_i) \in \mathbb{F}_p^{3\times \ell}$ depending only on the private inputs (r_i, ρ_i). Each of the three coordinates is thus an inner product between a *known* coefficient vector and the *bit-vector* $\overline{\boldsymbol{u}}(\pi)$. We then apply a DRE for private affine functions (Lem. 10) to create a DRE for these inner products. The DRE takes the form

$$\hat{h}_k(r, \overline{u}_k, \omega) = \{\overline{u}_k \cdot D_{i,j,k} + s_{i,j,k}\}_{i\in\{0,\ldots,n-1\}, j\in\{1,2,3\}} \quad \text{for} \quad k \in [5n+1]$$

where the shared secret input $r \in \mathbb{F}_q$ supplies the scalar (and hence r_i as the bits of r); $\overline{u}(\pi)$ is the decomposable input; $D_{i,j,k}$ is the (j,k)-th entry of $\boldsymbol{D}(r_i, \rho_i)$ as defined above; and the randomness is $(\rho, s) \in \mathrm{P} \times S$ (subject to $\sum_{k=1}^{5n+1} s_{i,j,k} = 0$ for all $i \in \{0,\ldots,n-1\}$, $j \in \{1,2,3\}$). The complete DRE construction is given in Thm. 4.

Step 3: The Boolean Circuit. The garbled circuit is formally defined in Constr. 2.

Garbling: Outputs a Lamport secret key (the encoding key)

$$\mathsf{ek} = \begin{pmatrix} L^0_{x,0} & \ldots & L^0_{x,n-1} & L^0_{y,0} & \ldots & L^0_{y,n-1} \\ L^1_{x,0} & \ldots & L^1_{x,n-1} & L^1_{y,0} & \ldots & L^1_{y,n-1} \end{pmatrix}$$

and the garbled circuit $\mathsf{ct}_{\mathsf{GC}}$ computed in three steps:

1. **Elliptic curve validation:** The garbled circuit first verifies that the input point $\pi = (x(\pi), y(\pi))$ lies on the curve E, i.e., it checks that $y(\pi)^2 \equiv x(\pi)^3 + 3 \pmod p$.

2. **Binary decomposition:** If the validation passes, the circuit computes the binary decomposition $\overline{\boldsymbol{u}}(\pi) \in \{0,1\}^{1+5n}$ of $\boldsymbol{u}(\pi) = (1, x(\pi), y(\pi), x(\pi)^2, y(\pi)^2, x(\pi)y(\pi))$. These two steps are implemented using a privacy-free Boolean garbled circuit. Let $L^{\overline{u}_k(\pi)}_{\overline{u},k}$ for $k \in [\ell]$ be the output labels of this step.

3. **DRE encoding:** With shared secret input $r \in \mathbb{F}_q$, for each $k \in \{1,\ldots,\ell\}$ where $\ell = 1 + 5n$, the garbled circuit comes hardwired with the DRE encoding from Thm. 4. Specifically, the garbled circuit contains an encryption of $\hat{h}_k(r, 0, \omega)$ under the 0-label $L^0_{\overline{u},k}$ and an encryption of $\hat{h}_k(r, 1, \omega)$ under the 1-label $L^1_{\overline{u},k}$ for $\overline{u}_k(\pi)$.

Evaluation: The evaluator, given the input labels

$$L = \left(L^{x(\pi)_0}_{x,0} \quad \ldots \quad L^{x(\pi)_{n-1}}_{x,n-1} \quad L^{y(\pi)_0}_{y,0} \quad \ldots \quad L^{y(\pi)_{n-1}}_{y,n-1}\right),$$

evaluates the Boolean circuits to obtain labels for $\overline{u}(\pi)$, and decrypts the appropriate encryptions to obtain $\hat{h}_k(r, \overline{u}_k(\pi), \omega)$ for each k. The evaluator then decodes the DRE as follows (Thm. 4):

- Recast $\hat{h}_k(r, \overline{u}_k(\pi), \omega)$ as $\{t_{i,j,k}\}_{i\in\{0,\ldots,n-1\}, j\in\{1,2,3\}}$. Then the Prover aggregates shares to cancel the $s_{i,j,k}$-terms: $\widehat{t}_{i,j} = \sum_{k=1}^{5n+1} t_{i,j,k}$.
- Interpreting $(\widehat{t}_{i,1}, \widehat{t}_{i,2}, \widehat{t}_{i,3})$ as the Jacobian coordinates of a masked point Q_i, the final output is obtained by the usual bit-weighted recombination $\mathrm{out} = \sum_i 2^i Q_i$. Since $Q_i = r_i\pi + \rho_i$ and $\sum_i 2^i \rho_i = 0$, then $\mathrm{out} = r\pi$.

Intuitively, the heavy work (group operations and most algebra) is pushed to the Prover and to the Verifier's offline preprocessing, while the online garbled circuit mainly performs validation, low-level arithmetic for feature extraction, and symmetric-key decryptions for table selection to output the precomputed DRE components.

5.2 Elliptic Curve Addition and Notation

Let $E/\mathbb{F}_p$ be the BN254 curve in short Weierstrass form

$$E : y^2 = x^3 + 3 \pmod{p},$$

where p is the BN254 base-field prime. Let $n = \lceil \log_2 p \rceil$ denote the bit length of $\mathbb{F}_p$ elements (so $x(\pi)$, $y(\pi)$ have n bits each). For BN254 we also have $n = \lceil \log_2 q \rceil$, the bit length of $\mathbb{F}_q^*$, where q is the BN254 group order.

Let π be a point on the curve, represented in affine coordinates $\pi := (x(\pi), y(\pi))$ with $x(\pi), y(\pi) \in \mathbb{F}_p$. Then we define $\boldsymbol{u}(\pi) = (1, x(\pi), y(\pi), x(\pi)^2, y(\pi)^2, x(\pi)y(\pi))$ and $\overline{\boldsymbol{u}}(\pi) \in \{0,1\}^\ell$ be the binary decomposition of $\boldsymbol{u}(\pi)$, where $\ell = 5n + 1$. More precisely, if $\boldsymbol{u}(\pi) = (u_0, u_1, u_2, u_3, u_4, u_5)$ where $u_0 = 1$, $u_1 = x(\pi)$, $u_2 = y(\pi)$, $u_3 = x(\pi)^2$, $u_4 = y(\pi)^2$, and $u_5 = x(\pi)y(\pi)$, then:

$$\overline{\boldsymbol{u}}(\pi) = (1, u_{1,0}, \ldots, u_{1,n-1}, u_{2,0}, \ldots, u_{2,n-1}, u_{3,0}, \ldots, u_{3,n-1}, u_{4,0}, \ldots, u_{4,n-1}, u_{5,0}, \ldots, u_{5,n-1})$$

where $(u_{i,0}, \ldots, u_{i,n-1})$ is the binary decomposition of u_i for $i \in \{1,2,3,4,5\}$ (i.e., $u_i = \sum_{j=0}^{n-1} u_{i,j} 2^j$). We use the notation $u_i(\pi)$ and $\overline{u}_i(\pi)$ to denote the ith component of $\boldsymbol{u}(\pi)$ and $\overline{\boldsymbol{u}}(\pi)$, respectively.

Further, we define $\boldsymbol{G} \in \mathbb{F}_p^{6\times(1+5n)}$ such that $\boldsymbol{G} \times \overline{\boldsymbol{u}}^T(\pi) = \boldsymbol{u}^T(\pi)$. The matrix $\boldsymbol{G}$ is structured as:

$$\boldsymbol{G} = \left(\begin{array}{c|ccccc} 1 & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \hline 0 & \mathbf{2} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ 0 & \mathbf{0} & \mathbf{2} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ 0 & \mathbf{0} & \mathbf{0} & \mathbf{2} & \mathbf{0} & \mathbf{0} \\ 0 & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{2} & \mathbf{0} \\ 0 & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{2} \end{array}\right)$$

where $\mathbf{0}$ denotes a row vector of n zeros, and $\mathbf{2}$ denotes the row vector $(2^0, 2^1, \ldots, 2^{n-1})$ of powers of 2. That is, the first column of $\boldsymbol{G}$ is $(1,0,0,0,0,0)^T$ (corresponding to the constant term $u_0 = 1$), and for each $i \in \{1,2,3,4,5\}$, columns $(i-1)n+2$ through $in+1$ contain the powers of 2 $(2^0, 2^1, \ldots, 2^{n-1})$ in row $i+1$, with zeros elsewhere. This ensures that $\boldsymbol{G} \times \overline{\boldsymbol{u}}^T(\pi)$ reconstructs the field elements from their binary decomposition.

The following lemmata Lem. 3, Lem. 4, Lem. 5, Lem. 6 formally define the elliptic curve addition and, further, $\phi + \delta\pi$ operation in Jacobian coordinates.

Lemma 3. *Let $\pi = (x(\pi), y(\pi)) \in E(\mathbb{F}_p) \setminus \{\mathcal{O}\}$ be given in affine coordinates, and let $\delta \in \{0,1\}$. Define $\boldsymbol{M}(\delta) \in \mathbb{F}_p^{3\times 6}$ by*

$$\boldsymbol{M}(\delta) = \begin{pmatrix} 1-\delta & \delta & 0 & 0 & 0 & 0 \\ 1-\delta & 0 & \delta & 0 & 0 & 0 \\ \delta & 0 & 0 & 0 & 0 & 0 \end{pmatrix}.$$

Then $\boldsymbol{M}(\delta) \times \boldsymbol{u}^T(\pi)$ is a Jacobian representation of the point $\delta\pi$: when $\delta = 1$ it equals $(x(\pi), y(\pi), 1)$, and when $\delta = 0$ it equals $(1,1,0)$, which represents $\mathcal{O}$ under the convention that any $(X, Y, 0)$ with $Y \neq 0$ encodes the point at infinity. Alternatively, $\boldsymbol{N}(\delta) \times \overline{\boldsymbol{u}}^T(\pi)$ is also equal to $\delta\pi$, where $\boldsymbol{N}(\delta) = \boldsymbol{M}(\delta) \times \boldsymbol{G}$.

The proof of this lemma is straightforward and is left to the reader.

Lemma 4. *Let $\phi, \pi \in E(\mathbb{F}_p) \setminus \{\mathcal{O}\}$ be two points in affine coordinates such that $x(\phi) \neq x(\pi)$. Then a Jacobian representation of $\phi + \pi$ is given by $\boldsymbol{A}(\phi) \times \boldsymbol{u}^T(\pi)$ where $\boldsymbol{A}(\phi) \in \mathbb{F}_p^{3\times 6}$ is a matrix that depends only on ϕ, defined by:*

$$\boldsymbol{A}(\phi) = \begin{pmatrix} 6 & x^2(\phi) & -2y(\phi) & x(\phi) & 0 & 0 \\ 9y(\phi) & 0 & -(y(\phi)^2 + 9) & 3x(\phi)y(\phi) & y(\phi) & -3x^2(\phi) \\ -x(\phi) & 1 & 0 & 0 & 0 & 0 \end{pmatrix}$$

Equivalently, Jacobian coordinates of $\phi + \pi$ are also given by $\boldsymbol{B}(\phi) \times \overline{\boldsymbol{u}}^T(\pi)$ where $\boldsymbol{B}(\phi) \in \mathbb{F}_p^{3\times(1+5n)}$ is defined as $\boldsymbol{B}(\phi) = \boldsymbol{A}(\phi) \times \boldsymbol{G}$.

Proof. Let $\phi = (x_1, y_1)$ and $\pi = (x_2, y_2)$ be affine points on

$$E/\mathbb{F}_p : \quad y^2 = x^3 + 3,$$

and assume $x_1 \neq x_2$. Define

$$Z \stackrel{\text{def}}{=} x_2 - x_1 \in \mathbb{F}_p^*.$$

Consider the following projective (Jacobian-style) representation of $\phi + \pi$ [Was08]:

$$\begin{aligned} X_3 &\stackrel{\text{def}}{=} (y_2 - y_1)^2 - (x_1 + x_2)\, Z^2, \\ Y_3 &\stackrel{\text{def}}{=} (y_2 - y_1)\,(x_1 Z^2 - X_3) - y_1 Z^3, \\ Z_3 &\stackrel{\text{def}}{=} Z. \end{aligned} \tag{26}$$

When $Z \neq 0$, the triple $(X_3 : Y_3 : Z_3)$ represents the affine point

$$\left(\frac{X_3}{Z_3^2}, \frac{Y_3}{Z_3^3}\right)$$

and one checks that it equals $\phi + \pi$ (this is the standard affine-addition formula written in projective coordinates; see, e.g., the explicit projective addition formulas for short Weierstrass curves).

We now expand (26) and simplify using the curve equations

$$y_1^2 = x_1^3 + 3, \qquad y_2^2 = x_2^3 + 3.$$

Derivation of Z_3. By definition, $Z_3 = Z = x_2 - x_1$, which is already linear in $\boldsymbol{u}(\pi)$.

Derivation of X_3. First expand:

$$\begin{aligned} X_3 &= (y_2 - y_1)^2 - (x_1 + x_2)(x_2 - x_1)^2 \\ &= (y_2^2 - 2y_1y_2 + y_1^2) - (x_1 + x_2)(x_2^2 - 2x_1x_2 + x_1^2) \\ &= y_2^2 - 2y_1y_2 + y_1^2 - \big(x_1^3 - x_1^2x_2 - x_1x_2^2 + x_2^3\big). \end{aligned}$$

Substitute $y_1^2 = x_1^3 + 3$ and $y_2^2 = x_2^3 + 3$ to obtain cancellation of the x_1^3 and x_2^3 terms:

$$X_3 = (x_2^3 + 3) - 2y_1y_2 + (x_1^3 + 3) - x_1^3 + x_1^2x_2 + x_1x_2^2 - x_2^3,$$

hence

$$X_3 = 6 + x_1^2x_2 - 2y_1y_2 + x_1x_2^2. \tag{27}$$

Derivation of Y_3. Start from (26):

$$Y_3 = (y_2 - y_1)\,(x_1Z^2 - X_3) - y_1Z^3.$$

Using $Z = x_2 - x_1$ and reducing modulo the curve equations gives

$$Y_3 \equiv -3x_1^2x_2y_2 + 3x_1x_2^2y_1 - y_1^2y_2 + y_1y_2^2 + 9(y_1 - y_2).$$

Regrouping this expression to match the $\boldsymbol{u}(\pi)$ basis yields:

$$Y_3 = 9y_1 \;-\; (y_1^2 + 9)\,y_2 \;+\; 3x_1y_1\,x_2^2 \;+\; y_1\,y_2^2 \;-\; 3x_1^2\,(x_2y_2). \tag{28}$$

Matrix form. Putting together the above, we have that $\boldsymbol{A}(\phi) \times \boldsymbol{u}^T(\pi)$ equals the Jacobian coordinates of $\phi + \pi$ where

$$\boldsymbol{A}(\phi) = \begin{pmatrix} 6 & x_1^2 & -2y_1 & x_1 & 0 & 0 \\ 9y_1 & 0 & -(y_1^2+9) & 3x_1y_1 & y_1 & -3x_1^2 \\ -x_1 & 1 & 0 & 0 & 0 & 0 \end{pmatrix}.$$

The alternative form $\boldsymbol{B}(\phi) \times \overline{\boldsymbol{u}}^T(\pi)$ follows immediately from $\boldsymbol{u}^T(\pi) = \boldsymbol{G} \times \overline{\boldsymbol{u}}^T(\pi)$ by setting $\boldsymbol{B}(\phi) = \boldsymbol{A}(\phi) \times \boldsymbol{G}$. This finishes the proof.

□

Lemma 5. *Let $\phi, \pi \in E(\mathbb{F}_p)$ be two points in affine coordinates such that $x(\phi) \neq x(\pi)$. Then the Jacobian coordinates of $\phi + \delta\pi$ where $\delta \in \{0,1\}$,[11] is given by $\boldsymbol{C}(\delta,\phi) \times \boldsymbol{u}^T(\pi)$ where $\boldsymbol{C}(\delta,\phi) \in \mathbb{F}_p^{3\times 6}$ is a matrix that depends only on δ, ϕ, defined by:*

$$\boldsymbol{C}(\delta,\phi) = \begin{pmatrix} 6\delta + (1-\delta)x(\phi) & \delta x^2(\phi) & -2\delta y(\phi) & \delta x(\phi) & 0 & 0 \\ 9\delta y(\phi) + (1-\delta)y(\phi) & 0 & -\delta(y(\phi)^2+9) & 3\delta x(\phi)y(\phi) & \delta y(\phi) & -3\delta x^2(\phi) \\ -\delta x(\phi) + (1-\delta) & \delta & 0 & 0 & 0 & 0 \end{pmatrix}$$

Equivalently, Jacobian coordinates of $\phi + \delta\pi$ are also given by $\boldsymbol{D}(\delta,\phi) \times \overline{\boldsymbol{u}}^T(\pi)$ where $\boldsymbol{D}(\delta,\phi) \in \mathbb{F}_p^{3\times(1+5n)}$ is defined as $\boldsymbol{D}(\delta,\phi) = \boldsymbol{C}(\delta,\phi) \times \boldsymbol{G}$.

Proof. The proof is analogous to the proof of Lem. 4. The only difference is that we need to consider the case where $\delta = 1$ and $\delta = 0$ separately. □

Lemma 6. *Given a point $\pi \in E(\mathbb{F}_p)$ in Jacobian coordinates (X, Y, Z) (which is not unique) we can obtain a uniform representation of the same point as $\mathsf{diag}(\lambda^2, \lambda^3, \lambda) \times (X, Y, Z)^T$ for λ sampled uniformly from $\mathbb{F}_p^*$ where $\mathsf{diag}(\lambda^2, \lambda^3, \lambda)$ is a diagonal matrix with the entries $\lambda^2, \lambda^3, \lambda$ on the diagonal. More precisely,*

$$\mathsf{diag}(\lambda) = \begin{pmatrix} \lambda^2 & 0 & 0 \\ 0 & \lambda^3 & 0 \\ 0 & 0 & \lambda \end{pmatrix}$$

The proof follows from the definitions of Jacobian representation and we skip the details.

5.3 Decomposable Randomized Encodings: Definitions and Preliminaries

Randomized encodings allow us to represent a function f by a simpler function $\hat{f}$ such that the encoding $\hat{f}(x;\rho)$ (where ρ is some randomness independent of x) reveals $f(x)$ but nothing else about x beyond what is revealed by $f(x)$ itself. This section presents the formal framework for decomposable randomized encodings, following the definitions (some of them verbatim) from Ishai [Ish13].

Definition 12 (Randomized encoding (Definition 3.1, [Ish13]))**.** *Let $X, Y, \hat{Y}, \mathrm{P}$ be finite sets and let $f : X \to Y$. A function $\hat{f} : X \times \mathrm{P} \to \hat{Y}$ is a* randomized encoding *of f if it satisfies:*

- *δ-**Correctness.** There exists a decoder (possibly randomized) $\mathsf{Dec} : \hat{Y} \to Y$ such that for all $x \in X$ and $\rho \in R$,*

  $$\Pr[\mathsf{Dec}(\hat{f}(x,\rho)) \neq f(x)] \leqslant \delta$$

  where randomness is over the choices of ρ and the decoder. We skip mentioning δ when it is a negligible function of the appropriate parameters.

[11] $\delta\pi$ denotes the point π if $\delta = 1$ and $\mathcal{O}$ if $\delta = 0$.

- *ϵ-**privacy.** There exists a randomized simulator* $\mathsf{Sim} : Y \to \hat{Y}$ *such that for all* $x \in X$*,*

$$\mathsf{Sim}(f(x)) \approx_\epsilon \hat{f}(x, \rho),$$

where $\rho \leftarrow \mathrm{P}$ *is uniform and* $\approx_\epsilon$ *denotes distributions that are statistically close. Again, we skip mentioning* ϵ *when it is a negligible function of the relevant parameters.*

Additionally, a randomized encoding is said to be *efficient* if its encoding and decoding complexities are polynomial in the size of the input. In the literature, various relaxations of the above definition are considered. For example, one may consider a randomized encoding that is only computationally or statistically private. However, we focus on the perfect privacy, as it suffices for our purposes.

Next we define the decomposability property and some useful properties of randomized encodings.

Definition 13 (Decomposable randomized encoding (DRE) (Definition 4.1, [Ish13]))**.** *For* $f : R \times X_1 \times \cdots \times X_n \to Y$*, a decomposable randomized encoding of* f *is one that has the form*

$$\hat{f}(r, x_1, \ldots, x_n, \rho) = (\hat{f}_1(r, x_1, \rho), \ldots, \hat{f}_n(r, x_n, \rho))$$

for some functions $\hat{f}_i : R \times X_i \times \mathrm{P} \to \hat{Y}_i$*, where* $r \in R$ *is a secret input that is shared across all decompositions (each part of the encoding may depend on* r*).*

The decomposability property is what makes DRE particularly useful: each input component x_i can be encoded independently (though the encodings may share a secret input s and randomness ρ), and the evaluation algorithm can reconstruct $f(x_1, \ldots, x_n)$ from these individual encodings. This structure is essential for the efficient garbled circuit constructions we present in the following sections.

Lemma 7 (Concatenation (Lemma 3.3, [Ish13]))**.** *Suppose* $\hat{f}_i(x, \rho_i)$ *is a randomized encoding of* $f_i(x)$ *for* $i = 1, \ldots, k$*. Then the function* $\hat{f}(x, (\rho_1, \ldots, \rho_k)) \stackrel{def}{=} (\hat{f}_1(x, \rho_1), \ldots, \hat{f}_k(x, \rho_k))$ *is a randomized encoding of* $f(x) \stackrel{def}{=} (f_1(x), \ldots, f_k(x))$*.*

Lemma 8 (Composition (Lemma 3.4, [Ish13]))**.** *Suppose* $\hat{f}(x, \rho)$ *is a randomized encoding of* $f(x)$ *and* $\hat{f}'((x, \rho), \rho')$ *is a randomized encoding of* $\hat{f}(x, \rho)$ *(viewing the latter as a deterministic function of* (x, ρ)*). Then* $\hat{f}''(x, (\rho, \rho')) \stackrel{def}{=} \hat{f}'((x, \rho), \rho')$ *is a randomized encoding of* $f(x)$*.*

5.4 Decomposable Randomized Encodings Constructions

Building on [FKN94], Ishai [Ish13] presents a simple decomposable encoding for summation in finite abelian groups. We state a slight generalization to *weighted summation* here.

Lemma 9 (DRE for weighted group summation (Generalization of Claim 4.2, [Ish13]))**.** *Let* $\mathcal{G}$ *be a finite abelian group and let* $f_{a_1,\ldots,a_n} : \mathcal{G}^n \to \mathcal{G}$ *be the group summation function* $f_{a_1,\ldots,a_n}(x_1, \ldots, x_n) = \sum_{i=1}^n a_i x_i$*.*[12] *Let* $R = \{(\rho_1, \ldots, \rho_n) \in \mathcal{G}^n : \sum_{i=1}^n a_i \rho_i = 0\}$*. Then the function* $\hat{f}_{a_1,\ldots,a_n} : \mathcal{G}^n \times R \to \mathcal{G}^n$ *defined by* $\hat{f}_{a_1,\ldots,a_n}((x_1, \ldots, x_n), (\rho_1, \ldots, \rho_n)) = (x_1 + \rho_1, \ldots, x_n + \rho_n)$ *is a decomposable encoding of* $f_{a_1,\ldots,a_n}$*.*

Proof. As in Ishai [Ish13], it is easy to verify that $\hat{f}_{a_1,\ldots,a_n}$ maps an input $x = (x_1, \ldots, x_n)$ to a uniformly random input $x' \in \mathcal{G}^n$ such that $f_{a_1,\ldots,a_n}(x') = f_{a_1,\ldots,a_n}(x)$. Thus, we can let $\mathsf{Dec} = f_{a_1,\ldots,a_n}$ and let $\mathsf{Sim}(y)$ output a random n-tuple in $f^{-1}_{a_1,\ldots,a_n}(y)$. □

[12] Where $+$ is the group operation and $a\rho = \overbrace{\rho + \ldots + \rho}^{a \text{ times}}$.

Remark 3. *We will actually use the above lemma on $f_{a_1,\ldots,a_n}(x_1,\ldots,x_n) = \sum_{i=1}^n a_i x_i$ with inputs $x_1 \ldots x_n$ in Jacobian representation which is not unique for every group element. Thus, we randomize the representations in the DRE. More specifically, $\hat{f}_{a_1,\ldots,a_n} : \mathbb{G}^n \times \mathbb{G}^n \times \mathbb{F}_p^{*n} \to \mathbb{G}^n$ defined by*

$$\hat{f}_{a_1,\ldots,a_n}((x_1,\ldots,x_n),(\rho_1,\ldots,\rho_n),(\lambda_1,\ldots,\lambda_n)) = (\mathsf{diag}(\lambda_1^2,\lambda_1^3,\lambda_1)\times(x_1+\rho_1)^T,\ldots,\mathsf{diag}(\lambda_n^2,\lambda_n^3,\lambda_n)\times(x_n+\rho_n)^T)$$

where $\mathsf{diag}(\lambda_i^2,\lambda_i^3,\lambda_i)$ *is the diagonal matrix with* $\lambda_i^2,\lambda_i^3,\lambda_i$ *on the diagonal as defined in Lem. 6.*

We will also need a DRE for private weights. However, we limit the inputs to be either 0 or 1 as it suffices for our purposes and keeps the DRE encoding function efficient.

Lemma 10 (DRE for private affine functions)**.** *Let $\mathcal{G}$ be a finite abelian group and let $f : \{\mathcal{G}\times\{0,1\}\}^n \to \mathcal{G}$ be the group summation function $f((a_1,x_1),\ldots,(a_n,x_n)) = \sum_{i=1}^n x_i a_i$.*[13] *Let $R = \{(\rho_1,\ldots,\rho_n) \in \mathcal{G}^n : \sum_{i=1}^n \rho_i = 0\}$. Then the function $\hat{f} : \{\mathcal{G}\times\{0,1\}\}^n \times R \to \mathcal{G}^n$ defined by $\hat{f}(((a_1,x_1),\ldots,(a_n,x_n)),(\rho_1,\ldots,\rho_n)) = (x_1a_1+\rho_1,\ldots,a_nx_n+\rho_n)$ is a decomposable encoding of f.*

Proof. The proof of this claim is identical to the proof of Lem. 9. This is a DRE with respect to the input blocks (a_i,x_i). It's decomposable because the i-th output depends only on (a_i,x_i) and ρ_i. □

5.5 DRE for the Scalar Multiplication

For the BN254 group, we define the function h in Def. 14, and construct a DRE for this function.

Definition 14. *Fix a nonzero point $\pi \in \mathbb{G}\setminus\{\mathcal{O}\}$. Let $h : \mathbb{F}_q \times \{0,1\}^\ell \to \mathbb{G}$ be the function defined by (with n, $\ell = 5n+1$ as in Sec. 5.2): the first argument is the shared secret input $r \in \mathbb{F}_q$, the second is the ℓ-tuple of bits $(\overline{u}_1(\pi),\ldots,\overline{u}_\ell(\pi))$ (the binary expansion of the point π), and*

$$h(r,\overline{u}_1(\pi),\ldots,\overline{u}_\ell(\pi)) \;\longmapsto\; r\pi.$$

The function h can be expressed as a matrix multiplication with a transition from field to group operations. Specifically, first, for each i, compute $\boldsymbol{N}(r_i)\times\overline{\boldsymbol{u}}^T(\pi)$ to obtain three field elements $(X,Y,Z) \in \mathbb{F}_p^3$ representing the Jacobian coordinates of the group element $r_i\pi$. Next, compute the weighted sum $\sum_{i=0}^{n-1} 2^i(r_i\pi)$ using elliptic curve addition,.

Theorem 4. *There exists a function $\hat{h}$ such that $\hat{h} : \mathbb{F}_q \times \{0,1\}^\ell \times (\mathrm{P}\times\Lambda\times S) \to \mathbb{F}_p^{3n\ell}$ is a DRE for h as defined in Def. 14, where the shared secret input $r \in \mathbb{F}_q$ is the first argument and $(\rho,\lambda,s) \in \mathrm{P}\times\Lambda\times S$ is the randomness used in the encoding.*

Proof. We start by describing the function $\hat{h}$ in more detail. Let $r \in \mathbb{F}_q$ be the shared secret input (first argument). The encoding has the form

$$\hat{h}\,(r,\overline{u}_1(\pi),\ldots,\overline{u}_\ell(\pi),\rho,\lambda,s) = (\hat{h}_1(r,\overline{u}_1(\pi),\rho,\lambda,s),\ldots,\hat{h}_\ell(r,\overline{u}_\ell(\pi),\rho,\lambda,s))$$

where $\hat{h}_k(r,\overline{u}_k(\pi),\rho,\lambda,s) = \{D_{i,j,k}\cdot\overline{u}_k(\pi)+s_{i,j,k}\}_{i,j}$, with r_i the bits of r, and $\boldsymbol{D}_i = \mathsf{diag}(\lambda_i^2,\lambda_i^3,\lambda_i)\times\boldsymbol{D}(r_i,\rho_i)$ for $i\in\{0,\ldots,n-1\}$, $j\in\{1,2,3\}$, $k\in\{1,\ldots,5n+1\}$, where $D_{i,j,k}$ is the $(j,k)^{th}$ entry in $\boldsymbol{D}_i$.

Here we set the randomness such that $\mathrm{P} = \{(\rho_0,\ldots,\rho_{n-1}) \in (\mathbb{G}\setminus\{\mathcal{O}\})^n : \sum_{i=0}^{n-1} 2^i\rho_i = \mathcal{O}\}$, $\Lambda = (\mathbb{F}_p^*)^n$ (with $\lambda = (\lambda_0,\ldots,\lambda_{n-1}) \in \Lambda$), and $S = \{(s_{i,j,k})_{i=0,j=1,k=1}^{n-1,3,\ell} \in \mathbb{F}_p^{3n\ell} : \forall i,j\ \sum_{k=1}^{\ell} s_{i,j,k} = 0\}$.

[13] Where $+$ is the group operation and $ax = 0$ if $x = 0$ and a otherwise.

Decoding Function. The decoding function $\mathsf{Dec} : \mathbb{F}_p^{3n\ell} \to \mathbb{G}$ on input $\gamma_1 \dots \gamma_\ell$ where $\gamma_k = \{t_{i,j,k}\}_{i,j}$ is either $\hat{h}_k(r, 0, \rho, \lambda, s)$ or $\hat{h}_k(r, 1, \rho, \lambda, s)$ is defined as follows:

1. For each $i \in \{0, \dots, n-1\}$ and $j \in \{1, 2, 3\}$, compute the j-th coordinate of the randomized Jacobian representation:
$$\hat{t}_{i,j} = \sum_{k=1}^{\ell} t_{i,j,k}$$
This correctly recovers the j-th coordinate since $\sum_{k=1}^{\ell} s_{i,j,k} = 0$ by the constraint on S.
2. For each $i \in \{0, \dots, n-1\}$, form the point $Q_i = (\hat{t}_{i,1}, \hat{t}_{i,2}, \hat{t}_{i,3}) \in \mathbb{F}_p^3$, which represents the group element $\mathsf{diag}(\lambda_i^2, \lambda_i^3, \lambda_i) \times (r_i\pi + \rho_i)$.
3. Apply the weighted group sum to recover $r\pi$:
$$\sum_{i=0}^{n-1} 2^i Q_i = r\pi$$
This correctly computes $r\pi$ since $\sum_{i=0}^{n-1} 2^i \rho_i = 0$ by the constraint on P.

Next we derive this DRE for h using Lem. 9 and Lem. 10. The scalar $r = r$ is the shared secret input; $\overline{u}(\pi)$ is the decomposable input.

Step 1: DRE for weighted group sum using Lem. 9. Let $r = r = \sum_{i=0}^{n-1} r_i 2^i \in \mathbb{F}_q$ be the binary decomposition of r. We want to compute $r\pi = \sum_{i=0}^{n-1} 2^i (r_i\pi)$. Define the function $f : \mathbb{G}^n \to \mathbb{G}$ by:

$$f(r_0\pi, \dots, r_{n-1}\pi) = \sum_{i=0}^{n-1} 2^i (r_i\pi) = r\pi$$

with weights $a_i = 2^i$ for $i \in \{0, \dots, n-1\}$. By Lem. 9 and Remark 3, there exists a DRE $\hat{f} : \mathbb{G}^n \times R_1 \to \mathbb{G}^n$ where $R_1 = \{(\rho_0, \dots, \rho_{n-1}) \in \mathbb{G}^n, (\lambda_0, \dots, \lambda_{n-1}) \in \mathbb{F}_p^{*n} : \sum_{i=0}^{n-1} 2^i \rho_i = 0\}$, where

$$\hat{f}((r_0\pi, \dots, r_{n-1}\pi), R_1 = ((\rho_0, \dots, \rho_{n-1}), (\lambda_0, \dots, \lambda_{n-1}))) = (\hat{f}_0(r_0\pi, R_1), \dots, \hat{f}_{n-1}(r_{n-1}\pi, R_1)$$

and for each i, $\hat{f}_i(r_i\pi, R_1) = \mathsf{diag}(\lambda_i^2, \lambda_i^3, \lambda_i) \times (r_i\pi + \rho_i)^T$. The decoding function is $\mathsf{Dec}_1 : \mathbb{G}^n \to \mathbb{G}$ defined by $\mathsf{Dec}_1((y_0, \dots, y_{n-1})) = \sum_{i=0}^{n-1} 2^i y_i$, which correctly computes $r\pi$ since $\sum_{i=0}^{n-1} 2^i \rho_i = 0$.

Step 2: DRE for $\mathsf{diag}(\lambda_i^2, \lambda_i^3, \lambda_i) \times (r_i\pi + \rho_i)^T$ using Lem. 10. For each $i \in \{0, \dots, n-1\}$, we need to encode the group element $r_i\pi + \rho_i$. Since $r_i \in \{0, 1\}$, we have:

$$r_i\pi + \rho_i = \begin{cases} \pi + \rho_i & \text{if } r_i = 1 \\ \rho_i & \text{if } r_i = 0 \end{cases}$$

For each i, let E_i be the event $x(\rho_i) = x(\pi)$. If $\neg E_i$ then we can apply Lem. 5 to express the randomized Jacobian coordinates of $\mathsf{diag}(\lambda_i^2, \lambda_i^3, \lambda_i) \times (r_i\pi + \rho_i)^T$ (setting with $\phi = \rho_i$ and $\delta = r_i$ in Lem. 5) as a linear function of $\boldsymbol{u}(\pi)$. Specifically:

$$\mathsf{diag}(\lambda_i^2, \lambda_i^3, \lambda_i) \times \boldsymbol{D}(r_i, \rho_i) \times \overline{\boldsymbol{u}}^T(\pi)$$

which can be written as $\boldsymbol{D}_i \times \overline{\boldsymbol{u}}^T(\pi)$ where $\boldsymbol{D}_i \stackrel{def}{=} \mathsf{diag}(\lambda_i^2, \lambda_i^3, \lambda_i) \times \boldsymbol{D}(r_i, \rho_i) \in \mathbb{F}_p^{3\times(1+5n)}$.

For each coordinate $j \in \{1,2,3\}$ (corresponding to X, Y, Z), let $\boldsymbol{D}_{i,j}$ be the j-th row of $\boldsymbol{D}_i$. Then, the j-th coordinate of the randomized Jacobian coordinates is:

$$\boldsymbol{D}_{i,j} \times \overline{\boldsymbol{u}}^T(\pi) = \sum_{k=1}^{1+5n} D_{i,j,k} \cdot \overline{u}_k(\pi)$$

where $D_{i,j,k}$ is the k-th entry of the row vector $\boldsymbol{D}_{i,j}$, and $\overline{u}_k(\pi)$ is the k-th component of $\overline{\boldsymbol{u}}(\pi)$. Now, for each $i \in \{0,\ldots,n-1\}$ and $j \in \{1,2,3\}$, we apply Lem. 10 to encode the function:

$$g_{i,j} : (\mathbb{F}_p \times \{0,1\})^{1+5n} \to \mathbb{F}_p$$

defined by $g_{i,j}((D_{i,j,1}, \overline{u}_1(\pi)), \ldots, (D_{i,j,1+5n}, \overline{u}_{1+5n}(\pi))) = \sum_{k=1}^{1+5n} D_{i,j,k} \cdot \overline{u}_k(\pi)$. Note that $D_{i,j,k}$ are known values that depend on r_i and the randomness ρ_i and λ_i from Step 1. The function $g_{i,j}$ takes as input pairs $(D_{i,j,k}, \overline{u}_k(\pi))$ where $D_{i,j,k} \in \mathbb{F}_p$ is a field element (the coefficient) and $\overline{u}_k(\pi) \in \{0,1\}$ is a bit. By Lem. 10, there exists a DRE $\hat{g}_{i,j} : ((\mathbb{F}_p \times \{0,1\}))^{1+5n} \times S_{i,j} \to \mathbb{F}_p^{1+5n}$ where $S_{i,j} = \{(s_{i,j,1}, \ldots, s_{i,j,1+5n}) \in \mathbb{F}_p^{1+5n} : \sum_{k=1}^{1+5n} s_{i,j,k} = 0\}$, defined by:

$$\hat{g}_{i,j}((D_{i,j,1}, \overline{u}_1(\pi)), \ldots, (D_{i,j,1+5n}, \overline{u}_{1+5n}(\pi))), S_{i,j} = (s_{i,j,1}, \ldots, s_{i,j,1+5n})) = \\ (D_{i,j,1} \cdot \overline{u}_1(\pi) + s_{i,j,1}, \ldots, D_{i,j,1+5n} \cdot \overline{u}_{1+5n}(\pi) + s_{i,j,1+5n})$$

The decoding function is $\mathsf{Dec}_{i,j} : \mathbb{F}_p^{1+5n} \to \mathbb{F}_p$ defined by $\mathsf{Dec}_{i,j}((y_1, \ldots, y_{1+5n})) = \sum_{k=1}^{1+5n} y_k$, which, conditioned on the fact that Lem. 5 can be applied (i.e. $\neg E_i$), correctly computes the correct j-th coordinate since $\sum_{k=1}^{1+5n} s_{i,j,k} = 0$.

Composition and final DRE. By Lem. 8 and Lem. 7, we can compose the DREs from Step 1 and Step 2 to obtain the final DRE $\hat{h}$ for h.

Correctness and Security. Our analysis holds conditioned on $\neg E_i$ for each $i \in \{0,\ldots,n-1\}$. Each ρ_i is marginally uniform over $\mathbb{G} \setminus \{\mathcal{O}\}$: for any two nonzero values, scaling every coordinate of a tuple in P by the unique nonzero scalar that sends the first value to the second gives a bijection between the corresponding fibers while preserving the weighted constraint. Since π is nonzero and G has odd prime order, $x(\rho_i) = x(\pi)$ exactly when $\rho_i \in \{\pi, -\pi\}$, so $\Pr[E_i] = 2/(|\mathbb{G}| - 1)$. A union bound, which does not require the events to be independent, therefore gives correctness and privacy errors $\delta, \epsilon \leqslant 2n/(|\mathbb{G}| - 1)$. □

5.6 Completing the Garbled Circuit

We wish to design an efficient Garbled Circuit that on public input an (on-chain) Lamport signature on $\pi = (x(\pi), y(\pi))$ outputs $r\pi$. Here $r \in \mathbb{F}_q^*$ is a private input known at the time garbling that is considered hard-coded in the circuit.

Construction 2. *The garbled circuit* GC *for the function* $f_r : \mathbb{G} \to \mathbb{G}$, $f_r(\pi) = r\pi$ *will be as follows:*

$\mathsf{Garble}(r) \to \mathsf{ct}_{GC}, \mathsf{ek}$*: Outputs a Lamport secret key*

$$\mathsf{ek} = \begin{pmatrix} L^0_{x,0} & \ldots & L^0_{x,n-1} & L^0_{y,0} & \ldots & L^0_{y,n-1} \\ L^1_{x,0} & \ldots & L^1_{x,n-1} & L^1_{y,0} & \ldots & L^1_{y,n-1} \end{pmatrix}$$

and the garbled circuit $\mathsf{ct}_{\mathsf{GC}}$ *computed in three steps:*

1. ***Elliptic curve validation:*** *The garbled circuit first verifies that the input point $\pi = (x(\pi), y(\pi))$ lies on the curve E, i.e., it checks that $y(\pi)^2 \equiv x(\pi)^3 + 3 \pmod{p}$.*

2. ***Binary decomposition:*** *If the validation passes, the circuit computes the binary decomposition $\overline{\boldsymbol{u}}(\pi) \in \{0,1\}^{1+5n}$ of $\boldsymbol{u}(\pi) = (1, x(\pi), y(\pi), x(\pi)^2, y(\pi)^2, x(\pi)y(\pi))$. If not, the circuit outputs the binary decomposition of $\boldsymbol{u}(g) = (1, x(g), y(g), x(g)^2, y(g)^2, x(g)y(g))$ where $g \in \mathbb{G}_1$ is some fixed point in the curve. These two steps are implemented using a privacy-free Boolean garbled circuit. Evaluating this circuit outputs the labels $L_{\overline{u},k}^{\overline{u}_k(\pi)}$ for $k \in [\ell]$.*

3. ***DRE encoding:*** *Let $r \in \mathbb{F}_q$ be the shared secret input (the scalar, known at garbling time). For each $k \in \{1, \ldots, \ell\}$ where $\ell = 1 + 5n$, the garbled circuit comes hardwired with the DRE encoding from Thm. 4. Specifically:*

   - *For each bit position $k \in [\ell]$ where $\ell = 1+5n$, the garbler precomputes $\hat{h}_k(r, 0, \omega)$ and $\hat{h}_k(r, 1, \omega)$.*
   - *The garbled circuit contains an encryption of $\hat{h}_k(r, 0, \omega)$ under the 0-label $L_{\overline{u},k}^0$ for $\overline{u}_k(\pi)$.*
   - *The garbled circuit contains an encryption of $\hat{h}_k(r, 1, \omega)$ under the 1-label $L_{\overline{u},k}^1$ for $\overline{u}_k(\pi)$.*

$\mathsf{Encode}(\mathsf{ek}, \pi) \to L$*: Outputs the Lamport signature of π*

$$L = \begin{pmatrix} L_{x,0}^{x_0} & \ldots & L_{x,n-1}^{x_{n-1}} & L_{y,0}^{y_0} & \ldots & L_{y,n-1}^{y_{n-1}} \end{pmatrix}$$

$\mathsf{Eval}(\mathsf{ct}_{GC}, L_\pi) \to r\pi$*: The evaluator, given the input labels L_π, evaluates the Boolean circuits to obtain labels for $\overline{u}(\pi)$, and decrypts the appropriate encryptions to obtain $\hat{h}_k(r, \overline{u}_k(\pi), \omega)$ for each k (with $r \in \mathbb{F}_q$ the shared secret input). The evaluator then decodes the DRE as follows (Thm. 4):*

- *We recast $\hat{h}_k(r, \overline{u}_k(\pi), \omega)$ as $\{t_{i,j,k}\}_{i\in[n],j\in\{1,2,3\}}$. Then the Prover aggregates shares to cancel the $s_{i,j,k}$-terms:*

$$\widehat{t}_{i,j} = \sum_{k=1}^{\ell} t_{i,j,k}.$$

- *Interpreting $(\widehat{t}_{i,1}, \widehat{t}_{i,2}, \widehat{t}_{i,3})$ as the Jacobian coordinates of a masked point Q_i, the final output is obtained by the usual bit-weighted recombination*

$$\text{out} = \sum_i 2^i Q_i.$$

Since $Q_i = r_i\pi + \rho_i$ and $\sum_i 2^i\rho_i = 0$, then $\text{out} = r\pi$.

Intuitively, the heavy work (group operations and most algebra) is pushed to the Prover and to the Verifier's offline preprocessing, while the online garbled circuit mainly performs validation, low-level arithmetic for feature extraction, and symmetric-key decryptions for table selection to output the precomputed DRE components.

Theorem 5. *Constr. 2 is an adaptively secure garbling scheme (Def. 2) for the function $f_r : \mathbb{G} \to \mathbb{G}$, $f_r(\pi) = r\pi$ in the random oracle model.*

The proof follows directly from Thm. 4 and the security of Yao's garbling scheme [LP09], specifically the authenticity property [BHR12b].[14] For adaptive security we rely on standard random oracle techniques [BHR12a] (namely equivocal encryption).

[14] For the optimization we propose, i.e. instead of letting the circuit compute the $\hat{h}_k(r, \overline{u}_k, \omega)$'s we precompute them and encrypt them under the corresponding labels, the simulator may just encrypt the simulated $\hat{h}_k$'s for the actual $\overline{u}_k$'s and provide encryption of 0 for the other bit.

Remark 4 (Optimizations for the Randomized Encoding)**.** *In practice we integrate the following optimizations:*

- *Several entries in the matrix $\boldsymbol{D}$ are* 0*. In our implementation the inner products only need to grow with the number of nonzero entries. This reduces garbled circuit size by a factor of* 1.87*.*
- *Rather than giving two ciphertexts encrypting $\hat{h}_i(r, 0, \rho, \lambda, s)$ and $\hat{h}_i(r, 1, \rho, \lambda, s)$ we just encrypt one of them and let the appropriate shifts be the output of a Pseudorandom Function (e.g. a random oracle) on the corresponding label, such that the latter can be computed locally. This reduces garbled circuit size by a factor of* 2*.*

Efficiency. We provide a theoretical estimation of our garbled circuit size. Since it is hard to theoretically estimate the size of the circuit for (1) and (2) of the construction (Elliptic Curve validation and Binary Decomposition) we postpone this for the experimental evaluation in Sec. 9. However, the main cost stems from (3), the DRE Encoding. Taking into account our optimizations, the latter consists of:

- For each $i = 0, \ldots, n-1$
  - $3n + 1$ $\log |\mathbb{F}|$-sized ciphertexts for the X coordinate.
  - $4n + 1$ $\log |\mathbb{F}|$-sized ciphertexts for the Y coordinate.
  - $n + 1$ $\log |\mathbb{F}|$-sized ciphertexts for the Z coordinate.

Here $\mathbb{F}$ is the field generated by the $E/\mathbb{F}_p$ group, therefore $|\mathbb{F}| = 254$ and $n = 254$. This gives us:

$$254 \cdot (8 \cdot 254 + 3) \cdot 254 \text{ bits} = 15.65 \text{ MiB}$$

This is validated empirically in Sec. 9.

6 BABE Protocol

6.1 Honest Setup Protocol

A BitVM-core protocol (Sec. 3) has two phases: a setup phase and a proving phase. The setup phase consists of off-chain interaction between the Prover and the Verifier (Fig. 3). The proving phase consists of on-chain interaction between the Prover and the Verifier and ends with the Prover posting a success transaction if he has a valid proof. To explain the core aspects of the protocol, we first describe the protocol assuming that the setup is run honestly by both Prover and Verifier (Algs. 1 and 2), but either party may be malicious in the proving phase. In Sec. 6.2, we discuss how the Prover and Verifier can verify that the setup was run correctly and abort if it was not.

6.1.1 Setup Phase

The setup phase protocol is described in Alg. 1. The statement $\boldsymbol{x}$ to be proven is given as input to the setup algorithms. For example, in the lending application from Sec. 1, when the borrower is the Prover, the statement is that the borrower repaid his loan on Ethereum. The Prover initiates the setup by sending his public key pk_P to the Verifier. The Verifier samples a secret msg and creates the witness encryption ciphertext ($\mathsf{ct}_{\mathsf{setup}}$ in Constr. 1) that doesn't depend on the proof. He also creates the garbled circuit (Constr. 2) to compute $\mathsf{ct}_{\mathsf{prove}}$, the remaining part of the ciphertext, generating the encoding key ek and the garbled circuit ciphertext $\mathsf{ct}_{\mathsf{GC}}$. The Verifier hashes the secret msg and the garbled circuit's encoding key

Algorithm 1 `Setup` algorithms (honest Prover and Verifier)

function $\texttt{Gen}(\mathcal{R})$
  return $\mathsf{Groth16}.\texttt{Gen}(\mathcal{R})$
end function

procedure $P_{\texttt{Setup}}(\mathsf{crs}, \boldsymbol{x})$ ▷ Run by Prover
  $(\mathsf{sk}_P, \mathsf{pk}_P) \leftarrow \mathsf{Sig}_{BTC}.\texttt{Gen}(1^\lambda)$ ▷ Sample signing key
  send (pk_P) to Verifier

  Upon receiving $(\mathsf{pk}_V, h_{\mathsf{msg}}, \mathsf{epk}, \mathsf{ct}_{\mathsf{setup}}, \mathsf{ct}_{\mathsf{GC}})$ from Verifier:
  $(\mathsf{lsk}_P, \mathsf{lpk}_P) \leftarrow \mathsf{LampSig}.\texttt{Gen}(1^\lambda)$ ▷ Sample Lamport key
  $(\mathcal{T}, \mathcal{S}) \leftarrow \mathsf{CreateTxSet}(\mathsf{pk}_P, \mathsf{pk}_V, \mathsf{lpk}_P, h_{\mathsf{msg}}, \mathsf{epk})$ ▷ See Alg. 3
  $\mathsf{presigs}_P \leftarrow \mathsf{SignTxs}_P(\mathsf{sk}_P, \mathcal{T})$ ▷ See Alg. 3
  send $(\mathsf{lpk}_P, \mathsf{presigs}_P)$ to Verifier

  Upon receiving $(\mathsf{presigs}_V)$ from Verifier:
  $\mathsf{st}_P \leftarrow (\mathsf{crs}, \boldsymbol{x}, \mathcal{T}, \mathsf{sk}_P, \mathsf{lsk}_P, \mathsf{presigs}_V, \mathsf{ct}_{\mathsf{setup}}, \mathsf{ct}_{\mathsf{GC}})$
  return $(\mathcal{S}, \mathsf{st}_P)$
end procedure

procedure $V_{\texttt{Setup}}(\mathsf{crs}, \boldsymbol{x})$ ▷ Run by Verifier
  Upon receiving (pk_P) from Prover:
  $(\mathsf{sk}_V, \mathsf{pk}_V) \leftarrow \mathsf{Sig}_{BTC}.\texttt{Gen}(1^\lambda)$ ▷ Sample signing key
  $\mathsf{msg} \leftarrow\!\!\$\ \mathcal{M}$, $r \leftarrow\!\!\$\ \mathbb{F}_q^*$ ▷ Sample secrets
  $\mathsf{ct}_{\mathsf{setup}} \leftarrow \mathsf{WE}.\mathsf{Enc}_{\mathsf{setup}}(\mathsf{crs}, \boldsymbol{x}, \mathsf{msg}, r)$ ▷ WE ciphertext (Constr. 1)
  $\mathsf{ct}_{\mathsf{GC}}, \mathsf{ek} \leftarrow \mathsf{Garble}(r)$ ▷ Garbled circuit ciphertext and encoding key (Sec. 5.6)
  $h_{\mathsf{msg}} \leftarrow \mathsf{Hash}_{BTC}(\mathsf{msg})$ ▷ Hash message for hashlock
  for $j \in \{1, \ldots, 2n\}, b \in \{0, 1\}$ **do** ▷ $n = \lceil \log_2 p \rceil$ is the number of bits for each coordinate of π_1
    $\mathsf{epk}_j^b \leftarrow \mathsf{Hash}_{BTC}(\mathsf{ek}_j^b)$ ▷ Hash input labels for hashlock
  end for
  send $(\mathsf{pk}_V, h_{\mathsf{msg}}, \mathsf{epk}, \mathsf{ct}_{\mathsf{setup}}, \mathsf{ct}_{\mathsf{GC}})$ to Prover

  Upon receiving $(\mathsf{lpk}_P, \mathsf{presigs}_P)$ from Prover:
  $(\mathcal{T}, \mathcal{S}) \leftarrow \mathsf{CreateTxSet}(\mathsf{pk}_P, \mathsf{pk}_V, \mathsf{lpk}_P, h_{\mathsf{msg}}, \mathsf{epk})$
  $\mathsf{presigs}_V \leftarrow \mathsf{SignTxs}_V(\mathsf{sk}_V, \mathcal{T})$
  $\mathsf{st}_V \leftarrow (\mathsf{crs}, \boldsymbol{x}, \mathcal{T}, \mathsf{sk}_V, \mathsf{ek}, \mathsf{presigs}_P)$
  send $(\mathsf{presigs}_V)$ to Prover
  return $(\mathcal{S}, \mathsf{st}_V)$
end procedure

ek to be used in the hashlock scripts. Then, both parties create a set of transaction skeletons[15] $\mathcal{T}$ shown in Fig. 4. The Prover pre-signs the transaction skeletons $\overline{\mathsf{tx}}_{\mathsf{ChallengeAssert}}$ and $\overline{\mathsf{tx}}_{\mathsf{Failure}}$ (ones that the Verifier may post during the proving phase). Similarly, the Verifier pre-signs the transaction skeleton $\overline{\mathsf{tx}}_{\mathsf{Success}}$ (that the Prover may post during the proving phase). At the end, the Prover stores the state st_P consisting of crs, $\boldsymbol{x}$, $\mathcal{T}$, his secret keys, the pre-signatures sent by the Verifier, the witness encryption ciphertext $\mathsf{ct}_{\mathsf{setup}}$, and the garbled circuit ciphertext $\mathsf{ct}_{\mathsf{GC}}$. The Verifier's state st_V consists of crs, $\boldsymbol{x}$, $\mathcal{T}$, his secret keys, and the pre-signatures sent by the Prover.

6.1.2 Proving Phase

The proving phase protocol is described in Alg. 2. In the proving phase, the Prover generates a Groth16 proof (π_1, π_2, π_3) using the witness for the statement $\boldsymbol{x}$ that was agreed upon during the setup phase. The proving phase involves the following transactions (shown in Fig. 4) posted on Bitcoin:

[15] Recall: transaction skeleton is transaction without transaction witness.

Algorithm 2 `Prove` algorithms

procedure $P_{\texttt{Prove}}(\mathsf{st}_P, \boldsymbol{w})$ ▷ Run by Prover
  Parse $\mathsf{st}_P = (\mathsf{crs}, \boldsymbol{x}, \mathcal{T}, \mathsf{sk}_P, \mathsf{lsk}_P, \mathsf{presigs}_V, \mathsf{ct}_{\mathsf{setup}}, \mathsf{ct}_{\mathsf{GC}})$ and $\mathsf{presigs}_V = (\sigma^V_{\mathsf{Success}})$
  $(\pi_1, \pi_2, \pi_3) \leftarrow \mathsf{Groth16.Prove}(\mathsf{crs}, \boldsymbol{x}, \boldsymbol{w})$
  $w_{\mathsf{Assert}} \leftarrow \mathsf{LampSig.Sign}(\mathsf{lsk}_P, \pi_1)$ ▷ Compute Lamport signature
  Post $\mathsf{tx}_{\mathsf{Assert}}$ with transaction witness w_{Assert}: call $\mathcal{F}_{\mathsf{BTC}}.\mathsf{WRITE}(\mathsf{tx}_{\mathsf{Assert}})$

  Upon seeing $\mathsf{tx}_{\mathsf{Assert}}$ and Δ_2 new blocks after $\mathsf{tx}_{\mathsf{Assert}}$ in $\mathcal{L}_P$:
  $w_{\mathsf{Success}} \leftarrow (\sigma^P_{\mathsf{Success}}, \sigma^V_{\mathsf{Success}})$ where $\sigma^P_{\mathsf{Success}} \leftarrow \mathsf{Sig}_{BTC}.\mathsf{Sign}(\mathsf{sk}_P, \overline{\mathsf{tx}}_{\mathsf{Success}})$
  Post $\mathsf{tx}_{\mathsf{Success}}$ with transaction witness w_{Success}: call $\mathcal{F}_{\mathsf{BTC}}.\mathsf{WRITE}(\mathsf{tx}_{\mathsf{Success}})$
  return 1

  Upon seeing $\mathsf{tx}_{\mathsf{ChallengeAssert}}$ in $\mathcal{L}_P$:
  Extract input labels L from the transaction witness of $\mathsf{tx}_{\mathsf{ChallengeAssert}}$
  $\mathsf{ct}_{\mathsf{prove}} \leftarrow \mathsf{EvalGC}(\mathsf{ct}_{\mathsf{GC}}, L)$ ▷ Evaluate garbled circuit (Sec. 5.6)
  $\mathsf{msg} \leftarrow \mathsf{WE.Dec}(\mathsf{crs}, (\boldsymbol{x}, \pi_1), (\mathsf{ct}_{\mathsf{prove}}, \mathsf{ct}_{\mathsf{setup}}), \boldsymbol{w})$ ▷ Must succeed for an honest Prover by u-robustness (Def. 8)
  $w_{\mathsf{WronglyChallenged}} \leftarrow (\sigma^P_{\mathsf{WronglyChallenged}}, \mathsf{msg})$ where $\sigma^P_{\mathsf{WronglyChallenged}} \leftarrow \mathsf{Sig}_{BTC}.\mathsf{Sign}(\mathsf{sk}_P, \overline{\mathsf{tx}}_{\mathsf{WronglyChallenged}})$
  Post $\mathsf{tx}_{\mathsf{WronglyChallenged}}$ with transaction witness $w_{\mathsf{WronglyChallenged}}$: call $\mathcal{F}_{\mathsf{BTC}}.\mathsf{WRITE}(\mathsf{tx}_{\mathsf{WronglyChallenged}})$
end procedure

procedure $V_{\texttt{Prove}}(\mathsf{st}_V)$ ▷ Run by Verifier
  Upon seeing $\mathsf{tx}_{\mathsf{Assert}}$ in $\mathcal{L}_V$:
  Parse $\mathsf{st}_V = (\mathsf{crs}, \boldsymbol{x}, \mathcal{T}, \mathsf{sk}_V, \mathsf{ek}, \mathsf{presigs}_P)$ and $\mathsf{presigs}_P = (\sigma^P_{\mathsf{ChallengeAssert}}, \sigma^P_{\mathsf{Failure}})$
  Extract π_1 and Lamport signature μ from the transaction witness of $\mathsf{tx}_{\mathsf{Assert}}$
  $L \leftarrow \mathsf{Encode}(\mathsf{ek}, \pi_1)$ ▷ Compute input labels (Sec. 5.6)
  $w_{\mathsf{ChallengeAssert}} \leftarrow (\sigma^P_{\mathsf{ChallengeAssert}}, \sigma^V_{\mathsf{ChallengeAssert}}, \mu, L)$ where $\sigma^V_{\mathsf{ChallengeAssert}} \leftarrow \mathsf{Sig}_{BTC}.\mathsf{Sign}(\mathsf{sk}_V, \overline{\mathsf{tx}}_{\mathsf{ChallengeAssert}})$
  Post $\mathsf{tx}_{\mathsf{ChallengeAssert}}$ with transaction witness $w_{\mathsf{ChallengeAssert}}$: call $\mathcal{F}_{\mathsf{BTC}}.\mathsf{WRITE}(\mathsf{tx}_{\mathsf{ChallengeAssert}})$

  Upon seeing $\mathsf{tx}_{\mathsf{ChallengeAssert}}$ and Δ_1 new blocks after $\mathsf{tx}_{\mathsf{ChallengeAssert}}$ in $\mathcal{L}_V$, if $\mathsf{tx}_{\mathsf{WronglyChallenged}} \notin \mathcal{L}_V$:
  $w_{\mathsf{Failure}} \leftarrow \big((\sigma^P_{\mathsf{Failure}}, \sigma^V_{\mathsf{Failure}}), (\sigma^V_{\mathsf{Failure}})\big)$ where $\sigma^V_{\mathsf{Failure}} \leftarrow \mathsf{Sig}_{BTC}.\mathsf{Sign}(\mathsf{sk}_V, \overline{\mathsf{tx}}_{\mathsf{Failure}})$
  Post $\mathsf{tx}_{\mathsf{Failure}}$ with transaction witness w_{Failure}: call $\mathcal{F}_{\mathsf{BTC}}.\mathsf{WRITE}(\mathsf{tx}_{\mathsf{Failure}})$
end procedure

1. Assert: used by the Prover to post the proof element π_1.
2. ChallengeAssert: used by the Verifier to post the input labels for the proof π_1. The Bitcoin script verifies that the input labels are for the same proof π_1 that the Prover posted.
3. If the Prover's proof is valid:
   (a) WronglyChallenged: The Prover evaluates the garbled circuit to compute $\mathsf{ct}_{\mathsf{prove}} = r\pi_1$, decrypts the secret msg (Constr. 1), then posts this transaction. The Bitcoin script requires the Prover to provide the decrypted secret msg.
   (b) Success: The Prover can post this transaction if and only if his proof is valid. This transaction can only be posted Δ_2 blocks after he posted Assert. This timelock gives the Verifier enough time to "stop" a malicious Prover.
4. If the Prover's proof is invalid:
   (a) Failure: This prevents the Prover from ever posting the Success transaction. It can only be posted Δ_1 blocks after the Verifier posted ChallengeAssert, which gives the Prover enough time to post WronglyChallenged if he decrypted the secret msg.

Algorithm 3 Helper functions for Algs. 1 and 2

function CREATETXSET($\mathsf{pk}_P, \mathsf{pk}_V, \mathsf{lpk}_P, h_{\mathsf{msg}}, \mathsf{epk}$)
  Construct transaction skeletons $\overline{\mathsf{tx}}_{\mathsf{Assert}}$, $\overline{\mathsf{tx}}_{\mathsf{ChallengeAssert}}$, $\overline{\mathsf{tx}}_{\mathsf{Failure}}$, $\overline{\mathsf{tx}}_{\mathsf{WronglyChallenged}}$, $\overline{\mathsf{tx}}_{\mathsf{Success}}$ as in App. A.1
  $\mathcal{T} := \{\overline{\mathsf{tx}}_{\mathsf{Assert}}, \overline{\mathsf{tx}}_{\mathsf{ChallengeAssert}}, \overline{\mathsf{tx}}_{\mathsf{Failure}}, \overline{\mathsf{tx}}_{\mathsf{WronglyChallenged}}, \overline{\mathsf{tx}}_{\mathsf{Success}}\}$
  $\mathcal{S} := \{\overline{\mathsf{tx}}_{\mathsf{Success}}\}$
  return $(\mathcal{T}, \mathcal{S})$
end function

function SIGNTXS$_P$($\mathsf{sk}_P, \mathcal{T}$)
  $\sigma^P_{\mathsf{ChallengeAssert}} := \mathsf{Sig}_{BTC}.\mathsf{Sign}(\mathsf{sk}_P, \overline{\mathsf{tx}}_{\mathsf{ChallengeAssert}})$
  $\sigma^P_{\mathsf{Failure}} := \mathsf{Sig}_{BTC}.\mathsf{Sign}(\mathsf{sk}_P, \overline{\mathsf{tx}}_{\mathsf{Failure}})$
  return $(\sigma^P_{\mathsf{ChallengeAssert}}, \sigma^P_{\mathsf{Failure}})$
end function

function SIGNTXS$_V$($\mathsf{sk}_V, \mathcal{T}$)
  $\sigma^V_{\mathsf{Success}} := \mathsf{Sig}_{BTC}.\mathsf{Sign}(\mathsf{sk}_V, \overline{\mathsf{tx}}_{\mathsf{Success}})$
  return $(\sigma^V_{\mathsf{Success}})$
end function

function VERIFYSIGS$_P$($\mathsf{pk}_V, \mathcal{T}, \mathsf{presigs}_V$)
  Parse $\mathsf{presigs}_V = (\sigma^V_{\mathsf{Success}})$
  Verify $\mathsf{Sig}_{BTC}.\mathtt{Verify}(\mathsf{pk}_V, \overline{\mathsf{tx}}_{\mathsf{Success}}, \sigma^V_{\mathsf{Success}})$
end function

function VERIFYSIGS$_V$($\mathsf{pk}_P, \mathcal{T}, \mathsf{presigs}_P$)
  Parse $\mathsf{presigs}_P = (\sigma^P_{\mathsf{ChallengeAssert}}, \sigma^P_{\mathsf{Failure}})$
  Verify $\mathsf{Sig}_{BTC}.\mathtt{Verify}(\mathsf{pk}_P, \overline{\mathsf{tx}}_{\mathsf{ChallengeAssert}}, \sigma^P_{\mathsf{ChallengeAssert}})$
  Verify $\mathsf{Sig}_{BTC}.\mathtt{Verify}(\mathsf{pk}_P, \overline{\mathsf{tx}}_{\mathsf{Failure}}, \sigma^P_{\mathsf{Failure}})$
end function

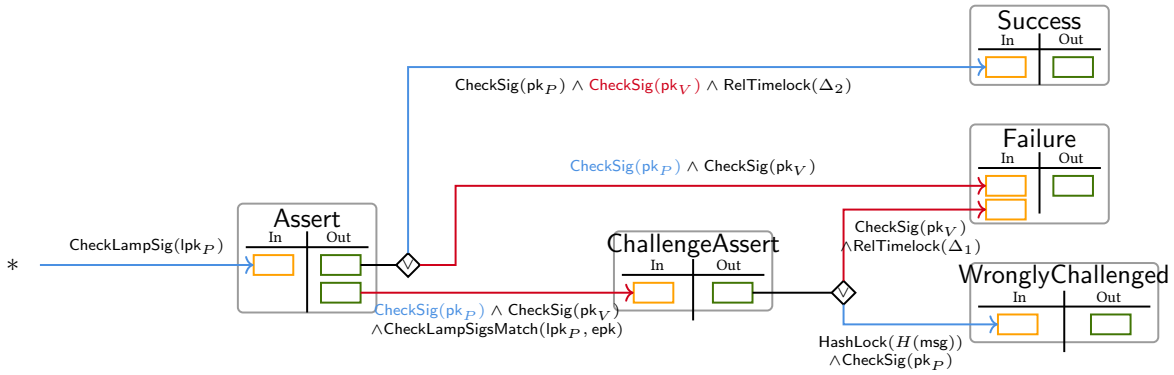

Figure 4: Illustration of the Bitcoin transaction graph for the BITVM-CORE protocol. An asterisk (∗) on an incoming arrow indicates that the source is an external UTXO (not part of the protocol). Gray boxes represent transactions. A transaction's inputs and outputs are represented by orange and green boxes, respectively, inside the transaction. Arrows connect one transaction's output used as an input by another transaction. The ◇ shape indicates that the transaction output can be used as input by any one out of multiple possible transactions. Locking scripts written on an arrow entering a transaction must be satisfied by the transaction witness. Arrows entering a transaction posted by the Prover are blue and arrows entering a transaction posted by the Verifier are red. A red $\mathsf{CheckSig}(\mathsf{pk}_V)$ indicates that the transaction taking that input is pre-signed by the Verifier during setup. Similarly, a blue $\mathsf{CheckSig}(\mathsf{pk}_P)$ indicates that the transaction is pre-signed by the Prover during setup.

6.1.3 Transaction Graph

The transactions posted during the proving phase form a graph where the outputs of one transaction are the inputs of another (this graph is shown in Fig. 4). The Assert transaction takes any external UTXO as input; this is indicated by an asterisk in Fig. 4. A detailed specification of the transactions is given in App. A.1. This section describes the locking scripts used and how they ensure the protocol's security. The transactions in Fig. 4 themselves do not move any funds; they only encode the logic of the SNARK verification protocol. Applications such as the trustless lending protocol in Sec. 1 can use a separate Deposit transaction to deposit Bitcoin and a Withdraw transaction that consumes both the Deposit output and the output of the Success transaction to move funds to the Prover only if his proof is deemed valid by the BABE protocol.

Pre-signed transactions. During the setup phase, the Prover and the Verifier pre-sign the set of allowed transactions that can be posted by the other party during the proving phase. For example, the Verifier pre-signs the Success transaction and gives this signature to the Prover during setup. Since the signatures are on the transaction skeleton, they can be signed during setup without knowing the transaction witnesses. This ensures the following:

- *Unilateral posting*: The Prover can post the Success transaction during the proving phase without having to depend on the Verifier to sign it.
- *Output binding*: The Prover cannot post any other transaction spending the same connector output because he does not have the Verifier's signature for such a transaction.
- *Input binding*: The pre-signed transaction Success commits to the hash of a specific Assert transaction skeleton because it uses a connector output of Assert as an input. This ensures that the Prover cannot post the Success transaction unless he first posts the correct Assert transaction.

In general, these properties ensure that only transactions from the set $\mathcal{T}$ created during the setup phase can be posted during the proving phase.

Timelocks. The Success transaction takes as its input the 0-th output of Assert, which carries the condition $\mathsf{RelTimelock}(\Delta_2)$ in its locking script. This ensures that the Prover cannot post the Success transaction until Assert is posted on chain and the Bitcoin chain has grown by at least Δ_2 blocks, which should give enough time for the Verifier to "stop" a malicious Prover. How does the Verifier "stop" the Prover?

Connector outputs. A connector output is an output of a transaction when there are multiple transactions that could potentially spend it. For example, in Fig. 4, the 0-th output of Assert is a connector output because it can be spent by either the Success or the Failure transaction. If the Verifier posts a Failure transaction within Δ_2 blocks of the Assert transaction, he stops the Prover from posting the Success transaction because one of the inputs of the Success transaction is no longer available.

Another connector output, the output of ChallengeAssert, ensures that it can be spent in one of two ways. Either the Prover posts the WronglyChallenged transaction, whose $\mathsf{HashLock}(H(\mathsf{msg}))$ script requires the Prover to have decrypted the secret msg. Or, the Verifier posts the Failure transaction if the Prover doesn't post the WronglyChallenged transaction within Δ_1 blocks of the ChallengeAssert transaction. Posting Failure stops the Prover from ever posting the Success transaction.

Looking ahead, we will need Δ_2 sufficiently larger than Δ_1 to ensure that the Prover has enough time to post WronglyChallenged before the Verifier can post Failure.

Commitment using Lamport signatures. The Assert transaction is used by the Prover to post π_1 on chain and the ChallengeAssert transaction is used by the Verifier to post the input labels for the proof π_1 on chain. The locking scripts in these transactions' inputs ensure that the Verifier indeed posts input labels for the same proof π_1 that the Prover posts.

The input of the Assert transaction carries a locking script$\mathsf{CheckLampSig}(\mathsf{lpk}_P)$ which means the Prover must provide his Lamport signature μ on some message π_1. The script does not care about the message π_1 itself, but only that the Prover has committed to some message.

Conveniently, the input labels of the garbled circuit are exactly the Lamport signatures of the input π_1 under the garbled circuit's encoding key ek. So, the input of the ChallengeAssert transaction requires the Verifier to provide a Lamport signature under the public key epk (corresponding to the encoding key ek). Not only that, the script $\mathsf{CheckLampSigsMatch}(\mathsf{lpk}_P, \mathsf{epk})$ requires the Verifier to also post a Lamport signature for the same message under the Prover's Lamport key lpk_P. This ensures that the Verifier Lamport signs the same π_1 that the Prover did, since the Verifier cannot forge a Lamport signature under lpk_P for any message other than π_1 without knowing the Prover's Lamport secret key. This approach to matching the labels was first proposed in [Che25].

Putting it all together. Together, the transaction graphs ensure that i) the Prover commits to a proof element π_1, ii) the Verifier posts input labels for the same proof π_1, iii) the Prover can post Success only if he decrypts msg, and iv) the Verifier can post Failure otherwise. We prove in Sec. 7 that these transactions together with the witness encryption and garbled circuit satisfy the BitVM-core security properties.

6.2 Verifying Setup Correctness

We augment the setup phase of the protocol to allow the Prover and Verifier to verify that the setup was run correctly. This verification is done off-chain either party can abort if the verification fails, without losing any money, because this is done before the Bitcoin is locked.

Verifying setup involves verifying the following:

1. Statement $\boldsymbol{x}$ (given as input to both parties' setup algorithms)
2. Transactions $\mathcal{T}, \mathcal{S}$ (created deterministically given both parties' Bitcoin public keys, Prover's Lamport public key, hashes of the input labels, and hash of the secret msg)
3. Pre-signatures exchanged between the Prover and the Verifier (verified using the signature verification algorithm)
4. Randomness used for encryption/garbling is independent of the proof element π_1 (required for the condition in Lem. 5). This is satisfied in many use-cases, e.g., when the relation requires a certain transaction to be finalized in Ethereum, the Verifier cannot compute π_1 at setup. Otherwise, this can be achieved by adding some randomness to the Prover's Groth16 witness.
5. Witness encryption ciphertext $\mathsf{ct}_{\mathsf{setup}}$
6. Garbled circuit ciphertext $\mathsf{ct}_{\mathsf{GC}}$ and hashes of the input labels

What remains is to verify that the witness encryption ciphertext $\mathsf{ct}_{\mathsf{setup}}$ and the garbled circuit were computed correctly with respect to the statement $\boldsymbol{x}$ and the secret msg. This verification, carried out off-chain, can be achieved using various techniques from the literature, such as zero-knowledge proofs or cut-and-choose. We adopt cut-and-choose for its relative simplicity and efficiency.

The cut-and-choose setup protocol is shown in Alg. 4 (Prover) and Alg. 5 (Verifier). In this protocol, the Verifier generates N_{CC} instances, each with a secret msg_i, independent randomness r_i, and a garbled

circuit for $f_i(\pi) \to r_i\pi$. The Verifier commits to all instances, then the Prover randomly selects a subset $\mathcal{I} \subset [N_{\mathsf{CC}}]$ of size M_{CC}. For every instance $i \notin \mathcal{I}$, the Verifier "opens" the instance by revealing the underlying secrets and the randomness used for both encryption and garbling. The Prover then recomputes the corresponding ciphertexts and labels to check that they match the Verifier's commitments. The Prover finalizes the remaining M_{CC} instances by storing their ciphertexts. After verifying correctness for $N_{\mathsf{CC}} - M_{\mathsf{CC}}$ opened instances, the probability that all M_{CC} finalized instances are faulty is at most $p_{\text{err}} = \binom{N_{\mathsf{CC}}}{M_{\mathsf{CC}}}^{-1}$.

In the proving phase (Alg. 6), the Verifier reveals the input labels for all M_{CC} finalized instances and the Prover can post the success transaction if he successfully decrypts the secret for any one of these M_{CC} instances—an event that succeeds except with probability p_{err}. In Sec. 9, we explore different parameter choices and the trade-offs in setup time, off-chain storage, and on-chain cost. We also discuss possible optimizations to reduce on-chain cost by not posting all input labels on chain.

7 Security Proof

7.1 Security Proof Assuming Honest Setup

To warm up, we prove that the protocol in Algs. 1 and 2 is secure in the proving phase assuming that the setup is run honestly by both Prover and Verifier. First, we define this honest-setup security.

Definition 15 (Honest-Setup u-Robustness)**.** *For all NP relations $\mathcal{R}$, all statements $\boldsymbol{x}$ and witnesses $\boldsymbol{w}$ such that $(\boldsymbol{x}, \boldsymbol{w}) \in \mathcal{R}$, all* PPT *adversarial Verifiers V^*, all rounds $r \in \mathbb{N}$, the following holds:*

$$\Pr\left[\begin{array}{l|l} (b \neq 1) \vee (\forall\, \mathsf{tx} \in \mathcal{S}, & \mathsf{crs} \leftarrow \mathtt{Gen}(\mathcal{R}) \\ \mathsf{tx} \notin \mathcal{L}_P^{r+u}) & : \ (\mathcal{S}, \mathsf{st}_P, \mathsf{st}_V) \leftarrow \mathsf{out}\langle P_{\mathtt{Setup}}(\mathsf{crs}, \boldsymbol{x}), V_{\mathtt{Setup}}(\mathsf{crs}, \boldsymbol{x})\rangle \\ & \ b \leftarrow \mathsf{out}_P(\langle P_{\mathtt{Prove}}(\mathsf{st}_P, \boldsymbol{w}), V^*(\mathsf{st}_V)\rangle_r) \end{array}\right] \leqslant\ 2^{-\kappa} + \mathsf{negl}(\lambda) \quad (29)$$

Assigning the proving interaction to $\tau_{\mathtt{Prove}}$ below denotes executing that interaction, including its ledger effects; its parties' outputs are not used.

Definition 16 (Honest-Setup Knowledge Soundness)**.** *For all NP relations $\mathcal{R}$, all* PPT *adversarial Provers P^*, there exists a* PPT *extractor $\mathcal{E}$ such that for every benign auxiliary input $\mathsf{aux} \in \{0,1\}^{\mathsf{poly}(\lambda)}$:*

$$\Pr\left[\begin{array}{l|l} \exists\, r \in \mathbb{N}, \mathsf{tx} \in \mathcal{S}: & \mathsf{crs} \leftarrow \mathtt{Gen}(\mathcal{R}) \\ \mathsf{tx} \in \mathcal{L}_V^r & \boldsymbol{x} \leftarrow P^*(\mathsf{crs}, \mathsf{aux}) \\ \wedge\, (\boldsymbol{x}, \boldsymbol{w}) \notin \mathcal{R} & : \ (\mathcal{S}, \mathsf{st}_P, \mathsf{st}_V) \leftarrow \mathsf{out}\langle P_{\mathtt{Setup}}(\mathsf{crs}, \boldsymbol{x}), V_{\mathtt{Setup}}(\mathsf{crs}, \boldsymbol{x})\rangle \\ & \ \tau_{\mathtt{Prove}} \leftarrow \langle P^*(\mathsf{st}_P), V_{\mathtt{Prove}}(\mathsf{st}_V)\rangle \\ & \ \boldsymbol{w} \leftarrow \mathcal{E}_{P^*}(\mathsf{View}_{P^*}) \end{array}\right] \leqslant\ 2^{-\kappa} + \mathsf{negl}(\lambda) \quad (30)$$

Note that the difference with respect to the BitVM-core security definitions (Sec. 3) is that setup is run by $P_{\mathtt{Setup}}$ and $V_{\mathtt{Setup}}$ instead of an adversarial Prover P^* or Verifier V^*. Other than that, we provide the outputs of the setup phase to the adversarial Prover P^* or Verifier V^* in the proving phase since they were not participating in the setup phase. We still allow the adversarial prover to choose the statement $\boldsymbol{x}$ for a stronger security property. In the case of honest setup, the setup doesn't abort by definition.

7.1.1 Proof of Honest-Setup u-Robustness

Theorem 6. *Assuming $\mathcal{F}_{\mathsf{BTC}}$ satisfies safety, u_{BTC}-liveness, and (τ, s)-chain growth for $s > u_{BTC}$, the protocol in Algs. 1 and 2 with $\Delta_1 > u_{BTC}$ satisfies honest-setup u-robustness where $u = \tau^{-1}(\Delta_2 + u_{BTC})$.*

Proof. For a given NP relation $\mathcal{R}$, let $\mathsf{crs} \leftarrow \mathsf{Groth16.Gen}(\mathcal{R})$. Let $(\boldsymbol{x}, \boldsymbol{w}) \in \mathcal{R}$ and let $\boldsymbol{x}, \mathcal{T}, \mathcal{S}, \mathsf{st}_P, \mathsf{st}_V$ be generated as per Alg. 1, where $\mathcal{S} = \{\mathsf{tx}_{\mathsf{Success}}\}$ (see Alg. 3 line 4). Suppose the Prover runs $P_{\mathtt{Prove}}$ as per Alg. 2 starting at round r.

We go through the steps of the $P_{\mathtt{Prove}}$ algorithm and show that with overwhelming probability, $\mathsf{tx}_{\mathsf{Success}} \in \mathcal{L}_P^{r+u}$, and we calculate the value u.

Alg. 2 line 3: Prover generates (π_1, π_2, π_3). By perfect correctness of Groth16 (Thm. 2, Def. 1), this is a valid proof, i.e., $\mathtt{Verify}(\mathsf{crs}, \boldsymbol{x}, (\pi_1, \pi_2, \pi_3)) = 1$.

Alg. 2 line 5: Prover calls $\mathcal{F}_{\mathsf{BTC}}.\mathsf{WRITE}(\mathsf{tx}_{\mathsf{Assert}})$ at round r when $\mathcal{L}_P^r$ has height h_0. $\mathsf{tx}_{\mathsf{Assert}}$ is valid because it contains a valid Lamport signature for the key lpk_P (Alg. 2 line 4). It is also unstoppable with respect to the state st_V because st_V does not contain the Lamport signing key lsk_P. Therefore, due to liveness, $\mathsf{tx}_{\mathsf{Assert}} \in \mathcal{L}_P[: h_1]$ where $h_1 \leqslant h_0 + u_{BTC}$.

Case 1: $\mathsf{tx}_{\mathsf{ChallengeAssert}} \notin \mathcal{L}_P[: h_1 + \Delta_2]$ for all rounds. When $h(\mathcal{L}_P) = h_1 + \Delta_2$, Prover calls $\mathcal{F}_{\mathsf{BTC}}.\mathsf{WRITE}(\mathsf{tx}_{\mathsf{Success}})$ (Alg. 2 line 7). $\mathsf{tx}_{\mathsf{Success}}$ is valid because i) its parents[16] $\mathsf{tx}_{\mathsf{Assert}}$ is in the ledger, ii) its transaction witness contains the Verifier's pre-signature obtained during the setup and the Prover's signature (Alg. 2 line 6), and iii) the timelock on its input (1) has expired. Moreover, $\mathsf{tx}_{\mathsf{Success}}$ is u_{BTC}-unstoppable when $\Delta_1 > u_{BTC}$ because i) $\mathsf{tx}_{\mathsf{Failure}}$ is not valid because if $\mathsf{tx}_{\mathsf{ChallengeAssert}}$ was included by the adversary at height $> h_1 + \Delta_2$, the timelock on input (1) of $\mathsf{tx}_{\mathsf{Failure}}$ would not expire by height $h_1 + \Delta_2 + u_{BTC}$, and ii) the adversary (not knowing sk_P) cannot produce a transaction witness for any other transaction spending either input of $\mathsf{tx}_{\mathsf{Success}}$. Therefore, due to liveness, $\mathsf{tx}_{\mathsf{Success}} \in \mathcal{L}_P[: h_2]$ where $h_2 \leqslant h_1 + \Delta_2 + u_{BTC}$.

Case 2: $\mathsf{tx}_{\mathsf{ChallengeAssert}} \in \mathcal{L}_P[h_3]$ for some $h_3 < h_1 + \Delta_2$ at some round.

1. Upon seeing $\mathsf{tx}_{\mathsf{ChallengeAssert}}$ in $\mathcal{L}_P$, Prover extracts the input labels posted by the Verifier in $\mathsf{tx}_{\mathsf{ChallengeAssert}}$ (Alg. 2 line 9). For $\mathsf{tx}_{\mathsf{ChallengeAssert}}$ to be valid, its transaction witness must satisfy the script $\mathsf{CheckLampSigsMatch}(\mathsf{lpk}_P, \mathsf{ek})$ which requires a valid Lamport signature under the Prover's key lpk_P and a valid Lamport signature under the Verifier's key ek for the same message. Since the adversary does not know lsk_P, he cannot produce a valid Lamport signature under lpk_P for any message other than π_1. Therefore, the extracted labels satisfy $L \leftarrow \mathsf{Encode}(\mathsf{ek}, \pi_1)$.

2. Prover evaluates the garbled circuit (Alg. 2 line 10) and due to correctness of the garbled circuit (Def. 2), $\mathsf{Eval}(\mathsf{ct}_{\mathsf{GC}}, L) = \mathtt{Enc}_{\mathsf{prove}}(\mathsf{crs}, \pi_1; r)$ (Constr. 1).

3. Prover decrypts the message (Alg. 2 line 11) and due to correctness of the witness encryption scheme, the Prover learns msg.

4. Prover constructs the transaction witness for the WronglyChallenged transaction (Alg. 2 line 12). This transaction is valid because i) its parent $\mathsf{tx}_{\mathsf{ChallengeAssert}}$ is in the ledger, and ii) its transaction witness contains the message msg. This transaction is u_{BTC}-unstoppable when $\Delta_1 > u_{BTC}$ because the adversary does not know sk_P and the timelock Δ_1 on its input (0) does not expire in u_{BTC} blocks. Therefore, due to liveness, $\mathsf{tx}_{\mathsf{WronglyChallenged}} \in \mathcal{L}_P[: h_4]$ where $h_4 \leqslant h_3 + u_{BTC}$.

5. This guarantees that for all $h > h_4$, $\mathsf{tx}_{\mathsf{Failure}} \notin \mathcal{L}_P[: h]$ because input (1) of $\mathsf{tx}_{\mathsf{Failure}}$ is no longer available. This holds in particular for $h \geqslant h_3 + \Delta_1$ since $h_4 \leqslant h_3 + u_{BTC} < h_3 + \Delta_1$. Also, we already know that for all $h < h_3 + \Delta_1$, $\mathsf{tx}_{\mathsf{Failure}} \notin \mathcal{L}_P[: h]$ because the timelock on input (1) of $\mathsf{tx}_{\mathsf{Failure}}$ does not expire at height $< h_3 + \Delta_1$. Therefore, for all h, $\mathsf{tx}_{\mathsf{Failure}} \notin \mathcal{L}_P[: h]$.

[16] A transaction tx is a parent of another transaction tx' one of the inputs of tx' is an output of tx.

6. The Prover calls $\mathcal{F}_{\mathsf{BTC}}.\mathsf{WRITE}(\mathsf{tx}_{\mathsf{Success}})$ when $h(\mathcal{L}_P) = h_1 + \Delta_2$. Following the same arguments as in Case 1, this transaction is valid and unstoppable, therefore by liveness, $\mathsf{tx}_{\mathsf{Success}} \in \mathcal{L}_P[: h_2]$ where $h_2 \leqslant h_1 + \Delta_2 + u_{BTC}$.

Due to the chain growth property (Thm. 3), $\mathsf{tx}_{\mathsf{Success}} \in \mathcal{L}_P^{r+u}$ where $u = \tau^{-1}(\Delta_2 + u_{BTC})$.

□

7.1.2 Proof of Honest-Setup Knowledge Soundness

To prove knowledge soundness, we first combine the witness encryption scheme Constr. 1 and the garbled circuit Constr. 2 and prove that no adversary given the ciphertexts of both and the input labels of the garbled circuit can decrypt the message without knowing a valid witness for the relation $\mathcal{R}$.

Lemma 11. *Let* $\mathtt{Gen}, \mathtt{Enc}_{\mathsf{setup}}, \mathtt{Enc}_{\mathsf{prove}}, \mathsf{Dec}$ *(Constr. 1) be an adaptively extractably one-way secure witness encryption scheme (Def. 11) and let* $\mathsf{Garble}, \mathsf{Encode}, \mathsf{Eval}$ *(Constr. 2) be adaptively private (Def. 2). Then in the random oracle model, for all NP relations* $\mathcal{R}$*, for all* PPT *adversaries* $\mathcal{A}$*, there exists a* PPT *extractor* $\mathcal{E}$ *such that*

$$\Pr\left[\begin{array}{l|l} \mathsf{Hash}_{BTC}(\mathsf{msg}') = h_{\mathsf{msg}} & \mathsf{crs} \leftarrow \mathtt{Gen}(\mathcal{R}) \\ \quad \wedge (\boldsymbol{x}, \boldsymbol{w}) \notin \mathcal{R} & \boldsymbol{x} \leftarrow \mathcal{A}(\mathsf{crs}) \\ & \mathsf{msg} \leftarrow\!\!\$\ \mathcal{M} \\ & r \leftarrow\!\!\$\ \mathbb{F}_q^* \\ & \mathsf{ct}_{\mathsf{setup}} \leftarrow \mathtt{Enc}_{\mathsf{setup}}(\mathsf{crs}, \boldsymbol{x}, \mathsf{msg}; r) \\ & \mathsf{ct}_{\mathsf{GC}}, \mathsf{ek} \leftarrow \mathsf{Garble}(r) \\ & h_{\mathsf{msg}} \leftarrow \mathsf{Hash}_{BTC}(\mathsf{msg}) \\ & \mathsf{epk}_j^b \leftarrow \mathsf{Hash}_{BTC}(\mathsf{ek}_j^b) \quad \forall j \in \{1, \ldots, 2n\}, b \in \{0,1\} \\ & \pi_1 \leftarrow \mathcal{A}(\mathsf{ct}_{\mathsf{setup}}, \mathsf{ct}_{\mathsf{GC}}, h_{\mathsf{msg}}, \mathsf{epk}) \\ & L \leftarrow \mathsf{Encode}(\mathsf{ek}, \pi_1) \\ & \mathsf{msg}' \leftarrow \mathcal{A}(L) \\ & \boldsymbol{w} \leftarrow \mathcal{E}(\mathsf{View}_{\mathcal{A}}) \end{array}\right] \leqslant \mathsf{negl}(\lambda)$$

Proof. Let $q_H = q_H(\lambda)$ be a polynomial upper bound on the number of Hash_{BTC} queries made in any of the experiments below, counting the final output as one additional query. We assume that the output length of Hash_{BTC} is λ bits and recall from Def. 11 that $1/|\mathcal{M}| = \mathsf{negl}(\lambda)$. We use domain separation between Hash_{BTC} and the random oracle used inside the witness-encryption construction.

Fix a stateful PPT adversary $\mathcal{A}$. Let

$$W := [\mathsf{Hash}_{BTC}(\mathsf{msg}') = h_{\mathsf{msg}}] \qquad \text{and} \qquad F := [(\boldsymbol{x}, \boldsymbol{w}) \notin \mathcal{R}].$$

We construct the extractor defining F below. All the probabilities that follow include its coins and refer to the same execution as W.

Hybrid H_0: the real experiment. This is the experiment in the lemma statement. We write msg_0 for its uniformly sampled encrypted message. Thus $h_{\mathsf{msg}} = \mathsf{Hash}_{BTC}(\mathsf{msg}_0)$, $(\mathsf{ct}_{\mathsf{GC}}, \mathsf{ek}) \leftarrow \mathsf{Garble}(r)$, and $L \leftarrow \mathsf{Encode}(\mathsf{ek}, \pi_1)$.

Hybrid H_1: simulated garbling. Let C_r be the circuit that, on input π_1, outputs

$$C_r(\pi_1) = \mathtt{Enc}_{\mathsf{prove}}(\mathsf{crs}, \pi_1; r) = r\pi_1.$$

In H_1 we instead compute

$$(\widetilde{\mathsf{ct}_{\mathsf{GC}}}, \mathsf{st}) \leftarrow \mathsf{Sim}_1(\mathsf{topo}(C_r)), \qquad \widetilde{L} \leftarrow \mathsf{Sim}_2(\mathsf{st}, \pi_1, C_r(\pi_1)).$$

The values $\{\mathsf{epk}_j^b\}_{j\in[2n],b\in\{0,1\}}$ are sampled as $4n$ independent uniform λ-bit strings. The oracle interface presented to $\mathcal{A}$ forwards ordinary queries to Hash_{BTC} and interposes at a revealed simulated label $\widetilde{L}_j$ by returning the already sampled value $\mathsf{epk}_j^{(\pi_1)_j}$.

We now justify this hybrid in the presence of the correlated ciphertext $\mathsf{ct}_{\mathsf{setup}}$. Suppose that a PPT predicate P distinguishes H_0 from H_1. Construct an adaptive-GC-privacy distinguisher $\mathcal{D}$ as follows. It generates crs, runs $\mathcal{A}$ to obtain $\boldsymbol{x}$, samples msg_0, r, and computes

$$\mathsf{ct}_{\mathsf{setup}} \leftarrow \mathtt{Enc}_{\mathsf{setup}}(\mathsf{crs}, \boldsymbol{x}, \mathsf{msg}_0; r), \qquad h_{\mathsf{msg}} \leftarrow \mathsf{Hash}_{BTC}(\mathsf{msg}_0).$$

It then submits the circuit C_r to its garbling challenger and retains $(r, \mathsf{ct}_{\mathsf{setup}}, h_{\mathsf{msg}})$ in its state. Upon receiving the challenge garbling, it gives $(\mathsf{ct}_{\mathsf{setup}}, \mathsf{ct}_{\mathsf{GC}}, h_{\mathsf{msg}}, \mathsf{epk})$ to $\mathcal{A}$, submits the resulting π_1 as its adaptive input, gives the returned labels to $\mathcal{A}$, completes the execution, and outputs P. Thus the same r is used in $\mathsf{ct}_{\mathsf{setup}}$ and hardwired in C_r in both branches. If the challenger is real, this produces H_0; if it is simulated, it produces H_1, subject only to the following random-oracle bookkeeping.

The privacy challenger does not reveal the unused encoding-key labels to $\mathcal{D}$. Sampling epk uniformly implicitly assigns their oracle images. Once the $2n$ selected labels are returned, $\mathcal{D}$ recognizes those points and interposes consistently. An inconsistency is observable only if a selected label was queried before it was returned, if a programmed point collides with a previous query, or if $\mathcal{A}$ queries one of the $2n$ unrevealed labels. The first two events have probability at most $\frac{(q_H+4n)^2}{2^\lambda}$. In the ideal garbling experiment the unrevealed labels retain λ bits of min-entropy, so their query probability is at most $2nq_H/2^\lambda$. If this probability were non-negligibly larger in the real experiment, the query transcript would distinguish real from simulated garbling. Equivalently, by correctness, possession of such a label together with the selected labels would permit evaluation on a second input and reveal a second value $C_r(\pi_1') = r\pi_1'$, whereas the adaptive simulator receives only $C_r(\pi_1)$. This uses adaptive privacy. Consequently, for every PPT predicate P of the complete execution,

$$\left|\Pr_{H_0}[P] - \Pr_{H_1}[P]\right| \leqslant \mathsf{negl}(\lambda)_{\mathsf{GC}} + \frac{2nq_H}{2^\lambda} + \frac{(q_H+4n)^2}{2^\lambda} = \mathsf{negl}(\lambda). \tag{31}$$

Hybrid H_2: an independent message hash. Starting from H_1, sample an independent $\mathsf{msg}_1 \leftarrow\!\!\$\ \mathcal{M}$ and replace $h_{\mathsf{msg}} = \mathsf{Hash}_{BTC}(\mathsf{msg}_0)$ by $h_{\mathsf{msg}} = \mathsf{Hash}_{BTC}(\mathsf{msg}_1)$. Everything else, including the encryption of msg_0, is unchanged. This hybrid can be run by an adversary in the one-way WE game: the WE challenger samples and encrypts the unknown msg_0, while the adversary itself samples the independent msg_1.

Couple H_1 and H_2 by sampling a common uniform value h and assigning it to the oracle point msg_0 in H_1 and to msg_1 in H_2. Apart from the negligible point-collision events already included above, the two executions are identical until $\mathcal{A}$ queries one of these two points. Let Q_i^+ be the event that $\mathcal{A}$ queries Hash_{BTC} at msg_i, or finally outputs $\mathsf{msg}' = \mathsf{msg}_i$ without having made that query. Thus the final output is treated as one additional oracle query for both messages.

Since msg_1 is independent of the ciphertext and all other inputs to $\mathcal{A}$, the standard random-oracle preimage bound gives

$$\Pr[Q_1^+] \leqslant \frac{q_H+1}{|\mathcal{M}|} + \frac{q_H+1}{2^\lambda} = \mathsf{negl}(\lambda). \tag{32}$$

It remains to bound Q_0^+ jointly with extraction failure. Queries made before $\mathsf{ct}_{\mathsf{setup}}$ is received hit the subsequently sampled uniform msg_0 with probability at most $q_H/|\mathcal{M}|$; we absorb this negligible term below. Index the remaining at most q_H oracle queries of $\mathcal{A}$ and its final output by $t \in [q_H + 1]$. For each t, define a stateful adversary $\mathcal{B}_t$ in the one-way WE experiment of Def. 11. It runs the H_2 simulation described above. At the t-th oracle query, it records the query point and stops the inner execution before answering that query; for $t = q_H + 1$ it records $\mathcal{A}$'s final output. It completes the required WE interaction using the π_1 already selected by $\mathcal{A}$, or an arbitrary fixed syntactically valid π_1 if the inner execution stopped earlier, and outputs the recorded string as its message guess. If the relevant string is msg_0, then $\mathcal{B}_t$ has recovered the message encrypted by its WE challenger.

By adaptive extractable one-way security, for every t there is an extractor $\mathcal{E}_t$ such that

$$\Pr\left[\mathcal{B}_t \text{ outputs } \mathsf{msg}_0 \ \wedge \ (\boldsymbol{x}, \mathcal{E}_t(\mathsf{View}_{\mathcal{B}_t})) \notin \mathcal{R}\right] \leqslant \mathsf{negl}(\lambda). \tag{33}$$

Define $\mathcal{E}$ to run all $q_H + 1$ extractors on their corresponding projected views and to return the first output that satisfies $\mathcal{R}$ (or an arbitrary string if none does). Therefore, F implies the failure of every $\mathcal{E}_t$. A union bound in H_2 yields

$$\Pr_{H_2}[Q_0^+ \wedge F] \leqslant (q_H + 1)\mathsf{negl}(\lambda) + \frac{q_H}{|\mathcal{M}|} = \mathsf{negl}(\lambda). \tag{34}$$

Stopping before the candidate query is answered is important: conditioned on no earlier Q_1^+, its prefix has exactly the same distribution in the above coupling of H_1 and H_2. Thus Eq. (34) also bounds the contribution of $Q_0^+ \wedge F$ when passing from H_1 to H_2.

For the concrete WE construction, we take each $\mathcal{E}_t$ to be the extractor constructed in Lem. 2. That extractor depends only on the public WE values and the forwarded random-oracle and generic-group transcripts; it does not depend on the private state of Sim_1 or Sim_2. Hence the combined extractor above is well-defined on the real view in H_0 as well.

Finally, in H_2 the hash target is the image of the independent message msg_1. Finding either msg_1 or a different preimage of the same random oracle output has probability at most

$$\Pr_{H_2}[W] \leqslant \frac{q_H + 1}{|\mathcal{M}|} + \frac{q_H + 1}{2^\lambda} = \mathsf{negl}(\lambda). \tag{35}$$

Using the identical-until-bad coupling, and then Eqs. (32), (34) and (35), we obtain

$$\Pr_{H_1}[W \wedge F] \leqslant \Pr_{H_2}[W \wedge F] + \Pr[Q_1^+] + \Pr[Q_0^+ \wedge F] + \mathsf{negl}(\lambda) \leqslant \mathsf{negl}(\lambda). \tag{36}$$

Instantiating Eq. (31) with the PPT predicate $W \wedge F$ and combining it with Eq. (36) proves

$$\Pr_{H_0}[W \wedge F] \leqslant \mathsf{negl}(\lambda).$$

□

Theorem 7. *Assuming $\mathcal{F}_{\mathsf{BTC}}$ satisfies safety, u_{BTC}-liveness, and (τ, s)-chain growth for $s > u_{BTC}$, the protocol in Algs. 1 and 2 with $\Delta_2 > \Delta_1 + 2u_{BTC}$ satisfies honest-setup knowledge soundness.*

Proof. For a given NP relation $\mathcal{R}$, let $\mathsf{crs} \leftarrow \mathsf{Groth16.Gen}(\mathcal{R})$ and let $\boldsymbol{x}, \mathcal{T}, \mathcal{S}, \mathsf{st}_P, \mathsf{st}_V$ be generated as per Alg. 1, where $\mathcal{S} = \{\mathsf{tx}_{\mathsf{Success}}\}$ (see Alg. 3 line 4). Suppose the Verifier runs $V_{\mathtt{Prove}}$ as per Alg. 2. Suppose that at some round r and some block height h, $\mathsf{tx}_{\mathsf{Success}} \in \mathcal{L}_V^r[h]$.

For $\mathsf{tx}_{\mathsf{Success}}$ to be valid, its ancestor $\mathsf{tx}_{\mathsf{Assert}}$ must be in the ledger. Suppose $\mathsf{tx}_{\mathsf{Assert}} \in \mathcal{L}_V^r[h_0]$. Due to the timelock on input (1) of $\mathsf{tx}_{\mathsf{Success}}$, $\mathsf{tx}_{\mathsf{Assert}}$ must have been included at least Δ_2 blocks earlier, i.e., $h_0 \leqslant h - \Delta_2$.

The Verifier, upon seeing $\mathsf{tx}_{\mathsf{Assert}} \in \mathcal{L}_V[h_0]$ at some round, called $\mathcal{F}_{\mathsf{BTC}}.\mathsf{WRITE}(\mathsf{tx}_{\mathsf{ChallengeAssert}})$ (Alg. 2 line 20). This transaction is valid because its transaction witness contains both parties' signatures, the Prover's Lamport signature and the Verifier's Lamport signature for the same message π_1 (Alg. 2 line 19). This transaction is unstoppable because the adversary does not know ek or sk_V. Due to liveness, $\mathsf{tx}_{\mathsf{ChallengeAssert}} \in \mathcal{L}_V[: h_1]$ where $h_1 \leqslant h_0 + u_{BTC}$.

When $h(\mathcal{L}_V) = h_1 + \Delta_1$, the Verifier calls $\mathcal{F}_{\mathsf{BTC}}.\mathsf{WRITE}(\mathsf{tx}_{\mathsf{Failure}})$ (Alg. 2 line 22). This transaction is valid because its transaction witness contains the Prover's pre-signature obtained during the setup and the Verifier's signature (Alg. 2 line 21), and the timelock Δ_1 on its input (1) has expired.

Case 1: $\mathsf{tx}_{\mathsf{Failure}} \in \mathcal{L}_V[: h_2]$ for $h_2 = h_1 + \Delta_1 + u_{BTC}$. This is a contradiction because if $\mathsf{tx}_{\mathsf{Failure}} \in \mathcal{L}_V[: h_2]$, and $\mathsf{tx}_{\mathsf{Success}} \in \mathcal{L}_V[: h]$, then the ledger $\mathcal{L}_V$ is invalid because both transactions contain a common input.

Case 2: $\mathsf{tx}_{\mathsf{Failure}} \notin \mathcal{L}_V[: h_2]$ for $h_2 = h_1 + \Delta_1 + u_{BTC}$. Due to liveness, this must mean that $\mathsf{tx}_{\mathsf{Failure}}$ is not u_{BTC}-unstoppable with respect to the ledger $\mathcal{L}_V[h_1 + \Delta_1]$ and the adversary's state $\mathsf{st} = (\mathsf{crs}, \boldsymbol{x}, \mathcal{T}, \mathcal{S}, \mathsf{st}_P, \mathsf{aux})$ for some auxiliary input aux. In particular, the adversary must have created a sequence of blocks such that $\mathsf{tx}_{\mathsf{Failure}}$ is invalid when placed in one of the blocks. Since the timelock Δ_1 has already expired, the adversary must have created a sequence of blocks containing a valid transaction tx' which shares an input with $\mathsf{tx}_{\mathsf{Failure}}$. The adversary cannot create any other transaction spending output (0) of $\mathsf{tx}_{\mathsf{Assert}}$ (which is input (0) of $\mathsf{tx}_{\mathsf{Failure}}$) because i) given $\Delta_2 > \Delta_1 + 2u_{BTC}$ and $h_1 \leqslant h_0 + u_{BTC}$,

$$h_1 + \Delta_1 + u_{BTC} \leqslant h_0 + \Delta_1 + 2u_{BTC} < h_0 + \Delta_2,$$

so the timelock Δ_2 will not expire by height $h_1 + \Delta_1 + u_{BTC}$, and ii) the adversary does not know sk_V and so cannot produce a valid signature for any transaction other than $\mathsf{tx}_{\mathsf{Success}}$. The adversary cannot create any other transaction spending the output of $\mathsf{tx}_{\mathsf{ChallengeAssert}}$ (which is input (1) of $\mathsf{tx}_{\mathsf{Failure}}$) through the leaf $\mathsf{CheckSig}(\mathsf{pk}_V) \wedge \mathsf{RelTimelock}(\Delta_1)$ because the adversary does not know sk_V. Therefore, the adversary can stop $\mathsf{tx}_{\mathsf{Failure}}$ only by spending the output of $\mathsf{tx}_{\mathsf{ChallengeAssert}}$ through the leaf $\mathsf{HashLock}(\mathsf{Hash}_{BTC}(\mathsf{msg})) \wedge \mathsf{CheckSig}(\mathsf{pk}_P)$. The transaction witness of such a transaction must contain msg' such that $\mathsf{Hash}_{BTC}(\mathsf{msg}') = \mathsf{Hash}_{BTC}(\mathsf{msg})$.

We make the invocation of Lem. 11 precise. Fix the adversarial Prover P^* and its benign auxiliary input aux. Construct a stateful adversary $\mathcal{A}_{P^*}$ for the experiment of Lem. 11. Given crs, it runs $\boldsymbol{x} \leftarrow P^*(\mathsf{crs}, \mathsf{aux})$ and outputs $\boldsymbol{x}$. Upon receiving $(\mathsf{ct}_{\mathsf{setup}}, \mathsf{ct}_{\mathsf{GC}}, h_{\mathsf{msg}}, \mathsf{epk})$ from the lemma's experiment, $\mathcal{A}_{P^*}$ completes an honest setup execution around these values. More precisely, it runs the honest $P_{\mathtt{Setup}}$ algorithm, samples the ordinary signing coins of the honest Verifier, constructs $\mathcal{T}$ and all required pre-signatures, and gives the resulting st_P to P^*. This has exactly the distribution of the honest setup in the theorem: the lemma's experiment has sampled the same msg, r, and has generated $\mathsf{ct}_{\mathsf{setup}}, \mathsf{ct}_{\mathsf{GC}}, h_{\mathsf{msg}}, \mathsf{epk}$ exactly as $V_{\mathtt{Setup}}$ does. Knowledge of ek is not needed to complete this simulation; its only later use by the honest Verifier is to compute the labels for the adaptively selected π_1.

The adversary $\mathcal{A}_{P^*}$ now runs P^* and emulates the honest Verifier and the ledger functionality. If during this execution a valid transaction spends the output of $\mathsf{tx}_{\mathsf{ChallengeAssert}}$ through $\mathsf{HashLock}(h_{\mathsf{msg}}) \wedge \mathsf{CheckSig}(\mathsf{pk}_P)$, then $\mathcal{A}_{P^*}$ reads the preimage msg' from that transaction's witness and outputs it in the WE–GC experiment. Otherwise it outputs an arbitrary string. Therefore, whenever the transaction tx' identified above exists, $\mathsf{Hash}_{BTC}(\mathsf{msg}') = h_{\mathsf{msg}}$.

Let $\mathcal{E}_{\mathcal{A}_{P^*}}$ be the extractor supplied by Lem. 11 for this wrapper adversary. Define the protocol extractor $\mathcal{E}_{P^*}$ by applying $\mathcal{E}_{\mathcal{A}_{P^*}}$ to the corresponding view embedded in View_{P^*}. This is well-defined for the extractor used in the proof of Lem. 11: as noted there, it depends only on the public setup values and the

random-oracle and generic-group transcripts forwarded from P^*, and not on the honest Verifier's secret signing key or on ek.

Let S denote the joint event that $\mathsf{tx}_{\mathsf{Success}}$ is included in the Verifier's ledger view and that $(\boldsymbol{x}, \mathcal{E}_{P^*}(\mathsf{View}_{P^*})) \notin \mathcal{R}$. Condition on safety, liveness, and signature unforgeability. Case 1 is impossible. In Case 2, the preceding transaction argument produces msg' satisfying the hashlock; hence, pointwise on this conditioned execution,

$$S \implies \left[\mathsf{Hash}_{BTC}(\mathsf{msg}') = h_{\mathsf{msg}} \ \wedge \ (\boldsymbol{x}, \mathcal{E}_{\mathcal{A}_{P^*}}(\mathsf{View}_{\mathcal{A}_{P^*}})) \notin \mathcal{R}\right].$$

The probability of the event on the right is negligible by Lem. 11. Adding the failure probabilities of safety, liveness, and signature unforgeability gives

$$\Pr[S] \leqslant 2^{-\kappa} + \mathsf{negl}(\lambda),$$

which is precisely Eq. (30).

□

Theorem 8. *Assuming $\mathcal{F}_{\mathsf{BTC}}$ satisfies safety, u_{BTC}-liveness, and (τ, s)-chain growth for $s > u_{BTC}$, the protocol in Algs. 1 and 2 with $\Delta_1 > u_{BTC}$ and $\Delta_2 > \Delta_1 + 2u_{BTC}$ satisfies honest-setup knowledge soundness and honest-setup u-robustness where $u = \tau^{-1}(\Delta_2 + u_{BTC})$.*

Proof. From Thm. 7 and Thm. 6. □

8 Extensions and Optimizations

8.1 Multiple Verifiers and Provers

In this paper, we defined BitVM-core as a two-party protocol. The BitVM-core definition and the BABE protocol can be easily extended to capture multiple Verifiers (as in BitVM2-core [LAA+25]) so that soundness holds as long as at least one Verifier is honest. Similarly, BABE can support multiple Provers so that the first Prover to generate a valid proof for his statement can withdraw the Bitcoin.

8.2 Optimistic Path

Leveraging techniques from BitVM2 [LAA+25], BABE can incorporate an optimistic path that bypasses the full proving phase when the Verifier can independently confirm the statement. For instance, if the statement is "Bob repaid his loan on Ethereum", the Verifier can directly check the finalized repayment transaction on the Ethereum blockchain. In such situations, neither a proof from the Prover nor the publication of input labels by the Verifier is necessary on Bitcoin, thus significantly reducing on-chain costs.

9 Evaluation

9.1 Honest Setup

We begin with the metrics in Fig. 2 for the honest-setup protocol. We also justify the values used for BitVM3 in Fig. 2.

On-chain cost. The BABE on-chain cost in Fig. 2 is based on the dispute-path transactions Assert, ChallengeAssert, and WronglyChallenged that are directly responsible for the on-chain proof verification, using the vSizes reported in Tab. 1. This is the basis for the \$56.90 on-chain cost entry reported for BABE in Fig. 2. We estimate USD cost using 1 sat ≈ 0.000955 USD (i.e., 1 BTC $\approx 95{,}500$ USD).

Table 1: On-chain cost for the honest-setup experiment in Fig. 2

Transaction	vSize (vB)	Fee (sat)	Feerate (sat/vB)	USD
Assert	9,240	20,611	2.23	19.68
ChallengeAssert	17,400	38,673	2.22	36.93
WronglyChallenged	149	302	2.02	0.29

The BitVM3 on-chain cost entry in Fig. 2 is based on the corresponding dispute-path transactions from a previously published BitVM3 on-chain experiment, normalized to the same fee rate as the BABE experiment.

Off-chain runtime breakdown. The setup time (174.90 ms) and decryption time (126.53 ms) reported for BABE in Fig. 2 are single-instance point estimates, according to our measurements. The detailed breakdown numbers below come from a representative run whose totals differ from these point estimates by at most $\approx 1.5\%$. To justify where this time goes, we decompose the dominant off-chain paths into protocol subroutines that correspond to steps in Algs. 1 and 2 and the constructions in Constr. 1 and Sec. 5.6.

On the setup path, the benchmark totals 177.398 ms and is dominated by generation of the DRE-selection tables used for the DRE encoding step in Sec. 5.6. (136.896 ms, 77.17%).

The remaining setup components are the privacy-free Boolean gadget at 16.808 ms (9.47%), that validates the curve point and derives the feature bits (Sec. 5.6). plus 18.029 ms (10.16%) of garbling overhead and 5.665 ms (3.19%) of non-garbling work (WE setup, Lamport-key hashes, and message commitments).

On the decryption path, the benchmark totals 125.834 ms and is dominated by evaluation of the garbled circuit (Alg. 2 and Sec. 5.6). (111.538 ms, 88.64%), plus 8.888 ms (7.06%) of overhead from label conversions/packing.

Off-chain storage (ciphertext size). The BABE storage entry in Fig. 2 is the serialized size of the per-instance off-chain setup artifact that the Prover retains through the protocol (cf. Sec. 6): the witness-encryption setup ciphertext $\mathsf{ct}_{\mathsf{setup}}$ from Constr. 1, the garbled-circuit ciphertext $\mathsf{ct}_{\mathsf{GC}}$ from Constr. 2 (used to derive $\mathsf{ct}_{\mathsf{prove}}$), and the associated hashes/commitments referenced by the Bitcoin scripts. Using compressed canonical serialization of the implementation artifact, this totals 22.16 MiB, according to our calculations.

The witness-encryption ciphertexts $(\mathsf{ct}_{\mathsf{setup}}, \mathsf{ct}_{\mathsf{prove}})$ contribute 480 B. The garbled-circuit ciphertext $\mathsf{ct}_{\mathsf{GC}}$ dominates and splits into 6.35 MiB for the Boolean-Yao gadgets that validate π_1 and derive the feature vector $\overline{\boldsymbol{u}}(\pi_1)$ (Sec. 5) and 15.77 MiB for the DRE-selection tables implementing the DRE encoding (Def. 13 and Thm. 4); the remaining ≈ 38 KB are Lamport-key hashes and message commitments plus serialization overhead.

The 15.77 MiB DRE term matches the sparsity-aware estimate: for $n = 254$ (Sec. 5), the DRE encoding step stores

$$n\big((3n+1)+(4n+1)+(n+1)\big) \;=\; n(8n+3) \;=\; 516{,}890$$

serialized field elements (Jacobian X, Y, Z with $(3n+1,\ 4n+1,\ n+1)$ field elements per bit). With 32-byte canonical serialization per field element, this is

$$516{,}890 \times 32 \;=\; 16{,}540{,}480\ \text{B} \;\approx\; 15.77\,\text{MiB}.$$

BitVM3 reference point (ciphertext size). The BitVM3 storage entry in Fig. 2 follows from the reported 2.7 billion non-free gates [Bit25b, WAA+26]; using half-gates garbling and interpreting this count as half-gates (one 16-byte ciphertext per non-free gate), this yields $2.7 \times 10^9 \times 16\,\mathrm{B} \approx 41{,}200\,\mathrm{MiB}$.

9.2 Cut-and-Choose Setup Verification

We use cut-and-choose to verify setup correctness against a malicious setup generator. We refer to Algs. 4 and 5 and App. B for the protocol definition.

Off-chain setup cost. Tab. 2 reports setup cost between two protocol roles (Prover and Verifier) executed on a single machine for different $(N_{\mathsf{CC}}, M_{\mathsf{CC}})$ choices; timings exclude network latency. We additionally report peak RAM usage and a breakdown of garbling and evaluation components.[17]

Table 2: BABE cut-and-choose setup cost for different parameter choices.

N_{CC}	M_{CC}	Setup time (min:s)	Peak RAM usage (GB)	Garbling (s)	Evaluation (s)
78	10	0:06	1.27	1.44	1.45
95	9	0:06	1.21	1.85	1.32
124	8	0:07	1.29	2.54	1.18
181	7	0:09	1.05	3.57	1.05
307	6	0:14	1.12	6.32	0.86
669	5	0:28	1.46	13.22	0.74
2,268	4	1:37	1.83	49.15	0.59
18,756	3	12:31	3.93	383.87	0.41

N_{CC}: total cut-and-choose instances. M_{CC}: finalized instances for evaluation. Setup time is wall-clock time for cut-and-choose setup (timings exclude network latency). Setup time includes both Prover and Verifier computation. Garbling time is measured on the garbler node for all N_{CC} instances (not summed across both roles). Evaluation time is total decoding time for all finalized instances using the on-chain input labels (Sec. 5.6). Peak RAM usage is the maximum RAM used during execution. Hardware: CPU: AMD Ryzen 7 7840U(16 CPU);.

Parameter choices and statistical security. We choose the $(N_{\mathsf{CC}}, M_{\mathsf{CC}})$ pairs in Tab. 2 so that the soundness error of the cut-and-choose setup $\binom{N_{\mathsf{CC}}}{M_{\mathsf{CC}}}^{-1}$ is at most 2^{-40}. In this step, the Verifier opens $N_{\mathsf{CC}} - M_{\mathsf{CC}}$ randomly chosen instances and retains M_{CC} unopened instances as finalized. Soundness can fail only if all opened instances are correct while all M_{CC} finalized instances are incorrect.

Comparison to BitVM3. As an open-sourced reference point for garbled-circuit-based Groth16 verification, we additionally report cut-and-choose setup costs for BitVM3 [Bit25b, WAA+26]. To make the comparison parameter-aligned, we use the same $(N_{\mathsf{CC}}, M_{\mathsf{CC}})$ pairs as in Tab. 2 and report setup time and the garbling/evaluation breakdown. Other garbled-circuit-based approaches do not provide public implementations and benchmarks at comparable levels of detail. We use $(N_{\mathsf{CC}}, M_{\mathsf{CC}}) = (181, 7)$ as the main

[17] These components correspond to the generation and use of the full protocol ciphertexts: in each cut-and-choose instance the Verifier produces the witness-encryption setup ciphertext $\mathsf{ct}_{\mathsf{setup}}$ (Constr. 1) and the garbled-circuit ciphertext $\mathsf{ct}_{\mathsf{GC}}$ (Constr. 2); the evaluation component measures the Prover's evaluation of $\mathsf{ct}_{\mathsf{GC}}$ on the on-chain input labels (Sec. 5.6) to derive the proving ciphertext $\mathsf{ct}_{\mathsf{prove}}$ (Sec. 6.1.2).

operating point for BitVM3, since the garbling time at this point is not very large while the on-chain footprint remains acceptable. We discuss augmenting BitVM3 with zk-SNARK-soldering in Sec. 9.3.

Table 3: BitVM3 cut-and-choose setup cost for different parameter choices (same $(N_{\mathsf{CC}}, M_{\mathsf{CC}})$ grid as Tab. 2).

N_{CC}	M_{CC}	Setup time (h:m:s)	Garbling (h:m:s)	Evaluation (h:m:s)
78	10	1:28:54	0:44:27	0:03:45
95	9	1:47:58	0:53:59	0:03:45
124	8	2:19:06	1:09:33	0:03:45
181	7	3:24:24	1:42:12	0:03:45
307	6	5:43:24	2:51:42	0:03:45
669	5	12:23:06	6:11:33	0:03:45
2,268	4	41:02:00	20:31:00	0:03:45

Times are in h:mm:ss format. Garbling time is from [Bit25b, WAA+26]; setup time is estimated as twice the garbling time. Evaluation time is expected to be negligible relative to garbling. Hardware: CPU: AMD Ryzen 7 7840U(16 CPU);.

Operating points. For BitVM3, we primarily use $(N_{\mathsf{CC}}, M_{\mathsf{CC}}) = (181, 7)$; for BABE we use $M_{\mathsf{CC}} = 4$ in the end-to-end evaluation (cf. Tabs. 2 and 6).

On-chain implication. In our end-to-end on-chain evaluation we focus on $M_{\mathsf{CC}} = 4$. Cut-and-choose only guarantees that at least one of the 4 instances is correct with high probability but the Prover doesn't know which one. Hence in the baseline, four sets of input labels must be posted on chain. This directly drives the on-chain footprint and motivates soldering.

9.3 Soldering (zk-SNARK-soldering)

zk-SNARK-soldering is an optimization that reduces the number of distinct on-chain input-label sets from M_{CC} down to one, by binding finalized instances to a base instance and proving correctness of this binding with a soldering zk-SNARK proof. Our prototype uses the SP1 zkVM as the proving backend, but the stack is interchangeable and could be replaced by any zk-SNARK or zk-STARK system.

Soldering idea (informal). Let $\mathcal{I}$ denote the set of finalized instances and let $b = \min(\mathcal{I})$ be a base instance. Let $\{L_{i,j,0}, L_{i,j,1}\}_{j=1}^{2n}$ denote the garbler's per-wire label pairs for instance i ($2n$ is the number of input bits, with n as in Sec. 5). For each $i \in \mathcal{I}\backslash\{b\}$, define the per-wire, per-bit deltas $\Delta_{i,j,\beta} := L_{i,j,\beta} \oplus L_{b,j,\beta}$ for $\beta \in \{0,1\}, j \in [2n]$. The soldering zk-SNARK proof attests that these deltas are consistent with the commitments fixed during cut-and-choose, enabling derivation of all finalized per-wire label pairs from the base instance's labels. Equivalently, the proof certifies that there exists a single collection of per-wire label pairs for the base instance and for all other finalized instances, and that the deltas $\{\Delta_{i,j,\beta}\}$ are exactly the XOR differences between these labels. This check is performed by verifying a single soldering zk-SNARK proof, and any verification failure aborts the protocol.

Soldering reduces on-chain cost because only the base instance's tag set is posted on-chain. All other finalized instances' per-wire label pairs can be derived off-chain from the base labels and the proven deltas $\{\Delta_{i,j,\beta}\}$, so the on-chain footprint no longer scales with M_{CC}.

We focus on $M_{\mathsf{CC}} = 4$ because soldering introduces additional zk-SNARK proving time during setup. In particular, approaches that treat each finalized instance separately lead to overhead that grows with

Table 4: BitVM3 soldering overhead (1019 labels) as a function of the number of finalized instances.

M_{CC}	Soldering time (s)
5	1529.11
6	1740.24
7	1977.10

Soldering time is the time to generate a zk-SNARK proof attesting to the correctness of input-label bindings across M_{CC} finalized instances. Hardware: CPU: AMD Ryzen 7 7840U(16 CPU);.

Table 5: BABE soldering overhead (508 labels) as a function of the number of finalized instances.

M_{CC}	Soldering time (s)
4	745.50
5	865.59
6	969.37

Soldering time is the time to generate a zk-SNARK proof attesting to the correctness of input-label bindings across M_{CC} finalized instances. Hardware: CPU: AMD Ryzen 7 7840U(16 CPU);.

Table 6: Off-chain timings at the operating point $M_{\mathsf{CC}} = 4$.

Configuration	Setup time for Verifier (s)	Evaluation time (s)
C&C baseline	49.15	0.59
C&C + zk-SNARK-soldering	49.15 + 745.50s	0.59

M_{CC}. We therefore focus on the $M_{\mathsf{CC}} = 4$ operating point in the end-to-end evaluation, and use $M_{\mathsf{CC}} = 5$ only as a reference point when discussing scaling.

We report soldering times for $M_{\mathsf{CC}} \in \{5, 6, 7\}$ (BitVM3) and $M_{\mathsf{CC}} \in \{4, 5, 6\}$ (BABE) in Tabs. 4 and 5, and compare the baseline and zk-SNARK-soldering configurations at $M_{\mathsf{CC}} = 4$ in Tab. 6.

BitVM3 with zk-SNARK-soldering. We can also augment BitVM3 with zk-SNARK-soldering to reduce the on-chain input-label footprint from M_{CC} sets to one. The soldering overhead depends on M_{CC} and the number of input labels, but is independent of the specific garbled circuit; thus the BABE soldering times in Tab. 5 serve as a lower bound, while the BitVM3 soldering times in Tab. 4 apply directly to BitVM3. Total setup time for BitVM3 with zk-SNARK-soldering is the sum of the C&C setup time from Tab. 3 and the corresponding soldering overhead.

Implementation. All reported measurements are obtained from our prototype implementation.

9.4 Verifiable Shamir Secret Sharing

C&C + VSSS (general idea). As a possible improvement over the C&C baseline, we consider compressing the on-chain input-label footprint using verifiable Shamir secret sharing (VSSS), inspired by Glock [Eag25] and the "efficient verifiable cut-and-choose" design notes [Lab25, BOB25b]. In the setup

phase, the Verifier secret-shares the Prover input labels across instances and publishes commitments to the shares, while binding each instance to the committed values (e.g., via hashes and nonce commitments). In the proving phase, the Prover interpolates from the committed shares to reconstruct the input labels for the M_{CC} finalized instances and proceeds as in the baseline. Conceptually, this targets the same bottleneck as zk-SNARK-soldering—reducing the number of distinct on-chain input-label sets from M_{CC} down to one—but replaces zk proving with interpolation and commitment checks.

BitVM3 baseline at $(N_{\mathsf{CC}}, M_{\mathsf{CC}}) = (181, 7)$. We continue to use $(N_{\mathsf{CC}}, M_{\mathsf{CC}}) = (181, 7)$ as the main BitVM3 operating point (cf. Tab. 3). At this point, the reported BitVM3 setup time is `3:24:24` (with garbling time `1:42:12`).[18] To isolate what VSSS adds, we report below the incremental VSSS overhead as a function of $(N_{\mathsf{CC}}, M_{\mathsf{CC}})$.

VSSS overhead across parameter sets. Tab. 7 reports the incremental overhead of adding the VSSS layer on top of cut-and-choose for different $(N_{\mathsf{CC}}, M_{\mathsf{CC}})$ choices.

These measurements come from a draft prototype intended to estimate the overhead under a mock-on-chain design (no end-to-end transaction integration). At the smallest operating point $(N_{\mathsf{CC}}, M_{\mathsf{CC}}) = (78, 10)$, the measured setup-time overhead is 32.96 s. Considering that N_{CC} garbling of BABE will take less than a second, we can estimate that the setup time will be around 40 seconds.

Table 7: Incremental overhead of VSSS over cut-and-choose for different parameter choices (prototype measurements).

N_{CC}	M_{CC}	Setup time overhead (s)	Peak RAM (GB)
78	10	32.96	0.98
95	9	40.86	1.20
124	8	53.84	1.57
181	7	83.93	2.31
307	6	159.23	3.92
669	5	332.61	8.56

Measurements run the VSSS layer standalone with the corresponding $(N_{\mathsf{CC}}, M_{\mathsf{CC}})$ parameters and 508 input bits. Setup time overhead is wall-clock time for setup between two roles (Prover and Verifier) executed on a single machine (timings exclude network latency). Hardware: CPU: AMD Ryzen 7 7840U(16 CPU);.

On-chain integration note. An end-to-end on-chain integration for VSSS still requires a concrete transaction design that ties the committed shares to the on-chain protocol logic (e.g., via adaptor signatures (ad-sig) and nonce commitments). As a result, we do not yet report on-chain costs for VSSS. Nevertheless, relative to prior BitVM3-centric experiments [BOB25b] we expect substantially lower on-chain footprint, since BABE requires far fewer evaluator input labels than BitVM3.

Acknowledgements

We thank the Babylon Labs team members for performing the on-chain experiments and for designing the illustrations. We thank Liam Eagen and Ying Tong Lai for early discussions about [EL26]. We thank Ertem Nusret Tas and Nikhil Vanjani for reviewing the security proofs of the paper. The work of first and second authors is supported in part by the AFOSR Award FA9550-24-1-0156, and research grants from the

[18] See the caption of Tab. 3 for how setup time is estimated from garbling time.

Bakar Fund, the Stellar Development Foundation, Supra Inc., and Byzantine Research Inc. The work of the fifth author is partially supported by the Stanford Input-Output Global Research Hub. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of these institutes.

References

[AAL+24] Lukas Aumayr, Zeta Avarikioti, Robin Linus, Matteo Maffei, Andrea Pelosi, Christos Stefo, and Alexei Zamyatin. BitVM: Quasi-turing complete computation on Bitcoin. Cryptology ePrint Archive, Report 2024/1995, 2024. 4

[ADJK25] Lukas Aumayr, Jesus Diaz, Dimitar Jetchev, and Aggelos Kiayias. Cardinal: Bridging bitcoin with ownership preservation. Cryptology ePrint Archive, Report 2025/2196, 2025. 18

[AFP25] Amit Agarwal, Rex Fernando, and Benny Pinkas. Efficiently-thresholdizable batched identity based encryption, with applications. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part III*, volume 16002 of *LNCS*, pages 69–100. Springer, Cham, August 2025. 8

[AIK04] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. Cryptography in NC^0. In *45th FOCS*, pages 166–175. IEEE Computer Society Press, October 2004. 9

[BBG05] Dan Boneh, Xavier Boyen, and Eu-Jin Goh. Hierarchical identity based encryption with constant size ciphertext. In Ronald Cramer, editor, *EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 440–456. Springer, Berlin, Heidelberg, May 2005. 11

[BC16] Olivier Blazy and Céline Chevalier. Structure-preserving smooth projective hashing. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part II*, volume 10032 of *LNCS*, pages 339–369. Springer, Berlin, Heidelberg, December 2016. 7, 21

[BCMS20] Benedikt Bünz, Alessandro Chiesa, Pratyush Mishra, and Nicholas Spooner. Recursive proof composition from accumulation schemes. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part II*, volume 12551 of *LNCS*, pages 1–18. Springer, Cham, November 2020. 10

[BCPR14] Nir Bitansky, Ran Canetti, Omer Paneth, and Alon Rosen. On the existence of extractable one-way functions. In David B. Shmoys, editor, *46th ACM STOC*, pages 505–514. ACM Press, May / June 2014. 19

[BCTV14] Eli Ben-Sasson, Alessandro Chiesa, Eran Tromer, and Madars Virza. Scalable zero knowledge via cycles of elliptic curves. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part II*, volume 8617 of *LNCS*, pages 276–294. Springer, Berlin, Heidelberg, August 2014. 10

[BFOQ25] Jan Bormet, Sebastian Faust, Hussien Othman, and Ziyan Qu. BEAT-MEV: Epochless approach to batched threshold encryption for MEV prevention. In Lujo Bauer and Giancarlo Pellegrino, editors, *USENIX Security 2025*, pages 3457–3476. USENIX Association, August 2025. 8

[BHR12a] Mihir Bellare, Viet Tung Hoang, and Phillip Rogaway. Adaptively secure garbling with applications to one-time programs and secure outsourcing. In Xiaoyun Wang and Kazue Sako, editors, *ASIACRYPT 2012*, volume 7658 of *LNCS*, pages 134–153. Springer, Berlin, Heidelberg, December 2012. 14, 33

[BHR12b] Mihir Bellare, Viet Tung Hoang, and Phillip Rogaway. Foundations of garbled circuits. In Ting Yu, George Danezis, and Virgil D. Gligor, editors, *ACM CCS 2012*, pages 784–796. ACM Press, October 2012. 13, 33

[Bit25a] Bitlayer. Bitvm bridge testnet | bitlayer. https://docs.bitlayer.org/docs/BitVMBridge/Multi-Chain/testnet/, 2025. Last accessed: 2024-01-08'. 1

[Bit25b] BitVM. Garbled snark verifier. https://github.com/BitVM/garbled-snark-verifier, 2025. GitHub repository. Last accessed: 2026-01-11. 6, 7, 48, 49

[BL20] Fabrice Benhamouda and Huijia Lin. Mr NISC: Multiparty reusable non-interactive secure computation. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part II*, volume 12551 of *LNCS*, pages 349–378. Springer, Cham, November 2020. 8

[BMM+21] Benedikt Bünz, Mary Maller, Pratyush Mishra, Nirvan Tyagi, and Psi Vesely. Proofs for inner pairing products and applications. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part III*, volume 13092 of *LNCS*, pages 65–97. Springer, Cham, December 2021. 10

[BN06] Paulo S. L. M. Barreto and Michael Naehrig. Pairing-friendly elliptic curves of prime order. In Bart Preneel and Stafford Tavares, editors, *SAC 2005*, volume 3897 of *LNCS*, pages 319–331. Springer, Berlin, Heidelberg, August 2006. 6

[BOB25a] BOB. Bob | bridge. https://app.gobob.xyz/en/bridge, 2025. Last accessed: 2024-01-08'. 1

[BOB25b] BOB. BOB lowers onchain costs for BitVM3 via cut-and-choose implementation to $10.91. https://gobob.xyz/blog/bob-lowers-onchain-costs-for-bitvm3, December 2025. BOB Blog. Last accessed: 2026-01-12. 50, 51

[Boy08] Xavier Boyen. The uber-assumption family (invited talk). In Steven D. Galbraith and Kenneth G. Paterson, editors, *PAIRING 2008*, volume 5209 of *LNCS*, pages 39–56. Springer, Berlin, Heidelberg, September 2008. 11

[BP15] Elette Boyle and Rafael Pass. Limits of extractability assumptions with distributional auxiliary input. In Tetsu Iwata and Jung Hee Cheon, editors, *ASIACRYPT 2015, Part II*, volume 9453 of *LNCS*, pages 236–261. Springer, Berlin, Heidelberg, November / December 2015. 19

[Cat25] Catalyst. Bitcoin prism. https://github.com/catalystsystem/bitcoinprism-evm, 2025. GitHub repository. Last accessed: 2026-01-14. 4

[CFH+22] Matteo Campanelli, Dario Fiore, Semin Han, Jihye Kim, Dimitris Kolonelos, and Hyunok Oh. Succinct zero-knowledge batch proofs for set accumulators. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022*, pages 455–469. ACM Press, November 2022. 10

[CFK24] Matteo Campanelli, Dario Fiore, and Hamidreza Khoshakhlagh. Witness encryption for succinct functional commitments and applications. In Qiang Tang and Vanessa Teague, editors, *PKC 2024, Part II*, volume 14602 of *LNCS*, pages 132–167. Springer, Cham, April 2024. 8

[CGPW25] Arka Rai Choudhuri, Sanjam Garg, Guru-Vamsi Policharla, and Mingyuan Wang. Practical mempool privacy via one-time setup batched threshold encryption. In Lujo Bauer and Giancarlo Pellegrino, editors, *USENIX Security 2025*, pages 3477–3495. USENIX Association, August 2025. 8

[Che25] Weikeng Chen. SoK: BitVM with succinct on-chain cost. Cryptology ePrint Archive, Report 2025/1253, 2025. 5, 39

[Cit25] Citrea. Citrea bridge | citrea. https://citrea.xyz/bridge, 2025. Last accessed: 2024-01-08'. 1

[DLT+24] Xinshu Dong, Orfeas Stefanos Thyfronitis Litos, Ertem Nusret Tas, David Tse, Robin Linus Woll, Lei Yang, and Mingchao Yu. Remote staking with economic safety. *CoRR*, abs/2408.01896, 2024. 6

[Dry17] Thaddeus Dryja. Discreet log contracts. https://static1.squarespace.com/static/6675a0d5fc9e317c60db9b37/t/66e4597b7f23866561a64a95/1726241147904/discreet+log+contracts+paper.pdf, 2017. Last accessed: 2026-01-12. 6

[Eag25] Liam Eagen. Glock: Garbled locks for Bitcoin. Cryptology ePrint Archive, Report 2025/1485, 2025. 5, 6, 50

[EL26] Liam Eagen and Ying Tong Lai. Argo MAC: Garbling with elliptic curve MACs. Cryptology ePrint Archive, Paper 2026/049, 2026. 1, 9, 23, 51

[FBFL25] Ariel Futoransky, Fadi Barbàra, Ramses Fernandez, and Gabriel Larotonda. OHMG: One hot modular garbling. Cryptology ePrint Archive, Report 2025/2338, 2025. 6

[FHAS24] Nils Fleischhacker, Mathias Hall-Andersen, and Mark Simkin. Extractable witness encryption for KZG commitments and efficient laconic OT. In Kai-Min Chung and Yu Sasaki, editors, *ASIACRYPT 2024, Part II*, volume 15485 of *LNCS*, pages 423–453. Springer, Singapore, December 2024. 8

[FKdP23] Dario Fiore, Dimitris Kolonelos, and Paola de Perthuis. Cuckoo commitments: Registration-based encryption and key-value map commitments for large spaces. In Jian Guo and Ron Steinfeld, editors, *ASIACRYPT 2023, Part V*, volume 14442 of *LNCS*, pages 166–200. Springer, Singapore, December 2023. 8

[FKN94] Uri Feige, Joe Kilian, and Moni Naor. A minimal model for secure computation. In *Proceedings of the 26th Annual ACM Symposium on Theory of Computing*, STOC '94, pages 554–563, New York, NY, USA, 1994. ACM. 9, 29

[GGKS25] Sanjam Garg, Aarushi Goel, Dimitris Kolonelos, and Rohit Sinha. Jigsaw: Doubly private smart contracts. Cryptology ePrint Archive, Report 2025/1147, 2025. 10

[GGPR13] Rosario Gennaro, Craig Gentry, Bryan Parno, and Mariana Raykova. Quadratic span programs and succinct NIZKs without PCPs. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 626–645. Springer, Berlin, Heidelberg, May 2013. 13

[GGSW13] Sanjam Garg, Craig Gentry, Amit Sahai, and Brent Waters. Witness encryption and its applications. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 467–476. ACM Press, June 2013. 6, 21

[GGW24] Sanjam Garg, Aarushi Goel, and Mingyuan Wang. How to prove statements obliviously? In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part X*, volume 14929 of *LNCS*, pages 449–487. Springer, Cham, August 2024. 10

[GHK+25] Sanjam Garg, Mohammad Hajiabadi, Dimitris Kolonelos, Abhiram Kothapalli, and Guru-Vamsi Policharla. A framework for witness encryption from linearly verifiable SNARKs and applications. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part III*, volume 16002 of *LNCS*, pages 504–539. Springer, Cham, August 2025. 8, 21

[GIKM00] Yael Gertner, Yuval Ishai, Eyal Kushilevitz, and Tal Malkin. Protecting data privacy in private information retrieval schemes. *J. Comput. Syst. Sci.*, 60(3):592–629, 2000. 5, 8

[GKL15] Juan A. Garay, Aggelos Kiayias, and Nikos Leonardos. The Bitcoin backbone protocol: Analysis and applications. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 281–310. Springer, Berlin, Heidelberg, April 2015. 14, 17

[GKP+13] Shafi Goldwasser, Yael Tauman Kalai, Raluca A. Popa, Vinod Vaikuntanathan, and Nickolai Zeldovich. How to run Turing machines on encrypted data. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 536–553. Springer, Berlin, Heidelberg, August 2013. 6, 20

[GKPW24] Sanjam Garg, Dimitris Kolonelos, Guru-Vamsi Policharla, and Mingyuan Wang. Threshold encryption with silent setup. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part VII*, volume 14926 of *LNCS*, pages 352–386. Springer, Cham, August 2024. 1, 7, 8, 21

[GMN22] Nicolas Gailly, Mary Maller, and Anca Nitulescu. SnarkPack: Practical SNARK aggregation. In Ittay Eyal and Juan A. Garay, editors, *FC 2022*, volume 13411 of *LNCS*, pages 203–229. Springer, Cham, May 2022. 10

[Gro16] Jens Groth. On the size of pairing-based non-interactive arguments. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 305–326. Springer, Berlin, Heidelberg, May 2016. 6, 12, 13

[Her18] Maurice Herlihy. Atomic cross-chain swaps. *arXiv preprint arXiv:1801.09515*, 2018. Last accessed: 2026-01-12. 6

[Hio22] Leona Hioki. Trustless bitcoin bridge creation with witness encryption. https://ethresear.ch/t/trustless-bitcoin-bridge-creation-with-witness-encryption/11953, February 2022. Ethereum Research. Last accessed: 2026-01-11. 7

[IK00] Yuval Ishai and Eyal Kushilevitz. Randomizing polynomials: A new representation with applications to round-efficient secure computation. In *41st FOCS*, pages 294–304. IEEE Computer Society Press, November 2000. 9

[IK02] Yuval Ishai and Eyal Kushilevitz. Perfect constant-round secure computation via perfect randomizing polynomials. In Peter Widmayer, Francisco Triguero Ruiz, Rafael Morales Bueno, Matthew Hennessy, Stephan Eidenbenz, and Ricardo Conejo, editors, *ICALP 2002*, volume 2380 of *LNCS*, pages 244–256. Springer, Berlin, Heidelberg, July 2002. 9

[Ish13] Yuval Ishai. Randomization techniques for secure computation. In Manoj Prabhakaran and Amit Sahai, editors, *Secure Multi-Party Computation*, volume 10 of *Cryptology and Information Security Series*, pages 222–248. IOS Press, 2013. 9, 24, 28, 29

[IW14] Yuval Ishai and Hoeteck Wee. Partial garbling schemes and their applications. In Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias, editors, *ICALP 2014, Part I*, volume 8572 of *LNCS*, pages 650–662. Springer, Berlin, Heidelberg, July 2014. 9

[KS08] Vladimir Kolesnikov and Thomas Schneider. Improved garbled circuit: Free XOR gates and applications. In Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfsdóttir, and Igor Walukiewicz, editors, *ICALP 2008, Part II*, volume 5126 of *LNCS*, pages 486–498. Springer, Berlin, Heidelberg, July 2008. 6, 14

[KST22] Abhiram Kothapalli, Srinath Setty, and Ioanna Tzialla. Nova: Recursive zero-knowledge arguments from folding schemes. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part IV*, volume 13510 of *LNCS*, pages 359–388. Springer, Cham, August 2022. 10

[LAA+25] Robin Linus, Lukas Aumayr, Zeta Avarikioti, Matteo Maffei, Andrea Pelosi, Orfeas Thyfronitis Litos, Christos Stefo, David Tse, and Alexei Zamyatin. Bridging Bitcoin to second layers via BitVM2. Cryptology ePrint Archive, Report 2025/1158, 2025. To appear in Usenix Security 2026. 1, 4, 5, 7, 18, 46

[Lab25] Alpen Labs. Efficient verifiable cut and choose for glock. https://hackmd.io/@alpen/B1QfSSO5gg, 2025. HackMD. Last accessed: 2026-01-12. 50

[Lam79] Leslie Lamport. Constructing digital signatures from a one-way function. Technical Report SRI-CSL-98, SRI International Computer Science Laboratory, October 1979. 5, 14, 15

[Lin23] Robin Linus. Bitvm: Compute anything on bitcoin. https://bitvm.org/bitvm.pdf, December 2023. Last accessed: 2024-01-08'. 4

[Lin24] Robin Linus. Bitvm 3s - garbled circuits for efficient computation on bitcoin. https://bitvm.org/bitvm3.pdf, 2024. Last accessed: 2024-01-08'. 1, 5

[LP07] Yehuda Lindell and Benny Pinkas. An efficient protocol for secure two-party computation in the presence of malicious adversaries. In Moni Naor, editor, *EUROCRYPT 2007*, volume 4515 of *LNCS*, pages 52–78. Springer, Berlin, Heidelberg, May 2007. 6

[LP09] Yehuda Lindell and Benny Pinkas. A proof of security of Yao's protocol for two-party computation. *Journal of Cryptology*, 22(2):161–188, April 2009. 13, 33

[Mau05] Ueli M. Maurer. Abstract models of computation in cryptography (invited paper). In Nigel P. Smart, editor, *10th IMA International Conference on Cryptography and Coding*, volume 3796 of *LNCS*, pages 1–12. Springer, Berlin, Heidelberg, December 2005. 11

[MLLP25] Varun Madathil, Arthur Lazzaretti, Zeyu Liu, and Charalampos Papamanthou. TACITA: Threshold aggregation without client interaction. Cryptology ePrint Archive, Report 2025/1579, 2025. 8

[Nak09] Satoshi Nakamoto. Bitcoin open source implementation of p2p currency. https://satoshi.nakamotoinstitute.org/posts/p2pfoundation/1/, 2009. Last accessed: 2024-01-08'. 4

[OKMZ25] Michele Orrù, George Kadianakis, Mary Maller, and Greg Zaverucha. Beyond the circuit: How to minimize foreign arithmetic in ZKP circuits. *CiC*, 2(1):23, 2025. 10

[OWWB20] Alex Ozdemir, Riad S. Wahby, Barry Whitehat, and Dan Boneh. Scaling verifiable computation using efficient set accumulators. In Srdjan Capkun and Franziska Roesner, editors, *USENIX Security 2020*, pages 2075–2092. USENIX Association, August 2020. 10

[PD16] Joseph Poon and Thaddeus Dryja. The bitcoin lightning network: Scalable off-chain instant payments. https://lightning.network/lightning-network-paper.pdf, 2016. Last accessed: 2026-01-12. 6

[PHGR13] Bryan Parno, Jon Howell, Craig Gentry, and Mariana Raykova. Pinocchio: Nearly practical verifiable computation. In *2013 IEEE Symposium on Security and Privacy*, pages 238–252. IEEE Computer Society Press, May 2013. 13

[Rub24] Jeremy Rubin. Delbrag. https://rubin.io/public/pdfs/delbrag.pdf, 2024. Last accessed: 2024-01-08'. 1, 5

[SGB24] István András Seres, Noemi Glaeser, and Joseph Bonneau. Short paper: Naysayer proofs. In Jeremy Clark and Elaine Shi, editors, *FC 2024, Part II*, volume 14745 of *LNCS*, pages 22–32. Springer, Cham, March 2024. 5

[Sho97] Victor Shoup. Lower bounds for discrete logarithms and related problems. In Walter Fumy, editor, *EUROCRYPT'97*, volume 1233 of *LNCS*, pages 256–266. Springer, Berlin, Heidelberg, May 1997. 11

[WAA+26] Robin Linus Woll, Ioannis Alexopoulos, Lukas Aumayr, Zeta Avarikioti, Matteo Maffei, and David Tse. BitVM3: Efficient bitcoin bridges via garbled circuits. Cryptology ePrint Archive, Report 2026/933, 2026. 5, 6, 7, 18, 48, 49

[Was08] Lawrence C. Washington. *Elliptic Curves: Number Theory and Cryptography*. Chapman and Hall/CRC, 2nd edition, 2008. 27

[Wik20] Bitcoin Wiki. Atomic swap. https://en.bitcoin.it/wiki/Atomic_swap, 2020. Last accessed: 2026-01-12. 6

[Wik21] Bitcoin Wiki. Hash time locked contracts. https://en.bitcoin.it/wiki/Hash_Time_Locked_Contracts, 2021. Last accessed: 2026-01-12. 6

[WNT20] Pieter Wuille, Jonas Nick, and Anthony Towns. Bip 0341, taproot: Segwit version 1 spending rules. https://en.bitcoin.it/wiki/BIP_0341, January 2020. Last accessed: 2024-01-08'. 15

[WOS+25] Anna P. Y. Woo, Alex Ozdemir, Chad Sharp, Thomas Pornin, and Paul Grubbs. Efficient proofs of possession for legacy signatures. In Marina Blanton, William Enck, and Cristina Nita-Rotaru, editors, *2025 IEEE Symposium on Security and Privacy*, pages 3291–3308. IEEE Computer Society Press, May 2025. 10

[Yao82] Andrew Chi-Chih Yao. Protocols for secure computations (extended abstract). In *23rd FOCS*, pages 160–164. IEEE Computer Society Press, November 1982. 5, 9, 13

[ZRE15] Samee Zahur, Mike Rosulek, and David Evans. Two halves make a whole - reducing data transfer in garbled circuits using half gates. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 220–250. Springer, Berlin, Heidelberg, April 2015. 6, 14

A Honest-Setup BABE Protocol Details

A.1 Transactions

The detailed specifications of the transactions are given in this section. Each transaction specifies the inputs, outputs, the locking scripts for each input and output, and the transaction witnesses.

A.1.1 Notation

We follow the transaction model of Sec. 2.5: a transaction is $\mathsf{tx} = (\mathsf{inputs}, \mathsf{tx_witnesses}, \mathsf{outputs})$. Each input is of the form $\mathsf{in} = (\mathsf{PrevTx}, \mathsf{outIndex}, \mathsf{leaf})$, where PrevTx is the previous transaction, outIndex is the index of the output in that transaction, and leaf is the leaf of the Taproot tree to be satisfied for that input. Each output is $\mathsf{out} = (a, \mathsf{lockScript})$ where a is the value in that UTXO and lockScript is the locking script. The i-th witness in tx_witnesses is the data (e.g., signatures, hash preimages) supplied to satisfy the i-th input's leaf script.

In the tables below, $*$ denotes that the corresponding field can be *any* value of the appropriate type. For example, in an input $(*, *, *)$, the first component can be any previous transaction, the second any output index, and the third any leaf (i.e., any UTXO and any leaf thereof). When a component is not $*$, it is fixed: e.g., $(\mathsf{tx}_{\mathsf{Assert}}, 0, \mathsf{leaf})$ means the input references the 0-th output of the Assert transaction and the spender must satisfy the leaf leaf.

Output locking scripts are often written as $\langle \mathsf{leaf}_0, \ldots, \mathsf{leaf}_{k-1} \rangle$. This denotes a *Taproot tree* (taptree): the UTXO can be spent by satisfying *one* of the leaves leaf_j. A transaction that spends such an output must specify which leaf it is satisfying and supply a witness that satisfies that leaf's script. The locking scripts CheckSig, RelTimelock, HashLock, CheckLampSig, CheckLampSigsMatch, and their combinations are as defined in Sec. 2.5.

A.1.2 Example: The Assert Transaction

We illustrate the notation and the role of inputs, outputs, and witnesses using the Assert transaction as an example.

Inputs and input script. The Assert transaction has a single input: $(*, *, \langle \mathsf{CheckLampSig}(\mathsf{lpk}_P) \rangle)$. The first two components are $*$: the input may reference *any* UTXO (any previous transaction and any output index). The third component fixes the *leaf* to be satisfied: $\langle \mathsf{CheckLampSig}(\mathsf{lpk}_P) \rangle$. So the UTXO being spent must have a locking script that is a Taproot tree containing the leaf $\mathsf{CheckLampSig}(\mathsf{lpk}_P)$ (e.g., a tree that includes this leaf among others). In practice, this is typically another UTXO set up so that its taptree includes this leaf. To spend it, the Prover must provide a witness that satisfies $\mathsf{CheckLampSig}(\mathsf{lpk}_P)$, i.e., a Lamport signature under lpk_P on some message (see Eq. (9)).

Witness. The Tx Witness for input (0) is $\mu_1, \ldots, \mu_{2n}$. These are the $2n$ components of the Lamport signature μ on the message (e.g., the proof element π_1; n is the bit length of $\mathbb{F}_p$ as in Sec. 5): for each bit of the message, the signer reveals the corresponding preimage from the Lamport key. This witness satisfies $\mathsf{CheckLampSig}(\mathsf{lpk}_P)$ and thus binds the Prover to a single message when posting the Assert transaction.

Outputs and their scripts. The Assert transaction has two outputs, both with amount 0.

- **Output (0)** has locking script $\langle \mathsf{RelTimelock}(\Delta_2) \wedge \mathsf{CheckSig}(\mathsf{pk}_P) \wedge \mathsf{CheckSig}(\mathsf{pk}_V),\ \mathsf{CheckSig}(\mathsf{pk}_P) \wedge \mathsf{CheckSig}(\mathsf{pk}_V) \rangle$. This is a Taproot tree with *two* leaves. The UTXO can be spent by satisfying *exactly one* of them:

  1. **First leaf** $\mathsf{RelTimelock}(\Delta_2) \wedge \mathsf{CheckSig}(\mathsf{pk}_P) \wedge \mathsf{CheckSig}(\mathsf{pk}_V)$: the spending transaction must be included at least Δ_2 blocks after the Assert transaction, and the witness must contain signatures σ_P, σ_V on the transaction skeleton under pk_P and pk_V. This path is used by the Success transaction.
  2. **Second leaf** $\mathsf{CheckSig}(\mathsf{pk}_P) \wedge \mathsf{CheckSig}(\mathsf{pk}_V)$: the witness must contain σ_P and σ_V (no timelock). This path is used by the Failure transaction.

Therefore, output (0) of Assert can be spent either (i) after Δ_2 blocks with both parties' Schnorr signatures, or (ii) immediately with both parties' Schnorr signatures.

- **Output (1)** has locking script $\langle \mathsf{CheckLampSigsMatch}(\mathsf{lpk}_P, \mathsf{lpk}_V) \wedge \mathsf{CheckSig}(\mathsf{pk}_V) \wedge \mathsf{CheckSig}(\mathsf{pk}_P) \rangle$. There is a single leaf, so this output can be spent only one way: the witness must satisfy $\mathsf{CheckLampSigsMatch}(\mathsf{lpk}_P, \mathsf{l}$ (Lamport signatures under lpk_P and lpk_V on the *same* message; see Eq. (10)) and provide σ_V, σ_P. This path is used by the ChallengeAssert transaction.

Deposit **Transaction**	
Inputs	(0) $(*, *, *)$
Outputs	(0) $(v, \langle \mathsf{CheckSig}(\mathsf{pk}_P) \wedge \mathsf{CheckSig}(\mathsf{pk}_V) \rangle)$
Tx Witness	(0) $*$

Assert **Transaction**	
Inputs	(0) $(*, *, \langle \mathsf{CheckLampSig}(\mathsf{lpk}_P) \rangle)$ (see Eq. (9))
Outputs	(0) $(0, \langle \mathsf{RelTimelock}(\Delta_2) \wedge \mathsf{CheckSig}(\mathsf{pk}_P) \wedge \mathsf{CheckSig}(\mathsf{pk}_V), \mathsf{CheckSig}(\mathsf{pk}_P) \wedge \mathsf{CheckSig}(\mathsf{pk}_V) \rangle)$ (1) $(0, \langle \mathsf{CheckLampSigsMatch}(\mathsf{lpk}_P, \mathsf{lpk}_V) \wedge \mathsf{CheckSig}(\mathsf{pk}_V) \wedge \mathsf{CheckSig}(\mathsf{pk}_P) \rangle)$ (see Eq. (10))
Tx Witness	(0) $\mu_1, \ldots, \mu_\ell$

ChallengeAssert **Transaction**	
Inputs	(0) $(\mathsf{tx}_{\mathsf{Assert}}, 1, \langle \mathsf{CheckLampSigsMatch}(\mathsf{lpk}_P, \mathsf{lpk}_V) \wedge \mathsf{CheckSig}(\mathsf{pk}_V) \wedge \mathsf{CheckSig}(\mathsf{pk}_P) \rangle)$ (see Eq. (10))
Outputs	(0) $(0, \langle \mathsf{RelTimelock}(\Delta_1) \wedge \mathsf{CheckSig}(\mathsf{pk}_V), \mathsf{HashLock}(\mathsf{Hash}_{BTC}(\mathsf{msg})) \wedge \mathsf{CheckSig}(\mathsf{pk}_P) \rangle)$
Tx Witness	(0) $L_1, \ldots, L_\ell, \mu_1, \ldots, \mu_\ell, \sigma_V, \sigma_P$

Failure **Transaction**	
Inputs	(0) $(\mathsf{tx}_{\mathsf{Assert}}, 0, \mathsf{CheckSig}(\mathsf{pk}_P) \wedge \mathsf{CheckSig}(\mathsf{pk}_V))$ (1) $(\mathsf{tx}_{\mathsf{ChallengeAssert}}, 0, \mathsf{RelTimelock}(\Delta_1) \wedge \mathsf{CheckSig}(\mathsf{pk}_V))$
Outputs	(0) $(0, \langle \mathsf{CheckSig}(\mathsf{pk}_V) \rangle)$
Tx Witness	(0) σ_P, σ_V (1) σ_V

WronglyChallenged **Transaction**	
Inputs	(0) $(\mathsf{tx}_{\mathsf{ChallengeAssert}}, 0, \mathsf{HashLock}(h_{\mathsf{msg}}) \wedge \mathsf{CheckSig}(\mathsf{pk}_P))$
Outputs	(0) $(0, \langle \mathsf{CheckSig}(\mathsf{pk}_P) \rangle)$
Tx Witness	(0) σ_P, msg

Success **Transaction**	
Inputs	(0) $(\mathsf{tx}_{\mathsf{Deposit}}, 0, \mathsf{CheckSig}(\mathsf{pk}_P) \wedge \mathsf{CheckSig}(\mathsf{pk}_V))$ (1) $(\mathsf{tx}_{\mathsf{Assert}}, 0, \mathsf{RelTimelock}(\Delta_2) \wedge \mathsf{CheckSig}(\mathsf{pk}_P) \wedge \mathsf{CheckSig}(\mathsf{pk}_V))$
Outputs	(0) $(v, \langle \mathsf{CheckSig}(\mathsf{pk}_P) \rangle)$
Tx Witness	(0) σ_P, σ_V
	(1) σ_P, σ_V

B Protocol for Malicious Security

Algorithm 4 `Setup` algorithms for malicious security (Prover)

procedure $P_{\mathtt{Setup},\mathsf{mal}}(\mathsf{crs}, \boldsymbol{x})$ ▷ Run by Prover
  $(\mathsf{sk}_P, \mathsf{pk}_P) \leftarrow \mathsf{Sig}_{BTC}.\mathtt{Gen}(1^\lambda)$ ▷ Sample signing key
  send (pk_P) to Verifier

  Upon receiving $\left(\mathsf{pk}_V, \{h_{\mathsf{msg},i}, \mathsf{epk}_i, h_{\mathsf{ct}_i}\}_{i=1}^{N_{\mathsf{CC}}}\right)$ from Verifier:
  Sample $\mathcal{I}$ as a uniformly random subset of $[N_{\mathsf{CC}}]$ of size M_{CC}
  send $\mathcal{I}$ to Verifier

  Upon receiving $\left(\{\mathsf{msg}_i, r_i, \mathsf{seed}_i\}_{i\in[N_{\mathsf{CC}}]\setminus\mathcal{I}}, \{\mathsf{ct}_{\mathsf{setup}\,i}, \mathsf{ct}_{\mathsf{GC}\,i}\}_{i\in\mathcal{I}}\right)$ from Verifier:
  for $i \in [N_{\mathsf{CC}}] \setminus \mathcal{I}$ **do** ▷ If any verification fails, abort
    $\mathsf{ct}_{\mathsf{setup}\,i} \leftarrow \mathsf{WE}.\mathtt{Enc}_{\mathsf{setup}}(\mathsf{crs}, \boldsymbol{x}, \mathsf{msg}_i, r_i)$ ▷ Constr. 1
    $(\mathsf{ct}_{\mathsf{GC}\,i}, \mathsf{ek}_i) \leftarrow \mathsf{Garble}(r_i; \mathsf{seed}_i)$ ▷ Sec. 5.6
    Verify $h_{\mathsf{msg},i} = \mathsf{Hash}_{BTC}(\mathsf{msg}_i)$
    for $j \in \{1, \ldots, 2n\}, b \in \{0, 1\}$ **do**
      Verify $(\mathsf{epk}_i)_j^b = \mathsf{Hash}_{BTC}((\mathsf{ek}_i)_j^b)$
    end for
  end for
  for $i \in [N_{\mathsf{CC}}]$ **do**
    Verify $h_{\mathsf{ct}_i} = \mathsf{RO}(\mathsf{ct}_{\mathsf{setup}\,i}, \mathsf{ct}_{\mathsf{GC}\,i})$
  end for
  Let $i_1 < \cdots < i_{M_{\mathsf{CC}}}$ be the elements of $\mathcal{I}$ and, for each $k \in [M_{\mathsf{CC}}]$, relabel $(h_{\mathsf{msg},k}, \mathsf{epk}_k, \mathsf{ct}_{\mathsf{setup}\,k}, \mathsf{ct}_{\mathsf{GC}\,k}) := (h_{\mathsf{msg},i_k}, \mathsf{epk}_{i_k}, \mathsf{ct}_{\mathsf{setup}\,i_k}, \mathsf{ct}_{\mathsf{GC}\,i_k})$
  $(\mathsf{lsk}_P, \mathsf{lpk}_P) \leftarrow \mathsf{LampSig}.\mathtt{Gen}(1^\lambda)$ ▷ Sample Lamport key
  $(\mathcal{T}, \mathcal{S}) \leftarrow \mathsf{CreateTxSetMalicious}\left(\mathsf{pk}_P, \mathsf{pk}_V, \mathsf{lpk}_P, \{h_{\mathsf{msg},k}, \mathsf{epk}_k\}_{k=1}^{M_{\mathsf{CC}}}\right)$ ▷ Alg. 7
  $\mathsf{presigs}_P \leftarrow \mathsf{SignTxs}_P(\mathsf{sk}_P, \mathcal{T})$ ▷ Alg. 3
  send $(\mathsf{lpk}_P, \mathsf{presigs}_P)$ to Verifier

  Upon receiving $(\mathsf{presigs}_V)$ from Verifier:
  $\mathsf{VerifySigs}_P(\mathsf{pk}_V, \mathcal{T}, \mathsf{presigs}_V)$ ▷ Alg. 3; if fails, abort
  $\mathsf{st}_P \leftarrow (\mathsf{crs}, \boldsymbol{x}, \mathcal{T}, \mathsf{sk}_P, \mathsf{lsk}_P, \mathsf{presigs}_V, \{\mathsf{ct}_{\mathsf{setup}\,k}, \mathsf{ct}_{\mathsf{GC}\,k}\}_{k=1}^{M_{\mathsf{CC}}})$
  return $(\mathcal{S}, \mathsf{st}_P)$
end procedure

The setup protocol using cut-and-choose, which achieves the BitVM-core properties in Defs. 8 to 10, is shown in Algs. 4 and 5.

The corresponding proving phase is shown in Alg. 6.

Algorithm 5 `Setup` algorithms for malicious security (Verifier)

procedure $V_{\texttt{Setup,mal}}(\mathsf{crs}, \boldsymbol{x})$ ▷ Run by Verifier

Upon receiving (pk_P) from Prover:

$(\mathsf{sk}_V, \mathsf{pk}_V) \leftarrow \mathsf{Sig}_{BTC}.\texttt{Gen}(1^\lambda)$ ▷ Sample signing key

for $i = 1, \ldots, N_{\mathsf{CC}}$ **do** ▷ Generate ciphertexts for cut-and-choose

$\mathsf{msg}_i \leftarrow\!\!\$\ \mathcal{M}$, $r_i \leftarrow\!\!\$\ \mathbb{F}_q^*$ ▷ Sample secrets

$\mathsf{ct}_{\mathsf{setup}\,i} \leftarrow \mathsf{WE}.\texttt{Enc}_{\mathsf{setup}}(\mathsf{crs}, \boldsymbol{x}, \mathsf{msg}_i, r_i)$ ▷ WE ciphertext (Constr. 1)

$\mathsf{seed}_i \leftarrow\!\!\$\ \{0,1\}^\lambda$ ▷ Seed for Garble

$\mathsf{ct}_{\mathsf{GC}\,i}, \mathsf{ek}_i \leftarrow \mathsf{Garble}(r_i; \mathsf{seed}_i)$ ▷ GC ciphertext and encoding key (Sec. 5.6)

$h_{\mathsf{msg},i} \leftarrow \mathsf{Hash}_{BTC}(\mathsf{msg}_i)$ ▷ Hash message for hashlock

for $j \in \{1, \ldots, 2n\}, b \in \{0,1\}$ **do**

$(\mathsf{epk}_i)_j^b \leftarrow \mathsf{Hash}_{BTC}((\mathsf{ek}_i)_j^b)$ ▷ Hash input labels for hashlock

end for

$h_{\mathsf{ct}_i} \leftarrow \mathsf{RO}(\mathsf{ct}_{\mathsf{setup}\,i}, \mathsf{ct}_{\mathsf{GC}\,i})$

end for

send $\left(\mathsf{pk}_V, \{h_{\mathsf{msg},i}, \mathsf{epk}_i, h_{\mathsf{ct}_i}\}_{i=1}^{N_{\mathsf{CC}}}\right)$ to Prover

Upon receiving $\mathcal{I}$ from Prover:

Verify $\mathcal{I} \subseteq [N_{\mathsf{CC}}]$ and $|\mathcal{I}| = M_{\mathsf{CC}}$; if verification fails, abort

send $\left(\{\mathsf{msg}_i, r_i, \mathsf{seed}_i\}_{i \in [N_{\mathsf{CC}}] \setminus \mathcal{I}}, \{\mathsf{ct}_{\mathsf{setup}\,i}, \mathsf{ct}_{\mathsf{GC}\,i}\}_{i \in \mathcal{I}}\right)$ to Prover

Let $i_1 < \cdots < i_{M_{\mathsf{CC}}}$ be the elements of $\mathcal{I}$ and, for each $k \in [M_{\mathsf{CC}}]$, relabel $(h_{\mathsf{msg},k}, \mathsf{epk}_k, \mathsf{ct}_{\mathsf{setup}\,k}, \mathsf{ct}_{\mathsf{GC}\,k}, \mathsf{ek}_k) := (h_{\mathsf{msg},i_k}, \mathsf{epk}_{i_k}, \mathsf{ct}_{\mathsf{setup}\,i_k}, \mathsf{ct}_{\mathsf{GC}\,i_k}, \mathsf{ek}_{i_k})$

Upon receiving $(\mathsf{lpk}_P, \mathsf{presigs}_P)$ from Prover:

$(\mathcal{T}, \mathcal{S}) \leftarrow \mathsf{CreateTxSetMalicious}\left(\mathsf{pk}_P, \mathsf{pk}_V, \mathsf{lpk}_P, \{h_{\mathsf{msg},k}, \mathsf{epk}_k\}_{k=1}^{M_{\mathsf{CC}}}\right)$ ▷ Alg. 7

$\mathsf{VerifySigs}_V(\mathsf{pk}_P, \mathcal{T}, \mathsf{presigs}_P)$ ▷ Alg. 3; if fails, abort

$\mathsf{presigs}_V \leftarrow \mathsf{SignTxs}_V(\mathsf{sk}_V, \mathcal{T})$ ▷ Alg. 3

$\mathsf{st}_V \leftarrow (\mathsf{crs}, \boldsymbol{x}, \mathcal{T}, \mathsf{sk}_V, \{\mathsf{ek}_k\}_{k=1}^{M_{\mathsf{CC}}}, \mathsf{presigs}_P)$

send $(\mathsf{presigs}_V)$ to Prover

return $(\mathcal{S}, \mathsf{st}_V)$

end procedure

Algorithm 6 `Prove` algorithms

procedure $P_{\texttt{Prove,mal}}(\mathsf{st}_P, \boldsymbol{w})$ ▷ Run by Prover
  Parse $\mathsf{st}_P = (\mathsf{crs}, \boldsymbol{x}, \mathcal{T}, \mathsf{sk}_P, \mathsf{lsk}_P, \mathsf{presigs}_V, \{\mathsf{ct}_{\mathsf{setup}_i}, \mathsf{ct}_{\mathsf{GC}\,i}\}_{i=1}^{M_{\mathsf{CC}}})$ and $\mathsf{presigs}_V = (\sigma^V_{\mathsf{Success}})$
  $(\pi_1, \pi_2, \pi_3) \leftarrow \mathsf{Groth16.Prove}(\mathsf{crs}, \boldsymbol{x}, \boldsymbol{w})$
  $w_{\mathsf{Assert}} \leftarrow \mathsf{LampSig.Sign}(\mathsf{lsk}_P, \pi_1)$ ▷ Compute Lamport signature
  Post $\mathsf{tx}_{\mathsf{Assert}}$ with transaction witness w_{Assert}: call $\mathcal{F}_{\mathsf{BTC}}.\mathsf{WRITE}(\mathsf{tx}_{\mathsf{Assert}})$

  Upon seeing $\mathsf{tx}_{\mathsf{Assert}}$ and Δ_2 new blocks after $\mathsf{tx}_{\mathsf{Assert}}$ in $\mathcal{L}_P$:
  $w_{\mathsf{Success}} \leftarrow (\sigma^P_{\mathsf{Success}}, \sigma^V_{\mathsf{Success}})$ where $\sigma^P_{\mathsf{Success}} \leftarrow \mathsf{Sig}_{BTC}.\mathsf{Sign}(\mathsf{sk}_P, \overline{\mathsf{tx}}_{\mathsf{Success}})$
  Post $\mathsf{tx}_{\mathsf{Success}}$ with transaction witness w_{Success}: call $\mathcal{F}_{\mathsf{BTC}}.\mathsf{WRITE}(\mathsf{tx}_{\mathsf{Success}})$
  return 1

  Upon seeing $\mathsf{tx}_{\mathsf{ChallengeAssert}}$ in $\mathcal{L}_P$:
  Extract input labels $\{L_i\}_{i=1}^{M_{\mathsf{CC}}}$ from the transaction witness of $\mathsf{tx}_{\mathsf{ChallengeAssert}}$
  For $i = 1, \ldots, M_{\mathsf{CC}}$: $\mathsf{ct}_{\mathsf{prove}_i} \leftarrow \mathsf{EvalGC}(\mathsf{ct}_{\mathsf{GC}\,i}, L_i)$ ▷ Evaluate garbled circuits (Sec. 5.6)
  For $i = 1, \ldots, M_{\mathsf{CC}}$: $\mathsf{msg}_i \leftarrow \mathsf{WE.Dec}(\mathsf{crs}, (\boldsymbol{x}, \pi_1), (\mathsf{ct}_{\mathsf{prove}_i}, \mathsf{ct}_{\mathsf{setup}_i}), \boldsymbol{w})$ ▷ At least one must succeed for an honest Prover by u-robustness (Def. 8)
  Find i such that $w_{\mathsf{WronglyChallenged}} \leftarrow (\sigma^P_{\mathsf{WronglyChallenged}}, \mathsf{msg}_i)$ where $\sigma^P_{\mathsf{WronglyChallenged}} \leftarrow \mathsf{Sig}_{BTC}.\mathsf{Sign}(\mathsf{sk}_P, \overline{\mathsf{tx}}_{\mathsf{WronglyChallenged}})$ is a valid transaction witness
  Post $\mathsf{tx}_{\mathsf{WronglyChallenged}}$ with transaction witness $w_{\mathsf{WronglyChallenged}}$: call $\mathcal{F}_{\mathsf{BTC}}.\mathsf{WRITE}(\mathsf{tx}_{\mathsf{WronglyChallenged}})$
end procedure

procedure $V_{\texttt{Prove,mal}}(\mathsf{st}_V)$ ▷ Run by Verifier
  Upon seeing $\mathsf{tx}_{\mathsf{Assert}}$ in $\mathcal{L}_V$:
  Parse $\mathsf{st}_V = (\mathsf{crs}, \boldsymbol{x}, \mathcal{T}, \mathsf{sk}_V, \{\mathsf{ek}_i\}_{i=1}^{M_{\mathsf{CC}}}, \mathsf{presigs}_P)$ and $\mathsf{presigs}_P = (\sigma^P_{\mathsf{ChallengeAssert}}, \sigma^P_{\mathsf{Failure}})$
  Extract π_1 and Lamport signature μ from the transaction witness of $\mathsf{tx}_{\mathsf{Assert}}$
  For $i = 1, \ldots, M_{\mathsf{CC}}$: $L_i \leftarrow \mathsf{Encode}(\mathsf{ek}_i, \pi_1)$ ▷ Compute input labels (Sec. 5.6)
  $w_{\mathsf{ChallengeAssert}} \leftarrow (\sigma^P_{\mathsf{ChallengeAssert}}, \sigma^V_{\mathsf{ChallengeAssert}}, \mu, \{L_i\}_{i=1}^{M_{\mathsf{CC}}})$ where $\sigma^V_{\mathsf{ChallengeAssert}} \leftarrow \mathsf{Sig}_{BTC}.\mathsf{Sign}(\mathsf{sk}_V, \overline{\mathsf{tx}}_{\mathsf{ChallengeAssert}})$
  Post $\mathsf{tx}_{\mathsf{ChallengeAssert}}$ with transaction witness $w_{\mathsf{ChallengeAssert}}$: call $\mathcal{F}_{\mathsf{BTC}}.\mathsf{WRITE}(\mathsf{tx}_{\mathsf{ChallengeAssert}})$

  Upon seeing $\mathsf{tx}_{\mathsf{ChallengeAssert}}$ and Δ_1 new blocks after $\mathsf{tx}_{\mathsf{ChallengeAssert}}$ in $\mathcal{L}_V$, if $\mathsf{tx}_{\mathsf{WronglyChallenged}} \notin \mathcal{L}_V$:
  $w_{\mathsf{Failure}} \leftarrow \big((\sigma^P_{\mathsf{Failure}}, \sigma^V_{\mathsf{Failure}}), (\sigma^V_{\mathsf{Failure}})\big)$ where $\sigma^V_{\mathsf{Failure}} \leftarrow \mathsf{Sig}_{BTC}.\mathsf{Sign}(\mathsf{sk}_V, \overline{\mathsf{tx}}_{\mathsf{Failure}})$
  Post $\mathsf{tx}_{\mathsf{Failure}}$ with transaction witness w_{Failure}: call $\mathcal{F}_{\mathsf{BTC}}.\mathsf{WRITE}(\mathsf{tx}_{\mathsf{Failure}})$
end procedure

Algorithm 7 Locking scripts for the protocol with malicious security (Algs. 4 and 5)

function CreateTxSetMalicious($\mathsf{pk}_P, \mathsf{pk}_V, \mathsf{lpk}_P, \{h_{\mathsf{msg},i}, \mathsf{epk}_i\}_{i=1}^{M_{\mathsf{CC}}}$)
  Define scripts:
  $\mathsf{ChallengeAssertScript} := \bigwedge_{j=1}^{2n} \Big[\Big(\bigwedge_{i=1}^{M_{\mathsf{CC}}} \mathsf{HashLock}((\mathsf{epk}_i)_j^0) \wedge \mathsf{HashLock}(\mathsf{lpk}_j^0) \Big)$
    $\vee \Big(\bigwedge_{i=1}^{M_{\mathsf{CC}}} \mathsf{HashLock}((\mathsf{epk}_i)_j^1) \wedge \mathsf{HashLock}(\mathsf{lpk}_j^1) \Big) \Big]$ ▷ Replaces CheckLampSigsMatch in Alg. 3
  $\mathsf{WronglyChallengedScript} := \bigvee_{i=1}^{M_{\mathsf{CC}}} (\mathsf{HashLock}(h_{\mathsf{msg},i}))$ ▷ Replaces $\mathsf{HashLock}(h_{\mathsf{msg}})$ in Alg. 3
  Construct transaction skeletons $\overline{\mathsf{tx}}_{\mathsf{Deposit}}$, $\overline{\mathsf{tx}}_{\mathsf{Assert}}$, $\overline{\mathsf{tx}}_{\mathsf{ChallengeAssert}}$, $\overline{\mathsf{tx}}_{\mathsf{Failure}}$, $\overline{\mathsf{tx}}_{\mathsf{WronglyChallenged}}$, $\overline{\mathsf{tx}}_{\mathsf{Success}}$ as in Sec. 6.1.3
  $\mathcal{T} := \{\overline{\mathsf{tx}}_{\mathsf{Deposit}}, \overline{\mathsf{tx}}_{\mathsf{Assert}}, \overline{\mathsf{tx}}_{\mathsf{ChallengeAssert}}, \overline{\mathsf{tx}}_{\mathsf{Failure}}, \overline{\mathsf{tx}}_{\mathsf{WronglyChallenged}}, \overline{\mathsf{tx}}_{\mathsf{Success}}\}$
  $\mathcal{S} := \{\overline{\mathsf{tx}}_{\mathsf{Success}}\}$
  return $(\mathcal{T}, \mathcal{S})$
end function