



Glimpse: On-Demand PoW Light Client with Constant-Size Storage for DeFi

*Giulia Scaffino, TU Wien and Christian Doppler Laboratory Blockchain Technologies
for the Internet of Things; Lukas Aumayr and Zeta Avarikioti, TU Wien;
Matteo Maffei, TU Wien and Christian Doppler Laboratory Blockchain Technologies
for the Internet of Things*

<https://www.usenix.org/conference/usenixsecurity23/presentation/scaffino>

**This paper is included in the Proceedings of the
32nd USENIX Security Symposium.**

August 9–11, 2023 • Anaheim, CA, USA

978-1-939133-37-3

**Open access to the Proceedings of the
32nd USENIX Security Symposium
is sponsored by USENIX.**

Glimpse: On-Demand PoW Light Client with Constant-Size Storage for DeFi

Giulia Scaffino^{1,2}, Lukas Aumayr¹, Zeta Avarikioti¹, and Matteo Maffei^{1,2}

¹TU Wien, {giulia.scaffino, lukas.aumayr, georgia.avarikioti, matteo.maffei}@tuwien.ac.at

²Christian Doppler Laboratory Blockchain Technologies for the Internet of Things

Abstract

Cross-chain communication is instrumental in unleashing the full potential of blockchain technologies, as it allows users and developers to exploit the unique design features and the profit opportunities of different existing blockchains. The majority of interoperability solutions are provided by centralized exchanges and bridge protocols based on a trusted majority, both introducing undesirable trust assumptions compared to native blockchain assets. Hence, increasing attention has been given to decentralized solutions: Light and super-light clients paved the way for chain relays, which allow verifying on a blockchain the state of another blockchain by respectively verifying and storing a linear and logarithmic amount of data. Unfortunately, relays turn out to be inefficient in terms of computational costs, storage, or compatibility.

We introduce *Glimpse*, an *on-demand bridge* that leverages a novel *on-demand light client* construction with only *constant* on-chain storage, cost, and computational overhead. *Glimpse* is *expressive*, enabling a plethora of DeFi and off-chain applications such as lending, pegs, proofs of oracle attestations, and betting hubs. *Glimpse* also remains *compatible* with blockchains featuring a limited scripting language such as the Liquid Network (a pegged sidechain of Bitcoin), for which we present a concrete instantiation. We prove *Glimpse* security in the Universal Composability (UC) framework and further conduct an economic analysis. We evaluate the cost of *Glimpse* for Bitcoin-like chains: verifying a simple transaction has at most 700 bytes of on-chain overhead, resulting in a *one-time* fee of \$3, only twice as much as a standard Bitcoin transaction.

1 Introduction

The blockchain landscape is fragmented into a plethora of blockchains presenting different technical features (scripting languages, consensus mechanisms, etc.) and attracting users for their unique properties (e.g., Bitcoin for its robust design, Monero and ZCash for their privacy, Ethereum for the support of DeFi applications, Algorand for its high throughput). Blockchain platforms already hold an impressive amount of

investments, users, and developers, who are often reluctant to migrate their assets and contracts to other chains. By providing interoperability solutions, centralized exchanges have enabled an appealing ecosystem of financial applications, such as trading different cryptocurrencies, collateral-based lending, and more. Unfortunately, centralized exchanges need to be trusted; additionally, they can be hacked, go bankrupt, or be fraudulent, in which case the users' money is at risk. The same holds true for solutions assuming a trusted majority of validators: e.g., an attacker managed to acquire five of the nine validation keys used in the Ronin bridge [1], stealing \$624M; attacks on Wormhole (\$320M), Nomad (\$200M) or Harmony (\$100M) and more, totaling over \$1.3B of stolen funds in the first 8 months of 2022 alone [2]. For these reasons, the design of decentralized interoperability solutions is crucial to unleashing the full potential of blockchain technologies.

The challenge of blockchain interoperability (or *cross-chain communication*) stems from the two functionalities it has to provide: (i) relaying information from a source ledger \mathcal{L}_S to a destination ledger \mathcal{L}_D , allowing a user of \mathcal{L}_D to verify that a transaction T_{X_S} has been included in \mathcal{L}_S without participating in \mathcal{L}_S 's consensus protocol (cross-chain verification), and (ii) atomic synchronization of transactions across different chains, e.g., in an atomic swap, a transaction T_{X_D} on \mathcal{L}_D succeeds if and only if T_{X_S} was previously posted on \mathcal{L}_S (cross-chain atomicity).

Related Work. Cross-chain verification is typically achieved by running either full nodes or light clients with linear storage overhead in \mathcal{L}_S 's length. The core idea of *light clients*, illustrated for the first time in Nakamoto's original paper [3] for Simplified Payment Verification (SPV), is to store and verify the block headers alone, as opposed to the whole block, and to verify which chain carries the most Proof-of-Work (PoW). The assumption underlying the security of light clients is that the majority of miners follow the consensus rules; therefore, the chain with the most PoW represents the honest chain. SPV-based light clients save storage (a Bitcoin block header is about 80B in size, whereas a block is about 1MB) but still require the relaying and processing of a *linear* amount of

information in the chain length, with an overhead of 60MB for Bitcoin and 4GB for Ethereum. Light clients and SPVs are the basis of chain relays [4–6], an expressive but expensive solution to the cross-chain verification problem. Chain relays verify and store every block header of \mathcal{L}_S within a smart contract on \mathcal{L}_D , thereby acting as light clients. The inefficiency of this construction, associated with the lack of incentives for relayers and the high maintenance costs, is arguably one of the reasons why relays are not used in practice.¹ Later on, *super-light clients* with logarithmic complexity were proposed (PoPoW [7], FlyClient [8]), but they either require constant PoW difficulty [7] or an hard fork in Bitcoin [8], and are thus not backward compatible.

More recently, Xie et al. [9] have introduced an Ethereum-compatible bridge with constant size storage, where a zk-SNARK proof guarantees that the blockchain has undergone a state update (either a single block or a batch of them). Each verified state update is stored within a stateful contract, and recent block headers of the source chain are relayed to and stored within the contract until a zero-knowledge proof is made available to finalize and verify the state update. Application-specific contracts can then rely on zkBridge to perform, e.g., SPV verifications. However, zkBridge still requires a linear amount of information relayed from the source to the destination chain. As chain relays, zkBridge still incurs in high maintenance costs and lacks of incentives for relayers, which continuously compute zk proofs and submit them to the bridge contract along with block headers. Moreover, zkBridge can only be deployed and used in destination chains supporting a quasi-Turing complete language, thus excluding important chains holding hundreds of millions of dollars.

Similarly to zkBridge, all current implementations based on the aforementioned client solutions require a quasi-Turing complete scripting language on the destination chain and are thus not compatible with blockchains with limited scripting capabilities, such as Bitcoin-based chains. The expressiveness of the scripting language is indeed one of the features setting apart different blockchains, with some (e.g., Ethereum) favoring the support of DeFi applications (albeit some smart contracts can be encoded in the Bitcoin scripting language too [10]) and others (e.g., Bitcoin) more conservatively arguing for a reduced trust base and easier script verification. For this reason, supporting blockchains with limited scripting capabilities is not only theoretically challenging but also a practically relevant research goal.

A different approach to the realization of bridges with constant size storage is *stateless SPV*, initially proposed by Prestwich [11] and implemented by Summa [12]. Stateless SPV also emerged within the Ethereum research community [13]. Recently, Barbára et al. [14] implemented, and for the first time formalized, stateless SPV within the BxTB cross-chain exchange. Instead of verifying all blockchain headers, the

¹For instance, the most popular Bitcoin relay on Ethereum [4] stopped its development in 2017 and the last transaction is from about 4 years ago.

	LC [4–6]	SLC [7,8]	zkBridge [9]	SSPV [10,13]	Glimpse
Information Relayed:	Linear	Logarithmic	Linear	Constant	Constant
Storage Overhead:	Linear	Logarithmic	Constant	Constant	Constant
Backward Compatibility:	Yes	No	Yes	Yes	Yes
\mathcal{L}_D q-Turing completeness:	Yes	Yes	Yes	Yes	No
Upfront Mining Secure:	Yes	Yes	Yes	No	Yes

Table 1: For light clients (LC), super-light clients (SLC), zk-Bridge, stateless SPV (SSPV), and Glimpse we show the amount of information relayed, storage overhead, backward compatibility, need for quasi-Turing complete scripting language on \mathcal{L}_D , and security w.r.t. upfront mining attacks.

idea is to perform a proof of inclusion *on-demand* for a specific transaction: for that, one needs to verify the headers of a sufficiently long subchain whose first block contains the transaction of interest. The authors conduct an economical security analysis, showing that stateless SPV suffices to discourage attacks on the system (i.e., to construct sufficiently long invalid subchains), as it would be economically more profitable to invest the mining power to honestly mine blocks. In this work, we introduce an attack against stateless SPV, which we call *upfront mining attack*: By knowing the to-be-verified transaction upfront, a malicious prover may produce a forged subchain beforehand, leveraging the fact that users on \mathcal{L}_D have no way to ensure that such proof corresponds to the suffix of the correct chain. Since there is no backward time constraint on performing an upfront mining attack, the attacker will eventually succeed in finding enough forged blocks regardless of their mining power and without needing to bribe any miners. This attack is not considered in stateless SPV nor in BxTB, and gives them a strictly weaker security notion than, e.g., SPV-based light clients, where this cannot happen due to the honest majority assumption. A summary comparison of the different clients for cross-chain verification can be found in Table 1.

Cross-chain atomicity, i.e., the second core functionality of interoperability, is typically implemented with *lock contracts*, such as Hashed TimeLock Contracts (HTLCs) and adaptor signatures [15]. These secret-based cryptographic techniques use a statement S that ties the authorization of a transaction T_{XD} on \mathcal{L}_D to the leakage of a secret witness s within a transaction T_{XS} posted on-chain \mathcal{L}_S . Lock contracts, however, have fundamental limitations: (i) they require both parties to monitor and actively participate in both chains (e.g., in the context of atomic swaps), (ii) their expressiveness is very limited: they require all transactions to be fully fixed upfront (and pre-signed by the party giving away the coins) since they must depend on the same secret, and (iii) they require one party choosing the secret at the start of the protocol, while the other party learns it later. As a result, lock contracts cannot be used in applications like lending or Proofs-of-Burn, where *the same party* needs to post transactions on both \mathcal{L}_S and \mathcal{L}_D *without the intervention of the other*. We expand on this in Appendix A, where we further discuss the related work.

To summarize, the present research landscape leaves the following research question open: *“Is it possible to design*

a secure solution for cross-chain verification which guarantees cross-chain atomicity, requires constant size storage, and makes use of limited scripting language on the target chain?"

Our Contributions. In this work, we positively answer the above question by presenting Glimpse, the first secure on-demand bridge that achieves *atomicity* and *constant size storage* by encoding a novel on-demand PoW light client in the scripting language of the destination chain.

To achieve constant-size overhead, our light client *assumes knowledge of the current PoW target* and is acting *on-demand, verifying only selected transactions*. In particular, Glimpse allows a *prover* and a *verifier* to establish a contract enforcing that if a specific set of transactions T_{X_S} are confirmed on a PoW \mathcal{L}_S within a given time after the contract is settled (we call this time frame the *contract lifetime*), then another set of transactions T_{X_D} can be published on \mathcal{L}_D . Technically, Glimpse is a contract living on \mathcal{L}_D , which receives from the prover a *proof* that T_{X_S} was included on \mathcal{L}_S with the desired number of confirmations, and enables T_{X_D} to appear on \mathcal{L}_D . Glimpse reconciles the *low on-chain costs and simple design* of lock contracts with the *expressiveness* of chain relays.

Glimpse builds on the notion of stateless SPV, but it refines and generalizes it in a number of ways. First, we propose a generic technique to *prevent upfront mining*, imposing prover and verifier to agree on a random value to be inserted in the to-be-verified transaction. Second, we generalize stateless SPV to *support applications in which part of the transaction T_{X_S} is not known a priori* but is instead determined at run-time (e.g., in lending, where for the loan payback transaction, the input it is not known beforehand, as the lent money can be used arbitrarily), as well as *applications requiring the synchronization of combinations of transactions on \mathcal{L}_S* . In particular, Glimpse allows to encode that if any set of transactions satisfying a logical formula expressed as Disjunctive Normal Form² (DNF), e.g., $T_{X_S} \vee T_{X_S}'$, is published on \mathcal{L}_S , then T_{X_D} can be published on \mathcal{L}_D . These generalizations allow us to encode a variety of DeFi applications, such as lending, pegs, wrapping/unwrapping of tokens, Proofs-of-Burn, verification of multiple oracle attestations, and layer-2 applications such as cross-chain virtual channels, payments, and betting hubs.

Third, we provide a construction that, for the first time, does not require quasi-Turing complete scripting languages on the destination chain, thus supporting Bitcoin-based blockchains such as the Liquid Network.³ The new DeFi protocols for the Liquid Network enabled by Glimpse are de-facto brought into Bitcoin, by pegged conversion of BTC into L-BTC tokens. We further show that only two opcodes are missing to directly support Glimpse on Bitcoin as a destination chain,

²A disjunctive normal form formula is a logical formula consisting of a disjunction of conjunctions; it can also be described as an OR of ANDs.

³The Liquid Network [16] is a Bitcoin sidechain supported by Blockstream [17] and other major Bitcoin stakeholders, that has anticipated all major upgrades in Bitcoin (SegWit [18], Taproot [19]). The Liquid Network plays a key role in the Bitcoin ecosystem, e.g., El Salvador’s Bitcoin bonds are Liquid Network security tokens [20].

$\mathcal{L}_S \backslash \mathcal{L}_D$	Bitcoin	Bitcoin Cash/SV	Litecoin	Liquid	Ethereum/EVM-chains
Bitcoin	-	$\times^{\dagger\dagger}$	\times^{\dagger}	✓	✓
Bitcoin Cash/SV	\times^{\dagger}	-	\times^{\dagger}	✓	✓
Litecoin	$\times^{\dagger,*}$	$\times^{\dagger\dagger,*}$	-	\times^*	✓
Ethereum PoW	$\times^{\dagger,*}$	$\times^{\dagger\dagger,*}$	$\times^{\dagger,*}$	\times^*	✓

Table 2: Popular Bitcoin-based and EVM-based Glimpse-compatible source (\mathcal{L}_S) and destination (\mathcal{L}_D) chains. \dagger : lack of string opcodes. $\dagger\dagger$: lack of Taproot. $*$: lack of crypto opcodes.

which adds a further motivation for their inclusion in the ongoing discussion within the community. Notably, Glimpse has full compatibility with Bitcoin as source chain, thus enabling, for the first time, Bitcoin-to-Ethereum cross-chain communication with constant storage, constant amount of relayed information, and no maintenance costs.

Table 2 provides a non-exhaustive list of popular chains for which we show the current compatibility when functioning as \mathcal{L}_S or \mathcal{L}_D for Glimpse.

Our further contributions are summarized below:

- We demonstrate the expressiveness of Glimpse by encoding a variety of DeFi and off-chain applications (Section 4).
- We formally analyze Glimpse in the UC framework, where we prove its atomicity properties (Section 5).
- We conduct an economic security analysis to quantify the costs of forgery attacks (affecting any light client) and censorship attacks (harming any timelock-based protocol). Specifically, Glimpse is secure against proof forgeries as long as the value *simultaneously locked on all valid Glimpse contracts* does not exceed a certain threshold (for concrete numbers, \$230M), which is comparable to the total value currently locked on popular bridges. To enhance security against censorship attacks on the destination chain, we further impose an upper bound on the value held by *each single Glimpse contract*, e.g., \$1.1M for Glimpse deployed on Ethereum (Section 6).
- We demonstrate the practicality of Glimpse by evaluating its on-chain costs in Ethereum- and Bitcoin-like chains, showing that, e.g., in Bitcoin the overall cost is at most \$3, around twice as much as ordinary transactions. We also further optimize it with Taproot [21, 22] (Section 7).

2 Background

The UTXO Transaction Model. Each user U is identified by a pair of digital keys (pk_U, sk_U) that are used to prove ownership over coins. A transaction $T_x = (\text{cntr}_{in}, \text{inputs}, \text{cntr}_{out}, \text{outputs}, \text{witnesses})$ is an atomic update of the blockchain state and is associated to a unique identifier $\text{txid} \in \{0, 1\}^{256}$ defined as the *hash $\mathcal{H}([T_x])$ of the transaction*, where $[T_x] := (\text{cntr}_{in}, \text{inputs}, \text{cntr}_{out}, \text{outputs})$ is the *body of the transaction*. Intuitively, a transaction maps a non-empty list of inputs to a non-empty list of newly created outputs, describing a redistribution of funds from the users identified in the inputs to those identified in the outputs.

$\text{cntr}_{in}, \text{cntr}_{out} \in \mathbb{N}_{>0}$ represent the number of elements in the inputs and outputs lists. Any input ζ in the list of inputs is an unspent output from an older transaction, defined by the tuple $\zeta := (\text{txid}, \text{outid})$, with $\text{txid} \in \{0, 1\}^{256}$ representing the hash of the old transaction containing the to-be-spent output, and $\text{outid} \in \mathbb{R}_{\geq 0}$ the index of such an output within the output list of the old transaction. These two fields uniquely identify the to-be-spent output. *witnesses* $\in \{0, 1\}^*$, also known as *scriptSig* or *unlocking script*, is a list of witnesses ω , i.e., the data that only the entity entitled to spend the output can provide, thereby authenticating and validating the transaction. Any output θ in the list of outputs is a pair $\theta := (\text{coins}, \phi)$ and can be consumed by at most one transaction (i.e., no double-spend). The amount of coins in an output θ is denoted by $\text{coins} \in \mathbb{R}_{>0}$, whereas the spendability of θ is restricted by the conditions in ϕ , also known as the *scriptPubKey* or *locking script*. Such conditions are modeled in the native scripting language of the blockchain and can vary from single-user $\text{OneSig}(\text{pk}_U)$ and multi-user $\text{MuSig}(\text{pk}_{U1}, \text{pk}_{U2})$ ownership, to time locks, hash locks, and more complex scripts.

Proof-of-Work Consensus. In a PoW blockchain, the probability that a node is selected as block proposer is proportional to its computational power. This is meant to hinder Sybil attacks since computational power is assumed hard to monopolize. Specifically, incentivized to win the reward in native assets, the nodes compete with each other to create, validate, and append new blocks to the ledger by solving a cryptographic puzzle that is hard to compute and easy to verify. The content of a block is summarized within a unique and cryptographically secured string that grants immutability to the blockchain: the *block header* $\text{header}(B) := (\text{ParentHash}, \text{MR}, \text{Timestamp}, \text{nBits}, \text{Nonce})$, where *ParentHash* is the hash of the previous block, *MR* is the root of the Merkle tree whose leaves are the transactions in *B*, *Timestamp* is the creation time of the block, *nBits* is a parameter for the target space, and *Nonce* a value that can be arbitrarily iterated to reach the PoW.

In particular, the nodes, called *miners*, repeatedly change the *Nonce* field of the block header until the hash of the header lies within a *target space* that is smaller (by several orders of magnitude) than the output space of the hash function. This is a necessary condition for the block to be *valid*. The size of the target space is parameterized by the total computational power of the network and is periodically adjusted to keep the expected *block time*, i.e., the time it takes to find a valid block, almost constant. We refer to the *target* as \mathcal{T} , and we say that a block *B* is valid when $\mathcal{H}(\text{header}(B)) < \mathcal{T}$. A miner is selected to propose the next block with probability proportional to the fraction of the network’s hashing power he controls. PoW blockchains periodically adjust the network difficulty to maintain an (almost) constant average block creation time, preventing uncontrolled inflation and network congestion.

3 Glimpse

We introduce Glimpse, a new *primitive for cross-chain communication* that allows participants to obtain *on-demand* the desired information about the state of a PoW source ledger \mathcal{L}_S on a destination ledger \mathcal{L}_D . Glimpse achieves this with only a *constant amount of data* (with respect to the source chain’s length), and assuming the *PoW target is known*.

In particular, Glimpse resembles challenge-response protocols: a *verifier* V challenges a *prover* P to prove the inclusion on \mathcal{L}_S of a *specific* set of transactions T_{X_S} . Depending on the outcome of the challenge, P and V want to publish on \mathcal{L}_D different pre-selected sets of transactions (T_{X_P} or T_{X_V}). To encode this, P and V first agree on the Glimpse specifics and on some consensus parameters of \mathcal{L}_S , then they deploy a *Glimpse contract* on \mathcal{L}_D . On ledger \mathcal{L}_S , an *issuer* I publishes the transaction set T_{X_S} , and an *untrusted relayer* R provides P with the necessary data to construct a *proof* \mathcal{P} , proving the occurrence of T_{X_S} on \mathcal{L}_S . If P submits a valid proof (response) to the Glimpse contract on \mathcal{L}_D , he can post T_{X_P} on \mathcal{L}_D . Else, V can post T_{X_V} after time T has elapsed.

3.1 Assumptions and Models

System Model. We assume a source ledger \mathcal{L}_S operating a PoW consensus. Glimpse relies on four parties: an issuer I that publishes T_{X_S} on \mathcal{L}_S , a prover P that proves the inclusion of T_{X_S} on \mathcal{L}_D , a relayer R (e.g., blockchain explorers, full nodes) that provides P with the necessary information to construct the proof \mathcal{P} , and a verifier V that guarantees contractual fairness. Depending on the application, parties can play several of these roles at once. E.g., in lending, Proofs-of-Burn, and backed assets (see Section 4.1), P can also play the roles of I and R , being incentivized to get reliable insights on \mathcal{L}_S ’s state.

We require P and V to each have a key pair (sk, pk) on \mathcal{L}_D , and I to have a key pair on \mathcal{L}_S . The Glimpse contract is deployed on \mathcal{L}_D and holds coins either coming from P , V , or from any other user of \mathcal{L}_D (this is application-specific). We assume \mathcal{L}_D to support the same hash function used by the consensus of \mathcal{L}_S , and both \mathcal{L}_S and \mathcal{L}_D to allocate the same domain for the hash function, to avoid oversize preimage attacks [23, 24]. Finally, \mathcal{L}_D needs to support the following functionalities: (i) Merkle proof verification, (ii) hash comparison, and (iii) block header and transaction body reconstruction. While (ii) is supported by default in most chains, (i) and (iii) can also be supported by Bitcoin-based chains by enabling a concatenation opcode (recently discussed within the Bitcoin community in the context of Speedy Covenants [25]).

Cryptographic Assumptions. We consider hash functions modeled as random oracles and digital signature schemes having Existential Unforgeability under Chosen Message Attack (EUF-CMA) security.

Communication Model. We assume there exist authenticated communication channels between the Glimpse parties, where all messages are delivered within a fixed time delay.

Cross-chain Communication (CCC) Model. Closely following [23], CCC protocols are usually articulated in three main phases: *Setup*, *Commit on \mathcal{L}_S* , and *Verify & Commit on \mathcal{L}_D* . The *Setup* phase parameterizes the involved blockchains, identifies the protocol participants, defines the timeline for the CCC protocol execution (if any), and specifies the transactions T_{X_S} and T_{X_D} to be synchronized. After a successful setup, in the *Commit on \mathcal{L}_S* phase, a publicly verifiable commitment to execute the CCC protocol, i.e., T_{X_S} , is posted on \mathcal{L}_S . In the *Verify & Commit on \mathcal{L}_D* phase, the commitment published on \mathcal{L}_S is relayed to \mathcal{L}_D , it is verified, and, upon successful verification, a publicly verifiable commitment, i.e., T_{X_D} , is posted on \mathcal{L}_D . An optional *Abort* phase reverts transaction T_{X_S} on \mathcal{L}_S in case the verification of the commitment failed or the commitment on \mathcal{L}_D is not posted.

A CCC protocol has to provide atomicity guarantees, which, for Glimpse, we articulate in a weak and a strong variant. *Weak atomicity* ensures that T_{X_D} appears on \mathcal{L}_D only if T_{X_S} has been already confirmed on \mathcal{L}_S . *Strong atomicity* ensures that T_{X_D} appears on \mathcal{L}_D if and only if T_{X_S} has been already confirmed on \mathcal{L}_S .

Let $\Delta_D \in \mathbb{N}$ be the *wait time parameter* of \mathcal{L}_D , i.e., the upper bound of the time it takes for a valid transactions to be included on \mathcal{L}_D , and let T_{X_S} and T_{X_D} be (sets of) transactions for \mathcal{L}_S and \mathcal{L}_D , respectively. We refer to n as the number of confirmation blocks that need to be mined on top of a block containing a transaction T_x , for T_x to be considered *stable* [26] on a PoW ledger.

Definition 1 (Weak Atomicity). A valid T_{X_D} is reported stable by honest players on \mathcal{L}_D at time t only if honest players of \mathcal{L}_S have reported T_{X_S} with at least n confirmations at time $t' \leq t - \Delta_D$: $T_{X_D} \in \mathcal{L}_D \implies T_{X_S} \in \mathcal{L}_S$.

Definition 2 (Strong Atomicity). A valid T_{X_D} is reported stable by honest players on \mathcal{L}_D at time t if and only if honest players of \mathcal{L}_S have reported T_{X_S} with at least n confirmations at time $t' \leq t - \Delta_D$: $T_{X_D} \in \mathcal{L}_D \iff T_{X_S} \in \mathcal{L}_S$. If either T_{X_S} or T_{X_D} is invalid and provided to honest players, then neither T_{X_S} nor T_{X_D} is reported stable on \mathcal{L}_S and \mathcal{L}_D , respectively.

Adversarial Model. I , R , P , and V are *mutually distrustful* parties, with *at least one between P and V being honest*. We let γ be the fraction of honest miners the blockchain can tolerate, where the specific value for γ depends on the underlying consensus. For this model, we formally prove in the UC framework that the Glimpse protocol UC-realizes an ideal functionality $\mathcal{F}_{W-Glimpse}$, and we show that weak atomicity holds (Section 5).

Then, we extend our model to incorporate rational participants or, in other words, participants exhibiting liveness during the whole duration of the protocol (as defined in the extended version of this paper [27], Appendix B.2): in this setting, assuming P has direct access to \mathcal{L}_S , we show that the Glimpse protocol UC-realizes an ideal functionality

$\mathcal{F}_{S-Glimpse}$, and strong atomicity holds (Section 5). Finally, in Section 6 we provide an economic analysis incorporating rational adversaries.

3.2 Protocol Overview

In the *Setup* phase, P and V cooperate in the creation of the Glimpse contract T_{X_G} , which hardcodes the PoW difficulty target \mathcal{T}_S of \mathcal{L}_S (*consensus parameter*), as well as the following *Glimpse specifics*: the hash of the to-be-verified transaction T_{X_S} , the contract lifetime T , the number n of confirmation blocks in the proof, and the funds' spending conditions.

P and V prepare transactions T_{X_P} and T_{X_V} , both spending the contract's funds, but in different ways; these two transactions are meant to be published by P and V respectively, and are commitments to how the coins must be distributed in case P provides a valid proof as a witness for T_{X_P} , or V reacts to the lack of such proof by publishing T_{X_V} after time T . P signs T_{X_V} and sends the signature to V , whereas V signs T_{X_P} and gives the signature to P . They exchange signatures over T_{X_G} and publish T_{X_G} on \mathcal{L}_D (if any additional party contributes funds to T_{X_G} , their signature over T_{X_G} is also required in order to publish T_{X_G}). Finally, they hand in T_{X_S} to I .

In the *Commit on \mathcal{L}_S* phase, I publishes T_{X_S} on \mathcal{L}_S .

In the *Verify & Commit on \mathcal{L}_D* phase, P queries R about the inclusion of T_{X_S} on \mathcal{L}_S asking for the necessary data to construct the proof \mathcal{P}^n (we describe how \mathcal{P}^n is constructed in Section 3.3). P publishes T_{X_P} on \mathcal{L}_D by using \mathcal{P}^n , V 's signature over T_{X_P} , and their own signature as witnesses. After time T , if the funds in T_{X_G} are still unspent, V publishes T_{X_V} on \mathcal{L}_D with P 's signature over T_{X_V} as well as their own as witnesses. The Glimpse instance is now closed and the funds are distributed as agreed in the *Setup* phase.

In Section 3.4, we will explore how this approach can be extended to generalize T_{X_S} as a set of transactions represented by a disjunctive normal form formula. Figure 1 depicts Glimpse protocol flow.

3.3 Designing the Proof

We show how the Glimpse proof \mathcal{P}^n is constructed by considering stateless SPV as a preliminary proposal. We identify its vulnerabilities to upfront mining attacks and gradually enhance the construction to attain the desired security properties, while also enhancing expressiveness and compatibility.

Stateless SPV. In stateless SPV [13, 14], users convince with a proof \mathcal{P}^n a quasi-Turing complete smart contract hosted on a destination chain \mathcal{L}_D that a transaction T_x has appeared on a PoW source chain \mathcal{L}_S . The proof \mathcal{P}^n consists of the header of the block containing T_x , the Merkle inclusion proof for the transaction within such block, and n subsequent confirmation block headers. The smart contract verifies the Merkle proof, checks that each of the $n + 1$ headers is a valid child of its parent, and ensures that all headers hold enough PoW, i.e., their hashes are smaller than the pre-defined target.

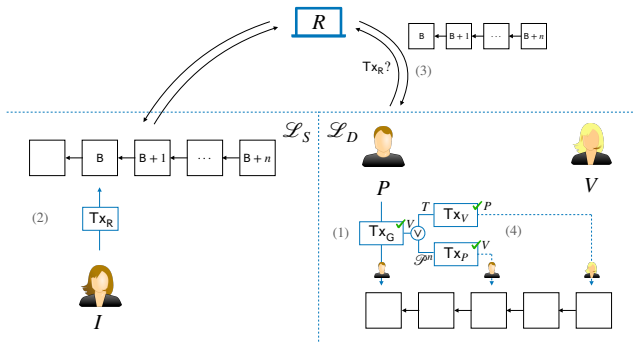


Figure 1: (1) Upon agreeing on the Glimpse specifics, P and V construct T_{X_G} , T_{X_P} , and T_{X_V} . P publishes T_{X_G} on \mathcal{L}_D . (2) I publishes T_{X_R} on \mathcal{L}_S . (3) P queries R to get the data to construct the proof \mathcal{P}^n for T_{X_R} . (4) P publishes T_{X_P} with \mathcal{P}^n as part of the witness. Else, if after T the funds in T_{X_G} are unspent, V publishes T_{X_V} .

Various proposals exist for stateless SPV, each with its specific requirements: some demand the contract to check certain fields of the to-be-verified transaction, to guarantee the transaction’s intended behavior; others only require a timelock T to limit the time window for the submission of the proof. We show that both requirements are necessary: the first to ensure the transaction is well formed and is indeed the one the parties have agreed upon during the *Setup*, the second to prevent a late submission of the proof (which would break security) and to avoid hostage situations where the funds in the contract are locked indefinitely. We observe that the number n of confirmation blocks in the proof must increase linearly with the timelock T , leading to short lived contracts for practicality.

Upfront Mining Attack. Current stateless SPV designs impose no restriction on T_x , exposing the construction to security risks. In particular, let us consider a malicious P (attacker) wishing to convince the smart contract on \mathcal{L}_D that T_x has been included on \mathcal{L}_S . If the attacker knows T_x beforehand, e.g., well before setting up the contract, they could start mining in advance in order to forge a proof. To illustrate this problem, we consider Bitcoin as source chain: Bitcoin is extended by approximately 53k blocks yearly. An attacker with 0.05% of the mining power of the network is expected to find around 6 blocks in less than three months. This means that the attacker can start forging such 6-block proof upfront and proceed by setting up the stateless SPV contract only when the forged proof is ready, e.g., three months later. In this case, the attacker foregoes any potential reward from honestly mining on the main chain, but can claim all the money the contract holds with certainty, regardless of their mining power, and without the need to bribe miners to forge the proof. Even worse, the attacker could set up multiple stateless SPV contracts with different users based on the same transaction T_x (e.g., a betting application) and use the same, upfront-mined fake proof

in all contracts, thereby cheating multiple users out of their funds at once.

Randomized T_x . Glimpse prevents P from launching an upfront mining attack by asking V to *randomize the transaction* T_x . Specifically, in the *Setup* phase, V samples a uniformly random string $r \xleftarrow{\$} \{0, 1\}^\lambda$ (λ is the security parameter) and plugs it into the body of T_x , producing T_{X_R} . This can be done, e.g., by adding an output of value 0 with spending condition $OP_RETURN(r)$.⁴ The hash of the randomized transaction $\mathcal{H}([T_{X_R}])$ is then hardcoded within T_{X_G} . Being unable to anticipate r , P cannot start forging \mathcal{P}^n upfront: their computational effort is now restricted to the time window T . Additionally, the random value serves as a *unique identifier*, preventing proof replay attacks.

We highlight that randomizing the transaction does not impose any trust assumption, and the verification remains constant-sized, achieving our design goals. We observe that security against upfront mining attacks restricts Glimpse to verifying transactions that will be included on \mathcal{L}_S “in the future”, i.e., only after the contract T_{X_G} has been set up. Past transactions cannot be randomized anymore and are thus vulnerable to upfront mining. This is a fundamental difference between Glimpse and light clients, which can verify any past transaction instead.

Improving Compatibility. Besides protecting from upfront mining attacks, we show that the Glimpse construction also forgoes the need for stateful smart contracts, opposed to [13, 14] which is designed for quasi-Turing complete contracts. Due to its simplicity, Glimpse can be deployed not only on quasi-Turing complete chains (e.g., Ethereum) but also on *Bitcoin-based chains* such as the Liquid Network. This is achieved by hardcoding the transaction fields and the verification logic in a transaction locking script, which is spendable by using signatures and \mathcal{P}^n as witness. We expand on this in Section 3.5.

3.4 Enhancing Expressiveness

Glimpse cannot yet encode sophisticated applications due to two shortcomings in expressiveness: First, inputs and outputs of T_{X_R} must be entirely known a priori, which prevents from using it in, e.g., cross-chain lending applications. Second, only single transaction verification is supported, prohibiting verification of, e.g., attestations from multiple oracles. To cater to such use cases, we augment Glimpse to verify (i) parameterized transactions, i.e., transactions which are *not fully known* during the initial *Setup* phase and (ii) *arbitrary combinations of transactions* on \mathcal{L}_S expressed as *disjunctive normal form* formulas.

Parameterized T_x . The core idea is that we encode in the *Setup* phase the abstract expression of a transaction’s spending condition (e.g., a signature) and not

⁴ OP_RETURN is a Bitcoin script opcode that marks a transaction output as unspendable and can be used to embed up to 80 bytes in a transaction.

the exact parameters (e.g., “whose” signature). From Section 2, recall our definition of transaction body $[Tx] := (cntr_{in}, inputs, cntr_{out}, outputs)$, where inputs and outputs are tuples $(txid, outid)$ and $(coins, \phi)$, respectively. Now, in Figure 2, we introduce the definition of *description* Desc of a transaction, which allows for parameterized inputs and outputs. Concretely, following Figure 2, txid, outid, and coins can either be static values or parameters x_i acting as *placeholders*. Similarly, to avoid fixing a priori specific parameters for the locking script, we say that ϕ can be a function f which encodes a family of scripts: f takes a fixed number of arguments (or parameters) z_i for a *known spending condition logic* and returns the desired parameterized locking script for the to-be-verified transaction. In other words, the parameterized locking script can be filled with concrete values, e.g., public key, script hash, *after* the *Setup* phase. The spending condition logic that f encodes must be already defined in the *Setup* phase: For instance, if the parties agree on $f(z)$ encoding any P2PKH (Pay-To-Public-Key-Hash), then the output of $f(z)$ is a P2PKH with a placeholder z in the place of the public key hash, and only after the *Setup* phase it accepts any public key hash as replacement for z . Further following Figure 2, inputs and outputs are lists of inputs and outputs as defined above, and $cntr_{in}$ and $cntr_{out}$ are the number of overall inputs and outputs, respectively. While the former can be parameterized, the latter must be known from the beginning to avoid miners interpreting transactions in an unintended way: we point to the extended version of this work [27] for a detailed discussion on this.

In the *Setup* phase, the parties now hardcode within the contract Tx_G the description Desc, the target \mathcal{T}_S , the lifetime T , and the proof size n . By replacing $\mathcal{H}([Tx_R])$ with Desc, Glimpse can now verify any Tx_R whose body has the same static data in Desc and any arbitrary realization (specified in a later point in time within the proof) of the parameterized ones. In other words, Desc defines the set of possible transactions Glimpse can accept, yet only one of them can be verified. For example, a Glimpse instance with a parameterized input in Desc, can accept and verify transactions with any input (i.e., any value for txid and outid), but with all other fields matching the ones specified in Desc. For security reasons, the random string sampled by V must always be included in Desc. With $[Tx_R] \leftarrow Desc$, we denote a transaction $[Tx_R]$ *compliant with a description* Desc. Given \mathcal{P}^n and the hardcoded Desc, the full transaction body can be reconstructed and hashed by the script of Tx_G .

Verification of DNF Formulas. Glimpse can efficiently verify any DNF formula \mathcal{F}_S over any k literals L_i , which, in our case, are either transactions or descriptions (see Figure 2). To accomplish this, instead of a single Tx_P , P and V create as many sets of transactions (Tx_T, Tx_F, Tx_P) as the number of conjunctive terms in the formula - these sets are kept off-chain. When I publishes on \mathcal{L}_S a combination of transactions specified by \mathcal{F}_S , P queries R , constructs the corresponding

txid	:=	$\{0, 1\}^{256} x_1$,
outid	:=	$\{0, 1\}^{32} x_2$
coins	:=	$\{0, 1\}^{64} x_3$
ϕ	:=	$f(z_1, \dots, z_n)$
inputs	:=	$[(txid, outid) inputs \cup [(txid, outid)]$
outputs	:=	$[(coins, \phi) outputs \cup [(coins, \phi)]$
$cntr_{in}, cntr_{out}$:=	$\{0, 1\}^{1-9}$
Desc	:=	$(cntr_{in}, inputs, cntr_{out}, outputs)$
L_i	:=	$Desc_i \neg Desc_i$
\mathcal{F}_S	:=	$(L_1 \wedge \dots \wedge L_k) \vee \dots \vee (L_1 \wedge \dots \wedge L_k)$
$\forall(x_1, \dots, x_3, z_1, \dots, z_n). (\mathcal{F}_S \iff Tx_D)$		

Figure 2: Enhanced expressiveness for the verification of parameterized transactions and DNF formulas.

proofs, and posts on \mathcal{L}_D the corresponding set $Tx_D := (Tx_T, Tx_F, Tx_P)$ of transactions. If P cheats by publishing an invalid set, i.e., falsely claiming a transaction was not published on \mathcal{L}_S , V can query R , disprove P , and publish Tx_V . We expand on this in Appendix B.

3.5 Compatibility

As anticipated in Table 2, PoW chains such as Bitcoin, Litecoin, Bitcoin Cash, Bitcoin SV, Ethereum PoW, etc., can be used as the source chain \mathcal{L}_S for Glimpse. As for which chains are supported as destination chain \mathcal{L}_D , we need to make some distinctions, because the more restrictive is the scripting language of \mathcal{L}_D , the fewer \mathcal{L}_S may be compatible.

In particular, Glimpse requires \mathcal{L}_D to support the hash function used in the PoW consensus of \mathcal{L}_S . This strict requirement already rules out some combinations, see the lack of cryptographic primitives in Table 2. For instance, Bitcoin-based chains do not have opcodes for computing Keccak and Scrypt hash functions which are used in Ethereum PoW and Litecoin, respectively. As a result, Ethereum PoW and Litecoin cannot act as source chains for Glimpse contracts deployed on Bitcoin-based chains. On the other hand, EVM-based chains (e.g., Ethereum, Polygon, Binance Smart Chain, Avalanche) as well as chains allowing for a quasi-Turing complete scripting language, can always act as destination chains for Glimpse, no matter what the selected source chain is.

Now, we discuss how the particular design of Glimpse makes it compatible with the Liquid Network, the Bitcoin sidechain. For extended discussion and examples, we refer to the full version of this paper [27], Appendix D.

Liquid Network (Full Compatibility). The Liquid Network inherits its design from Bitcoin, while offering an enriched scripting language. Specifically, certain opcodes essential to Glimpse are disabled in Bitcoin but enabled on the Liquid Network. These include: (i) string concatenation (`OP_CAT`), which can be used for Merkle proof verification, block header reconstruction, and transaction body reconstruction, and (ii) `OP_SUBSTR`, which allows splitting strings. This is necessary for comparing hashes (i.e., compare a block header hash to

(V) can spend. Then, P signs $[T_{xV}]$ producing $\sigma_P([T_{xV}])$ and sends to V the Glimpse specifics: $(Desc, T_1, \mathcal{T}_S, n, \alpha, \text{scriptG}, [T_{xG}], [T_{xP}], [T_{xV}], \sigma_P([T_{xV}]))$.

Upon receiving the message from P , V checks the correctness and well-formedness of the Glimpse specifics and checks if $\sigma_P([T_{xV}])$ is a valid signature. Upon successful verification, V signs T_{xP} and sends the signature to P . Upon receiving $\sigma_V([T_{xP}])$ from V , P checks the validity of the signature and, if valid, P signs $[T_{xG}]$ and publishes T_{xG} on \mathcal{L}_D with witness $\omega = \sigma_P([T_{xG}])$.

Commit on \mathcal{L}_S . When P wants to pay back the loan on \mathcal{L}_S , P posts $T_{xPayback}$ such that $[T_{xPayback}] \leftrightarrow Desc$, meaning that $[T_{xPayback}]$ is equal to $Desc$ apart from x , which in $[T_{xPayback}]$ is replaced by an arbitrary input controlled by P .

Verify & Commit on \mathcal{L}_D . P monitors \mathcal{L}_S checking for $T_{xPayback}$ inclusion with at least n confirmations. P constructs \mathcal{P}^n by (i) taking the concrete realization x^R of the parameter x , (ii) retrieving the header of the block B including $T_{xPayback}$, (iii) constructing the Merkle proof (MP) of $T_{xPayback}$ inclusion in B , and (iv) fetching the first n confirmation block headers of B . Formally, $\mathcal{P}^n := (x^R, MP, \text{header}(B), \text{confHeaders}_n)$, as shown in the pseudocode of Figure 4. P signs $[T_{xP}]$ and gets back their α coins by publishing T_{xP} on \mathcal{L}_D with witness $\omega = (\mathcal{P}^n, \sigma_P([T_{xP}]), \sigma_V([T_{xP}]))$.

After T_1 , if the output of T_{xG} is still unspent, V signs $[T_{xV}]$ and publishes T_{xV} with witness $\omega = (\sigma_P([T_{xV}]), \sigma_V([T_{xV}]))$, claiming the α coins in Glimpse. It is crucial for V to publish T_{xV} right after time T_1 , otherwise P could maliciously claim the funds by publishing T_{xP} after T_1 . On the contrary, T_1 prevents T_{xV} from being valid before T_1 .

To ease readability, so far, we have omitted the loan interest rate, which can be easily taken into account in the money distribution of T_{xP} ; for instance, in Figure 4, one can set $\text{outcomeP} = 0.95$ (thereby leaving the 5% of interest to V).

Liquidation. If the asset price on \mathcal{L}_S drops below a predefined liquidity threshold, V must be able to claim P 's collateral before T_1 . For this, we assume there exists a trusted oracle O on \mathcal{L}_D that regularly publishes a transaction T_{xO} with the real-time price of assets on \mathcal{L}_S ; for instance, O can be a Discreet Log Contract-based [28] or a voting-based [29] oracle. If O is not trusted, we can consider a set of k independent oracles, with the promise that if a large enough number of oracles agree on the same price, the liquidation is granted by verifying a DNF formula over the oracle transactions' descriptions. Oracles do not need to cooperate, nor have a common transaction structure. For simplicity, we discuss the case of a single trusted O whose T_{xO} is described by, e.g., $Desc_O := (1, [\zeta_i], 2, [\theta_r, (0, \text{OP_RETURN}(\text{real-time-price}))])$.

We note that $\theta_r := (0, \text{OP_RETURN}(r))$ includes the randomness, which now must be taken from \mathcal{L}_D itself, so that Glimpse participants can (only for a short time window!) anticipate it and include it in $Desc_O$: for example, r can be the hash of

the transaction (or block) including the last price update published by the oracle. It is V 's responsibility to ensure $Desc_O$ embeds the most recent random string.

Setup($Desc, T_1, \mathcal{T}_S, n$):

1. P sends $Desc := (1, [(x)], 1, [(\alpha', \text{OneSig}(\text{pk}_V))])$ to V .
2. Received $Desc$, V samples a random $r \xleftarrow{\$} \{0, 1\}^\lambda$ and sends $Desc := (1, [(x)], 2, [(\alpha', \text{OneSig}(\text{pk}_V)), (0, \text{OP_RETURN}(r))])$ to P .
3. Let $\alpha := \theta_P \cdot \text{coins}$ and let ζ_P point to an unspent output of P .
4. Let $[T_{xG}] := (1, [\zeta_P], 1, [\theta_G := (\alpha, \text{scriptG}(Desc, T_1, \mathcal{T}_S, n, (P, V)))])$.
5. Let ζ_G point to θ_G . Let $[T_{xP}] := (1, [\zeta_G], 2, [(\alpha \cdot \text{outcomeP}, \text{OneSig}(\text{pk}_P)), (\alpha \cdot (1 - \text{outcomeP}), \text{OneSig}(\text{pk}_V))])$ and $[T_{xV}] := (1, [\zeta_G], 1, [(\alpha, \text{OneSig}(\text{pk}_V))])$.
6. P computes $\sigma_P([T_{xV}])$ and sends $(Desc, T_1, \mathcal{T}_S, n, \alpha, \text{scriptG}, [T_{xG}], [T_{xP}], [T_{xV}], \sigma_P([T_{xV}]))$ to V .
7. Upon receiving $(Desc, T_1, \mathcal{T}_S, n, \alpha, \text{scriptG}, [T_{xG}], [T_{xP}], [T_{xV}], \sigma_P([T_{xV}]))$, V checks the Glimpse parameters. If they are correct and well-formed, V signs $[T_{xP}]$ and sends $\sigma_V([T_{xP}])$ to P .
8. P verifies if $\sigma_V([T_{xP}])$ is a valid signature. Upon successful verification, P signs $[T_{xG}]$ and posts T_{xG} with witness $\omega := \sigma_P([T_{xG}])$.

Commit on \mathcal{L}_S ($Desc$): P posts $T_{xPayback}$ s.t. $[T_{xPayback}] \leftrightarrow Desc$.
Verify & Commit on \mathcal{L}_D ($Desc, T_1, n$):

1. If $T_{xPayback}$ has n confirmations, P constructs \mathcal{P}^n (as shown below) and signs $[T_{xP}]$. P posts T_{xP} with witness $\omega := (\mathcal{P}^n, \sigma_P([T_{xP}]), \sigma_V([T_{xP}]))$.
2. After T_1 , if θ_G is still unspent, V signs $[T_{xV}]$ and posts T_{xV} with witness $\omega := (\sigma_P([T_{xV}]), \sigma_V([T_{xV}]))$.

Construct $\mathcal{P}^n(Desc, n)$ for $T_{xPayback}$:

1. Isolate x_i^R , i.e., the realization of every parameter x_i in $Desc$.
2. Fetch the header of the block B including $T_{xPayback}$ and compute the Merkle proof MP for $T_{xPayback}$.
3. Retrieve the first n confirmation headers after B .
4. Return $\mathcal{P}^n := (x_i^R, MP, \text{header}(B), \text{confHeaders}_n)$.

scriptG($Desc, T_1, \mathcal{T}_S, n, (P, V)$) outputs a locking script that:

- Upon receiving $\omega = (\mathcal{P}^n, \sigma_P([T_{xP}]), \sigma_V([T_{xP}]))$, does the following:
 1. If x^R is a valid realization of x , i.e., it matches the expected format and length, reconstruct $[T_{xPayback}]$. Else, return \perp .
 2. Compute $\mathcal{H}([T_{xPayback}])$.
 3. Given $\mathcal{H}([T_{xPayback}])$ and $\text{header}(B)$, verify the Merkle proof MP . If successfully verifies, reconstruct the header of B . Else, return \perp .
 4. If the hashes of B and of the n confirmation blocks are smaller than \mathcal{T}_S , go to the next step. Else, return \perp .
 5. If $(\sigma_P([T_{xP}]), \sigma_V([T_{xP}]))$ are valid signatures of P and V over $[T_{xP}]$, unlock the coins. Else, return \perp .
- Upon receiving $\omega = (\sigma_P([T_{xV}]), \sigma_V([T_{xV}]))$, if the current time $t > T_1$, unlock the coins. Else, return \perp .

Figure 4: Glimpse pseudocode for cross-chain lending.

To include liquidation, scriptG has to additionally incorporate the following logic: if, before T_1 , O attests the collateral price on \mathcal{L}_S below a predefined liquidity threshold, V can claim the collateral by publishing a liquidation transaction $T_{\times L} := (1, [\zeta_G], 1, [(\alpha, \text{OneSig}(\text{pk}_V))])$ s.t. $[T_{\times L}] \leftrightarrow \text{Desc}_O$, with witness \mathcal{P}'_O , being \mathcal{P}_O the proof for $T_{\times O}$, i.e., the most recent oracle attestation. The liquidation transaction is constructed in the *Setup* phase, during which P signs it and gives their signature to V .

Our construction enables the first form of trustless peer-to-peer lending on chains having limited scripting capabilities. We note that in Bitcoin-based chains funds cannot be pooled, resulting in borrowers having to seek for lenders. To facilitate matching the demand and supply, we suggest setting up dedicated communication channels or platforms.

4.1 Other Applications

Backed Assets. We refer to *backed assets* as assets issued on a ledger \mathcal{L}_D that are backed by a cryptocurrency or another asset on a ledger \mathcal{L}_S . This category includes assets that are issued on sidechains and backed on parent chains, such as Liquid tokens L-BTC backed by BTC, but also encompasses wrapped tokens, for example, WBTC in Ethereum backed by BTC in Bitcoin.

Sidechains are blockchains tightly bound to a pre-existing parent blockchain with the purpose of enabling or extending some features. Users can easily move funds from the parent chain to the sidechain (and vice versa) through verifiable two-way pegs: assets are locked in an address of the parent chain (sidechain) and are then released on the sidechain (parent chain), ready to be used. Let us remove the liquidation mechanism from the lending protocol in Section 4 and assume V can create assets on \mathcal{L}_D : with these two caveats, the same construction can be used to encode trustless pegs, where V issues new assets on the sidechain and locks them in the Glimpse contract (rather than giving a loan), and P can get them by proving to have sent some assets to the peg address on the parent chain. Similarly, P can get back the funds on the parent chain by proving to have returned the coins to the peg address on the sidechain.

Equivalently, Glimpse can be also used to wrap and unwrap tokens. Wrapped tokens are digital assets that represent other underlying assets, typically from a different blockchain. They are issued (wrapping) on \mathcal{L}_D when the corresponding original tokens have been locked on \mathcal{L}_S , and they are then destroyed or locked (unwrapping) to release the original ones.

Proofs-of-Burn. Proofs-of-Burn prove that a certain amount of cryptocurrency or other valuable assets has been burnt, i.e., has been sent to an unspendable address or a special smart contract where they become permanently irretrievable. Proofs-of-Burn are used, e.g., as a bootstrapping mechanism to get assets on a new chain. In Proofs-of-Burn Glimpse is used as for backed assets, with the only difference being that funds

are moved unidirectionally.

Proofs of Oracle Attestations. Let us assume there exist on \mathcal{L}_S k different oracles posting information about real-world events (e.g., real-time prices for currencies). On \mathcal{L}_D , a user wants to verify k oracle attestations for a specific event. In this case, Glimpse can be used to verify the formula $\mathcal{F}_S = (\text{Desc}_1 \wedge \dots \wedge \text{Desc}_k)$, with Desc_i being the description of the transaction published by the i -th oracle O_i . We note that oracles do not have to cooperate, nor to operate on the same chain.

Off-chain Glimpse and its Applications. State channels [30–32] and generalized channels [15] enable Payment Channel Networks (PCNs) that offer off-chain the same functionality as the underlying chain. We can thus host Glimpse on PCNs, thereby remarkably improving its performance and scalability, and enabling a new range of applications on layer-2. In particular, instead of posting the contract on \mathcal{L}_D , Glimpse can be encoded in a standard channel update, be kept off-chain, and posted on-chain only to resolve disputes. In this way, the contract can be iteratively updated off-chain within subsequent channel updates, allowing for changes of the Glimpse specifics on the run. For instance, one could efficiently and securely extend the contract lifetime by updating the randomness in the contract as well as inside the to-be-verified transaction,⁶ while accordingly adjusting the number of confirmation blocks for the proof. E.g., if initially $T = 2$ hour and $n = 10$, one can update the randomness and extend T by 1 hour minutes only asking for $n = 5$ confirmation blocks: security is preserved and the proof verification cost is decreased. Overall, hosting Glimpse contracts on PCNs dramatically improves its practicality (parameters changed on the run to accommodate users' needs) and its cost (no on-chain transactions in the optimistic case).

Application-wise, payment channel hubs may employ Glimpse contracts to set up betting hubs, where users connected to the hub bet on a certain $T_{\times S}$ being published on-chain within an absolute time T . If $T_{\times S}$ is published, the users and the hub reflect the correct outcome in a channel update within time $t < T$. If any party (user or hub) misbehave, the counterparty can post the contract on-chain (thereby closing the channel) and enforce the correct outcome. Glimpse also provides an out-of-the-box solution for enabling off-chain applications synchronized to an on-chain event. Examples include cross-chain payments based on [30, 33] or cross-chain virtual channels based on [34].

5 Security in the UC Framework

We model Glimpse in the synchronous Global Universal Composability (GUC) framework [35], closely following prior work [15, 31, 32]. First, we state that according to the

⁶We note that this can be done only if the particular application for which Glimpse is used allows for updates of the randomness in the not-yet-posted transaction to be verified.

basic assumptions in Section 3.1, Glimpse achieves weak atomicity. Next, we show that by additionally assuming liveness of the parties and direct access to \mathcal{L}_S for P (and V , in case of DNF verification), Glimpse achieves strong atomicity.

We use a global clock \mathcal{G}_{Clock} [35] and authenticated channels with guaranteed delivery \mathcal{G}_{GDC} [32] to model time and communication. We assume a static corruption model, where the adversary decides which set of parties to corrupt before the execution of the protocol. We use the instantiation of the functionality \mathcal{G}_{Ledger} defined in [36] to model a ledger \mathcal{L} , where the parameters are chosen such that the ledger achieves both *liveness* and *consistency* as defined in [26]. We define two similar ideal functionalities $\mathcal{F}_{W-Glimpse}$ and $\mathcal{F}_{S-Glimpse}$ (see the extended version of this work [27], Appendix B.1), formalizing our desired properties of *weak atomicity* and *strong atomicity* in the general case of DNF verification, respectively. More concretely, the ideal functionality is parameterized over two ledgers \mathcal{L}_S or \mathcal{L}_D . After P and V parties have registered to it, the functionality ensures that the respective transactions are posted on \mathcal{L}_S or \mathcal{L}_D , such that *weak atomicity* or *strong atomicity* holds.

We then formally model our Glimpse protocol Π in the UC framework (see [27], Appendix B.2), and prove that Π realizes $\mathcal{F}_{W-Glimpse}$ or $\mathcal{F}_{S-Glimpse}$ depending on the underlying assumptions. In a nutshell, this is done by designing an ideal world adversary (or simulator) \mathcal{S} and showing that no probabilistic polynomial time *environment* can computationally distinguish between interacting with the real world protocol Π in the presence of an adversary \mathcal{A} and the ideal functionality in the presence of a simulator \mathcal{S} . In other words, \mathcal{S} translates any attack on the protocol into an attack on the ideal functionality, which intuitively means that Π is “as secure”, i.e., has the same properties, as $\mathcal{F}_{W-Glimpse}$ or $\mathcal{F}_{S-Glimpse}$. This is formalized in [27], Appendix B. In Appendix C of [27], we formally prove Theorems 1 and 2, which make use of Definitions 2 to 8 of [27], Appendix B.2. The definitions underlined in the theorems can be found in [27], Appendix B.2.

Theorem 1. *Given the functionalities \mathcal{G}_{Clock} and \mathcal{G}_{GDC} , the protocol Π is instantiated with two ledger instantiations \mathcal{G}_{Ledger} for \mathcal{L}_S and \mathcal{L}_D , and has strictly randomized inputs. Let $\Delta_D \in \mathbb{N}$ be the wait time of \mathcal{L}_D and $gen\mathcal{P}$ a proof generation function that is T-sound. Then, the protocol Π UC-realizes the ideal functionality $\mathcal{F}_{W-Glimpse}$.*

Theorem 2. *Given the functionalities \mathcal{G}_{Clock} and \mathcal{G}_{GDC} , the protocol Π is instantiated with two ledger instantiations \mathcal{G}_{Ledger} for \mathcal{L}_S and \mathcal{L}_D , and has strictly randomized inputs. Let $\Delta_D \in \mathbb{N}$ be the wait time of \mathcal{L}_D and $gen\mathcal{P}$ a proof generation function that is complete and T-sound. **All parties have direct access to \mathcal{L}_S and \mathcal{L}_D , and they exhibit liveness.** Then, the protocol Π UC-realizes the ideal functionality $\mathcal{F}_{S-Glimpse}$.*

6 Economic Security Analysis

We now extend our analysis to incorporate rational players. As for light clients, the security of Glimpse relies on the assumption that the underlying chains operate under a well-designed incentive mechanism that ensures an honest majority of miners, thereby guaranteeing consistency and liveness.

However, introducing cross-chain or cross-layer applications brings forth new external profit opportunities for miners. If the total value locked in the application exceeds the amount of reward offered by the internal incentive mechanism, miners may stop adhering to the protocol rules with the purpose of maximizing their profit. This is true for *all cross-chain and cross-layer applications and bridges* [37]: atomic swaps [38], chain relays [6], payment channels [39], bridges [9, 40], etc. In particular, Glimpse and chain relays equally suffer from forgery attacks, where miners use their computational power to forge a fake subchain of blocks (suffix) rather than mining honestly.

In this section, we study the cost for an adversary bribing miners to mount a proof forgery attack, thereby compromising the security of Glimpse. Furthermore, we study the cost for an adversary bribing miners to mount a censorship attack towards, e.g., $\mathbb{T}_{\times P}$ or $\mathbb{T}_{\times V}$, thereby compromising the security of Glimpse and violating the liveness of the underlying blockchain. For these attacks, we define the *secure parameter space* specific to Glimpse.

6.1 Proof Forgery Attack

In a *proof forgery attack* a malicious prover (attacker) bribes the miners of the source chain to convince them forging a fake proof, i.e., an invalid extension of the longest chain. With such a fake proof, the attacker can fool the Glimpse contract and steal the funds in it. This attack, however, is not limited to Glimpse, but threatens light clients and chain relays as well. Light clients operate under the assumption that the majority of the mining power is in the hands of honest miners, resulting in a good chain quality [26]. If this assumption is broken, a light client can also be fooled to accept a fake n -block suffix. This attack becomes profitable if the total value locked in Glimpse (or in any cross-chain applications relying on a light client or chain relay)⁷ is larger than the profit miners make when mining honestly.

Attack Strategy. Let us consider a powerful attacker consisting of all provers having an *active* Glimpse contract on the *same* \mathcal{L}_S but (potentially) *different* \mathcal{L}_D . The attacker bribes the miners of \mathcal{L}_S to forge a proof for all these Glimpse contracts at once. The bribe consists of the coins held by all the active Glimpse contracts over the considered \mathcal{L}_D . The miners following the attack can optimize their computational effort by forging a *single proof for all the contracts*: they can create a

⁷Without creating any fork on the main chain, corrupted miners can create a fake light client (or chain relay) suffix that will not be part of the blockchain as, e.g., it contains invalid transactions which are rejected by full nodes.

single fake block B^f including all the transactions the attacker wants them to include (the transactions the Glimpse contracts are conditioned on), and then mine n confirmation blocks on top of B^f in time T , with n and T being averaged over the active Glimpse instances.

Total Value Locked in Active Glimpse Contracts. To understand when this attack constitutes a real threat, we first need to know the economic resources the attacker has at their disposal for bribing the miners. This is given by the *total value locked* in all the *simultaneously active* Glimpse contracts operating over the same \mathcal{L}_S . We denote the total value with α_T . Note that α_T must also take into account similar active cross-chain protocols relying on the same proof design [41], i.e., checking that enough PoW has been done within a fixed time frame.

Computing α_T requires monitoring all the destination chains which support Glimpse, and look for the *active* Glimpse contracts (we recall that, for practicality, Glimpse contracts are meant to be active only for a short time). As this could be unpractical, we propose possible alternative solutions: honest parties may use a public bulletin board where they announce their Glimpse contracts (we assume at least one between P and V is honest, see Section 3.1) or, alternatively, use some heuristics to estimate α_T , e.g., computing the total value for the most used Glimpse destination chain and multiplying it by the number of compatible chains. For the analysis that follows, we assume honest parties are able to compute α_T . We leave a more rigorous analysis of how one can compute α_T in the face of such a powerful adversary as future work.

Model. We require the number n of confirmation blocks for Glimpse to be at least equal to the minimum number of blocks for which the probability of a temporary fork (or ordinary block reorganization) is negligible. With this, we prevent a malicious P fooling the contract by submitting a proof taken from an orphaned branch of \mathcal{L}_S . Finally, we require the probability of n being larger than the number of honest blocks mined for \mathcal{L}_S in T to be negligible. With this, we exclude P from being unable to construct a proof because n is larger than the number of blocks honestly appended to \mathcal{L}_S over the time window T .

To study the proof forgery attack, we need to consider the expected gain for honest miners ($\mathbb{E}[\mathbf{H}]$) and the expected gain for corrupted miners mounting the forgery attack ($\mathbb{E}[\mathbf{F}]$). The first represents the money miners would get from following the protocol rules (block rewards and transaction fees). The second is the profit from the attack, i.e., the value of the bribe that miners would get from the attacker if they successfully forge the proof. The miners' total profit is therefore given by $\mathbb{E}[\mathbf{F}] - \mathbb{E}[\mathbf{H}]$. If the total profit is positive, then the attacker (and cooperating miners) is incentivized to mount the attack.

Expected Gain for Honest Miners. Let R be the block reward (in USD) on \mathcal{L}_S , $\mathbb{E}[\mathbf{B}]$ the number of expected blocks

on \mathcal{L}_S in T , and $\mu_r \leq 1 - \gamma$ ⁸ the attacker's *relative* mining power on \mathcal{L}_S (which gives the attacker's probability of finding a valid block). The expected gain for honest miners is:

$$\mathbb{E}[\mathbf{H}] = R \cdot \mathbb{E}[\mathbf{B}] \cdot \mu_r. \quad (1)$$

Expected Gain for Corrupted Miners. The expected gain for corrupted miners depends on n , on their mining power, and on the fluctuation (in USD) of the bribe value during the time window T of the attack. Let μ be the corrupted miners' hashing power (in hashes per second); then, the number of hashes computed in T is given by $N := \mu \cdot T$. We consider N repeated, independent, and equally distributed hash evaluations, and we let P_T be the probability of finding a hash smaller than a target \mathcal{T} . In time T , miners will be able to forge a proof consisting of n confirmation blocks with (binomial) probability given by:

$$P_{n,T} = 1 - \sum_{k=0}^n \binom{N}{k} P_T^k (1 - P_T)^{N-k}. \quad (2)$$

Let α_T be the bribe the attacker offers to miners, and δ be the total expected percentage price drop (over all different \mathcal{L}_D) of the bribe after time T . The expected gain for corrupted miners is given by:

$$\mathbb{E}[\mathbf{F}] = \alpha_T \cdot P_{n,T} \cdot (1 - \delta). \quad (3)$$

We observe that if corrupted miners hold a significant share of the computational power of the network, this attack becomes detectable, undermining users' trust in the chain.

Secure Parameter Space. Glimpse is economically secure when it is more profitable for miners to honestly mine blocks rather than mounting a proof forgery attack. In other words, Glimpse is secure when the total profit for the attacker (and cooperating miners) is negative, i.e., when $\mathbb{E}[\mathbf{F}] < \mathbb{E}[\mathbf{H}]$. This inequality yields the *upper bound* for the total value α_T locked in *simultaneously active* Glimpse contracts over the same \mathcal{L}_S :

$$\alpha_T < \frac{R \cdot \mathbb{E}[\mathbf{B}] \cdot \mu_r}{P_{n,T} \cdot (1 - \delta)}. \quad (4)$$

Before opening a new Glimpse contract of value α over \mathcal{L}_S , an honest party must first compute the total value α_T locked in all the active Glimpse over \mathcal{L}_S , and make sure that the new total value locked, i.e., $\alpha + \alpha_T$, fulfills the inequality in Equation (4). We stress that α_T is the *upper bound at each point in time* and that, for practicality, Glimpse contracts are meant to have a short lifetime. Put differently, Glimpse is a dynamic and fast protocol that can move large amounts of money capped by α_T at each point in time. The value α_T may considerably fluctuate depending on the underlying chains, e.g., the block reward of the source chain, and on the market conditions, e.g., the values (in USD) of the assets involved.

⁸ $\gamma \in [0, 1]$ is the fraction of honest miners the blockchain can tolerate.

Example. Let us assume all active Glimpse contracts have Bitcoin as \mathcal{L}_S and the Liquid Network as \mathcal{L}_D . These contracts have, on average, $n = 5$ and $T = 1$ hour. We consider a Bitcoin target having 19 leading zeros (the largest in 2022), an attacker controlling 23% of the total hashing power (largest mining pool in 2022), and the average prices for BTC and L-BTC at November 2022. Without price drop, $\alpha_T \sim 230$ million USD as of January 2023: this upper bound compares to the one for other bridges, e.g., Gravity [40, 42]. However, while in Glimpse funds are locked for a short time and α_T is the maximum at each point in time, other bridges (e.g., Gravity) have worst performances, as funds are locked for long periods.

6.2 Censorship Attack

Glimpse makes use of a timelock and, as any other protocol relying on timelocks [38, 39, 43], it suffers from *censorship attacks*, i.e., attacks on the liveness of the underlying chains. We investigate how an attack on the liveness of the destination chain harms Glimpse, and we estimate the cost of mounting such an attack.

In a censorship attack, a malicious verifier (attacker) bribes block proposers (validators in Proof-of-Stake, or miners, in PoW) to not include a specific transaction (in our case, e.g., $T_{\times P}$) on chain. Rational proposers, however, will only censor \mathcal{L}_D if they gain something from doing so. Therefore, the economic security of Glimpse depends on the conditions making this attack profitable for V (and the proposers).

Closely following [38], we define the *bribing game* as a Markov game running in $T + 1$ sequential stages, where a stage is the period between two blocks. In each stage, the block proposer chooses between censoring the transaction pointed out by V (playing the game), or including the transaction in the block (refusing to play the game). The bribing game is *safe* if, after eliminating the strictly dominated strategies, the only action left in stage one for each block proposer is to *refuse* the bribery and include the transaction.

The analysis stems from these assumptions: (i) block proposers are rational, i.e., they always try to maximize their profit and, if they can choose, they always follow dominant strategies; (ii) block proposers do not create forks; (iii) the probability μ_r of each player to be selected as block proposer is publicly known and is constant during the attack; (iv) the attacker and the victim are not block proposers; (v) all block proposers can see timelocked transactions that will be valid in the future; (vi) the Glimpse lifetime T is a timelock expressed in number of blocks; finally, (vii) block rewards and fees generated outside the Glimpse protocol are constant and do not affect the attack.

A *weak block proposer* as a player whose probability to be selected as the next block proposer is $\mu_r < \frac{f}{\alpha}$. We let μ_w be the sum of the probabilities of all weak block proposers in the system, f be the fee of the to-be-censored transaction, and α the value of the bribe, i.e., the economic value hold by the (*single!*) Glimpse contract under attack. Let $\alpha > f$. As

proved in [38], the following theorem holds:

Theorem 3. *The bribing game is safe if there is at least one block proposer such that $\mu_r < \frac{f}{\alpha}$ (weak block proposer) and*

$$T > \frac{\log \frac{f}{\alpha}}{\log(1 - \mu_w)}. \quad (5)$$

Secure Parameter Space. Glimpse is secure from censorship attacks when $\alpha < \frac{f}{\mu_r}$ and T fulfills Equation (5).

Example. For instance, considering a \$2 transaction fee in Ethereum and the lowest probability to be selected as block proposer being $1.8 \cdot 10^{-6}$ [44], each Glimpse contract in Ethereum can hold at most 1.1 million USD. With $\mu_w = 1.5\%$ and a quite high fee-to-bribe ratio we have: $T > 25$ with $\frac{f}{\alpha} = 0.7$, $T > 15$ with $\frac{f}{\alpha} = 0.8$, and $T > 7$ with $\frac{f}{\alpha} = 0.9$.

7 Evaluation

On-chain Costs for EVM Chains. We now consider Ethereum, but a similar discussion also applies to any EVM-based chain. In Ethereum, the cost of a transaction is measured in *gas*: every computation consumes an amount of gas proportional to its complexity, and the data stored on-chain consume an amount of gas proportional to its byte length. The computational cost of Glimpse stems from the proof verification, which consists of a Merkle proof verification, a transaction body and a block header reconstruction, and hash comparisons (Figure 4). A Merkle tree with k leaves has a Merkle proof of size $O(\log_2 k)$. This leads to Merkle proof verification cost scaling logarithmically in the number of transactions in a block, as shown in Table 3. Each of the n confirmation blocks in \mathcal{P}^n yields an overhead of 36k gas for the hash comparison.

Besides these computational costs, the Glimpse contract has to store the PoW target of the source chain as well as the hash of the to-be-verified transaction(s) - or the transaction description(s). In Ethereum the data are stored in 32-bytes slots and for each slot 20k gas are consumed: this leads to 40k gas storage cost for the target and the transaction hash, or to approximately 188k gas in the case of target and a ~ 300 -bytes description. We have implemented an open-source cost evaluation which can be found in a Github repository [45].

Glimpse has lower on-chain costs compared to optimized relay solutions, such as Ethrelay [6] and zkRelay [5] (presented in Appendix A), and zkBridge. With Ethrelay, each block header submission results in an average cost of 280k gas, whereas the inclusion of a transaction is verified via SPV combined with an advanced search algorithm that checks for main chain membership. For relatively recent blocks, this leads to a gas consumption of 110k gas. With zkRelay, the submission of a batch of blocks costs 522k gas, including the verification of the zero-knowledge proof and the storage costs. The proof verification alone results in 351k gas. To verify

that a transaction has been included in a block on the source blockchain, users have to provide the relay contract with a two Merkle proofs: one for verifying the block inclusion in the batch, and one for verifying the transaction inclusion in the block. With zkBridge, each zk proof verification alone already results in 230k gas.

While the costs of relays and zkBridge for verifying transaction inclusion in a block are analogous to the ones of Glimpse, the crucial difference lies in the maintenance costs. Chain relays continuously incur in high maintenance costs for *relaying, verifying, and storing the full list of block headers of \mathcal{L}_S* ; similarly, zkBridge requires relayers to continuously relay block headers, compute zk proofs, and submit the headers and the corresponding proofs on-chain. Glimpse, on the other hand, does not have any maintenance cost because of its on-demand nature. With Glimpse, the verification of a single fully known transaction via a proof comprising of 5 confirmation blocks has an upper bound of 330k gas: we stress that this is a *one-time* fee, compared to the continuous 280k gas for each block header submission of Ethrelay and 522k gas per batch submission of zkRelay.

On-chain Costs for Bitcoin-based Chains. In Bitcoin-based chains, transaction fees are usually proportional to the size in bytes of the transaction. In Bitcoin, for instance, this results in a few satoshi per byte as of November 2022.

To cope with the limited scripting capabilities, Glimpse parties can use Taproot, which using Merkelized Abstract Syntax Trees (MAST) [46] allows to commit to multiple scripts within a single one, i.e., a Pay-To-Taproot (P2TR) script which contains the root of the tree. The size of a MAST for Glimpse depends on (i) the number of undefined inputs and outputs in Desc, because their well-formedness needs to be checked with dedicated opcodes, (ii) the number of confirmation blocks in \mathcal{P}^n , because their hashes have to be compared with the PoW target, (iii) the number of transactions within the block, because it affects the number of levels in the tree, and (iv) the size of Desc, being the description hard-coded in the script. For example, for a single to-be-verified transaction Tx of ~ 350 bytes, 6 confirmation blocks, and one parameterized input or output, the upper bound for the MAST is 10MB. For DNF formulas, parties need to compute and exchange the MAST for each literal. For more details, see the extended version of this work [27].

We theoretically estimate the size of transactions T_{X_G} , T_{X_P} , and T_{X_V} on the Liquid Network, where Taproot and all necessary string opcodes are available. Assuming T_{X_G} has two P2PKH inputs and one P2TR output, the transaction size is approximately 350 bytes. Assuming T_{X_P} has one P2TR input and two P2PKH outputs, the size is again roughly 350 bytes. Instead, assuming T_{X_V} has one P2TR input and two P2PKH outputs, its size is of about 200 bytes. Concretely, in November 2022, users' fees for T_{X_G} and T_{X_P} would amount to \$1.5 each, whereas for T_{X_V} to \$0.84. The total cost would be at most 3\$, similar to the costs of standard Bitcoin transactions.

No. Tx in B	$\mathcal{P}^n=0$	$\mathcal{P}^n=5$
$2^6+1 \leq Tx \leq 2^7$	92k gas	273k gas
$2^7+1 \leq Tx \leq 2^8$	95k gas	276k gas
$2^8+1 \leq Tx \leq 2^9$	99k gas	280k gas
$2^9+1 \leq Tx \leq 2^{10}$	103k gas	283k gas
$2^{10}+1 \leq Tx \leq 2^{11}$	106k gas	287k gas
$2^{11}+1 \leq Tx \leq 2^{12}$	111k gas	291k gas

Table 3: On-chain costs for EVM chains for proof verification with $n = 0$ and $n = 5$, for different number of Tx in B.

Computational Overhead. The computational overhead is minimal: parties need to create and verify signatures, and construct proofs. If the destination chain is the Liquid Network, parties also need to construct and verify the MAST. All these operations can be performed using commodity hardware.

Communication Overhead. We now discuss the communication overhead for Glimpse. In the *Setup* phase, parties need to exchange T_{X_G} , T_{X_P} and T_{X_V} , as well as the descriptions of the to-be-verified transactions. Verifying a DNF formula with m literals requires the parties to exchange $4 \cdot 2^m$ transactions and the respective signatures. For Bitcoin-based chains supporting Taproot, parties also need to exchange the MAST, which roughly amounts to 10MB.

8 Conclusion

We present Glimpse, an *on-demand light client for cross-chain communication* that only requires constant-size storage. Glimpse enables many applications such as lending, Proofs-of-Burn, proofs of oracle attestations, and off-chain applications, while remarkably retaining low on-chain costs and *compatibility* with chains having limited scripting capabilities. The security and atomicity properties of Glimpse are proven within the UC framework. By conducting an economic analysis which considers rational players, we provide the secure parameter space for Glimpse.

Acknowledgments

The work was partially supported by CoBloX Labs, by the European Research Council (ERC) under the European Union's Horizon 2020 research (grant agreement 771527-BROWSEC), by the Austrian Science Fund (FWF) through the SpyCode SFB project F8510-N and the project CoRaF (grant agreement 2020388), by the Austrian Federal Ministry for Digital and Economic Affairs, the National Foundation for Research, Technology and Development and the Christian Doppler Research Association through the Christian Doppler Laboratory Blockchain Technologies for the Internet of Things (CDL-BOT), and by the WWTF through the project 10.47379/ICT22045.

References

- [1] "Ronin attack shows cross-chain crypto is a 'bridge' too far," 2022. [Online]. Available: <https://rb.gy/hvo01>
- [2] "Hackers have stolen \$1.4 billion this year using crypto

- bridges. Here's why it's happening," 2022, <https://shorturl.at/yGJT3>.
- [3] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," 2009, <http://bitcoin.org/bitcoin.pdf>.
- [4] "BTC Relay," <https://github.com/ethereum/btcrelay>.
- [5] M. Westerkamp and J. Eberhardt, "zkRelay: Facilitating sidechains using zkSNARK-based chain-relays," in *IEEE European Symposium on Security and Privacy Workshops*, 2020.
- [6] P. Fraunthaler, M. Sigwart, C. Spanring, M. Sober, and S. Schulte, "ETHRelay: A cost-efficient relay for Ethereum-based blockchains," in *IEEE International Conference on Blockchain*. IEEE, 2020.
- [7] A. Kiayias, A. Miller, and D. Zindros, "Non-interactive Proofs of Proof-of-Work," in *Financial Cryptography and Data Security FC*, 2020.
- [8] B. Bünz, L. Kiffer, L. Luu, and M. Zamani, "FlyClient: Super-light clients for cryptocurrencies," in *IEEE Symposium on Security and Privacy, SP*, 2020.
- [9] T. Xie, J. Zhang, Z. Cheng, F. Zhang, Y. Zhang, Y. Jia, D. Boneh, and D. Song, "zkBridge: Trustless cross-chain bridges made practical," in *ACM SIGSAC Conference on Computer and Communications Security, CCS*, 2022.
- [10] M. Bartoletti and R. Zunino, "BitML: A calculus for bitcoin smart contracts," in *ACM SIGSAC Conference on Computer and Communications Security, CCS*, 2018.
- [11] "How to validate Bitcoin payments in Ethereum (for only 700k gas!)," 2018, <https://medium.com/summa-technology/cross-chain-auction-technical-f16710bfe69f>.
- [12] "Summa," <https://github.com/summa-tx/bitcoin-spv>.
- [13] J. Prestwich, "Non-atomic swaps," 2019, <https://ethresear.ch/t/stateless-spv-proofs-and-economic-security/5451>.
- [14] F. Barbàra and C. Schifanella, "BxTB: cross-chain exchanges of bitcoins for all Bitcoin wrapped tokens," in *Fourth International Conference on Blockchain Computing and Applications, BCCA*, 2022.
- [15] L. Aumayr, O. Ersoy, A. Erwig, S. Faust, K. Hostáková, M. Maffei, P. Moreno-Sanchez, and S. Riahi, "Generalized channels from limited blockchain scripts and adaptor signatures," in *Asiacrypt*, 2021.
- [16] "The Liquid Network," <https://blockstream.com/liquid/>.
- [17] "Blockstream," <https://blockstream.com>.
- [18] "What the heck is SegWit," 2020, <https://medium.com/bitbees/what-the-heck-is-segwit-3f58b7352b1c>.
- [19] "Bitcoin's Taproot upcoming upgrade and how it matters to the network," 2021, <https://tokenize.exchange/blog/article/bitcoin-taproot-upcoming-upgrade>.
- [20] "Salvador Bitcoin Bonds," 2021, <https://rb.gy/fcku5>.
- [21] "Taproot: SegWit version 1 spending rules," <https://github.com/bitcoin/bips/blob/master/bip-0341.mediawiki>, 2020.
- [22] "Validation of Taproot scripts," <https://github.com/bitcoin/bips/blob/master/bip-0342.mediawiki>, 2020.
- [23] A. Zamyatin, M. Al-Bassam, D. Zindros, E. Kokoris-Kogias, P. Moreno-Sanchez, A. Kiayias, and W. J. Knottenbelt, "SoK: Communication across distributed ledgers," in *Financial Cryptography and Data Security: 25th International Conference, FC 2021*, 2021.
- [24] "BSIP 64: Optional HTLC preimage length and add hash160 algorithm," <https://github.com/bitshares/bsips/issues/163>.
- [25] "Bitcoindev Speedy covenants (OP_CAT2)," <https://lists.linuxfoundation.org/pipermail/bitcoin-dev/2022-May/020434.html>, 2022.
- [26] J. A. Garay, A. Kiayias, and N. Leonardos, "The Bitcoin backbone protocol: Analysis and applications," in *Advances in Cryptology - EUROCRYPT*. Springer, 2015.
- [27] G. Scaffino, L. Aumayr, Z. Avarikioti, and M. Maffei, "Glimpse: On-demand PoW light client with constant-size storage for DeFi," *Cryptology ePrint Archive*, Paper 2022/1721, 2022, <https://eprint.iacr.org/2022/1721>.
- [28] T. Dryja, "Discreet Log Contracts," <https://adiabat.github.io/dlc.pdf>.
- [29] M. Sober, G. Scaffino, C. Spanring, and S. Schulte, "A voting-based blockchain interoperability oracle," in *IEEE International Conference on Blockchain*, 2021.
- [30] A. Miller, I. Bentov, R. Kumaresan, and P. McCorry, "Sprites and state channels: Payment networks that go faster than lightning," in *FC 2019: Financial Cryptography and Data Security*.
- [31] S. Dziembowski, S. Faust, and K. Hostáková, "General State Channel Networks," in *Computer and Communications Security, CCS*, 2018.
- [32] S. Dziembowski, L. Eckey, S. Faust, J. Hesse, and K. Hostáková, "Multi-party Virtual State Channels," in *Advances in Cryptology - EUROCRYPT*, 2019.

- [33] L. Aumayr, P. Moreno-Sanchez, A. Kate, and M. Maffei, “Blitz: Secure Multi-Hop Payments Without Two-Phase Commits,” in *USENIX Security Symposium*, 2021.
- [34] L. Aumayr, P. M. Sanchez, A. Kate, and M. Maffei, “Breaking and Fixing Virtual Channels: Domino Attack and Donner,” in *NDSS*, 2023.
- [35] R. Canetti, Y. Dodis, R. Pass, and S. Walfish, “Universally composable security with global setup,” in *Theory of Cryptography*, 2007.
- [36] C. Badertscher, U. Maurer, D. Tschudi, and V. Zikas, “Bitcoin as a transaction ledger: A composable treatment,” in *Advances in Cryptology – CRYPTO 2017*.
- [37] “Vitalik Buterin on cross-chain applications,” <https://rb.gy/hvo01>, 2022.
- [38] T. Nadahalli, M. Khabbazian, and R. Wattenhofer, “Timelocked bribing,” in *Financial Cryptography and Data Security*, N. Borisov and C. Diaz, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2021.
- [39] Z. Avarikioti, L. Thyfronitis, and S. Orfeas, “Suborn channels: Incentives against timelock bribes,” in *Financial Cryptography and Data Security*, I. Eyal and J. Garay, Eds. Springer International Publishing, 2022.
- [40] “Gravity bridge,” <https://github.com/Gravity-Bridge/Gravity-Docs>.
- [41] “Summa proofs are not composable,” 2019. [Online]. Available: <https://medium.com/@dionyziz/summa-proofs-are-not-composable-57b87825f428>
- [42] “Value locked in Ethereum L1 bridges,” 2023, <https://www.theblock.co/data/scaling-solutions/scaling-overview/value-locked-of-ethereum-l1-bridges>.
- [43] I. Tsabary, M. Yechieli, A. Manuskin, and I. Eyal, “MAD-HTLC: Because HTLC is crazy-cheap to attack,” in *IEEE Symposium on Security and Privacy, SP*, 2021.
- [44] “Beacon scan,” <https://beaconscan.com/statistics>.
- [45] “Glimpse Github,” <https://github.com/Glimpse-CrossChainPrimitive/Glimpse>.
- [46] “Merkelized Abstract Syntax Tree (MAST),” <https://bitcoinops.org/en/topics/mast/>.
- [47] A. Zamyatin, D. Harz, J. Lind, P. Panayiotou, A. Gervais, and W. Knottenbelt, “XCLAIM: Trustless, interoperable, cryptocurrency-backed assets,” in *2019 IEEE Symposium on Security and Privacy (SP)*, 2019.
- [48] TierNolan, “Alt chains and atomic transfers,” 2013.
- [49] M. Herlihy, “Atomic cross-chain swaps,” *CoRR*, 2018. [Online]. Available: <http://arxiv.org/abs/1801.09515>
- [50] J. Xu, D. Ackerer, and A. Dubovitskaya, “A game-theoretic analysis of cross-chain atomic swaps with htcls,” *CoRR*, 2020. [Online]. Available: <https://arxiv.org/abs/2011.11325>
- [51] J. Gugger, “Bitcoin-Monero cross-chain atomic swap,” Cryptology ePrint Archive, 2020. [Online]. Available: <https://eprint.iacr.org/2020/1126>
- [52] S. Thyagarajan, G. Malavolta, and P. Moreno-Sanchez, “Universal atomic swaps: Secure exchange of coins across all blockchains,” in *2022 IEEE Symposium on Security and Privacy (SP) (SP)*, 2022.
- [53] P. Hoenisch, S. Mazumdar, P. Moreno-Sanchez, and S. Ruj, “Lightswap: An atomic swap does not require timeouts at both blockchains,” Cryptology ePrint Archive, 2022. [Online]. Available: <https://eprint.iacr.org/2022/1650>
- [54] “Submarine swap in Lightning Network,” <https://wiki.ion.radar.tech/tech/research/submarine-swap>, 2021.
- [55] “What is atomic swap and how to implement it,” <https://www.axiomadev.com/blog/what-is-atomic-swap-and-how-to-implement-it/>.
- [56] M. Westerkamp and M. Diez, “Verilay: A verifiable Proof of Stake chain relay,” in *IEEE International Conference on Blockchain and Cryptocurrency, ICBC*, 2022.
- [57] T. Bugnet and A. Zamyatin, “XCC: Theft-resilient and collateral-optimized cryptocurrency-backed assets,” Cryptology ePrint Archive, 2022.
- [58] “Bitcoin Wiki: Payment channels,” 2018, https://en.bitcoin.it/wiki/Payment_channels.
- [59] G. Malavolta, P. Moreno-Sanchez, C. Schneidewind, A. Kate, and M. Maffei, “Anonymous multi-hop locks for blockchain scalability and interoperability,” in *Network and Distributed System Security Symposium, NDSS*, 2019.
- [60] S. Dziembowski, L. Eckey, S. Faust, and D. Malinowski, “Perun: Virtual payment hubs over cryptocurrencies,” in *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2019, pp. 106–123.
- [61] L. Aumayr, O. Ersoy, A. Erwig, S. Faust, K. Hostáková, M. Maffei, P. Moreno-Sanchez, and S. Riahi, “Bitcoin-Compatible Virtual Channels,” in *IEEE Symposium on Security and Privacy*, 2021.

A Further Related Work

The idea of chain relays first appeared with BTC Relay [4], realizing a Bitcoin relay on Ethereum. BTC Relay verifies and stores Bitcoin block headers; the costs the relayers had to bare for keeping the relay up-to-date are high (linear in the total number of blocks within the blockchain) and not compensated by user’s fees.

Westerkamp et al. [5] introduced zkRelay which batches multiple headers. Their validity is verified off-chain and proven on-chain via zkSNARKs. zkRelay has constant verification costs and releases the target ledger from processing and storing every single block header of the source blockchain. Although the on-chain costs are lower than for BTC Relay, a maintenance overhead for the off-chain computation and for on-chain storage remain. Furthermore, the users’ costs for transaction inclusion verification are doubled, as both the block inclusion in the batch and the transaction inclusion in the block have to be verified.

Frauenthaler et al. [6] propose Ethrelay, a relay adopting an optimistic approach: Block headers are optimistically accepted and only validated on-demand. The computational costs per header are cut out, but the storage costs persist.

Zamyatin et al. [47] propose XCLAIM, a framework for trustless and efficient cross-chain exchanges. XCLAIM exhibits functionalities for issuing, transferring, swapping and redeeming cryptocurrency-backed assets securely on existing blockchains. To make the protocol non-interactive, the XCLAIM implementation operating between Bitcoin and Ethereum makes use of a chain relay on Ethereum, specifically of the implementation of BTC Relay. The relay costs are shared among all users of XCLAIM, with decreasing costs for very active users.

Gravity [40] is a bidirectional bridge solution between Ethereum and the Cosmos ecosystem. The Gravity bridge has two main components: a Solidity smart contract deployed on Ethereum and a Cosmos SDK blockchain module. Users deposit assets on one side of the bridge (e.g., Cosmos) and a token representation is minted on the other side of the bridge (e.g., Ethereum), and vice versa. Gravity relies on 2/3 of a set of 140 validators to sign transactions attesting on Cosmos deposits on the Ethereum side and vice versa. To join as a validator, one has to stake assets, which are slashed upon detected misbehavior. Gravity assumes an honest super majority of validators.

Another conceptually and technically different solution for cross-chain communication is atomic swaps, which likely originated from a forum user TierNolan [48] and was later analyzed by, e.g., Herlihy [49] or Xu et al. [50]. Atomic swaps allow multiple parties to exchange assets across multiple blockchains in a distributed and coordinated manner. Different constructions have been proposed by [51–53].

Table 4 compares Glimpse to other state-of-the-art cross-chain solutions. With ①, ②, ③ we denote three classes of

	Commit on \mathcal{L}_S			Ver.&Comm. on \mathcal{L}_D		Expressiveness
	Ass.	SR	Consensus	Ass.	SR	
Universal Atomic Swap [52]	Sync	①	Any	Sync	①	Secret-based logic
HTLC-based Swap [49, 54, 55]	Sync	①	Any	Sync	①	Secret-based logic
Glimpse	Sync	①	PoW	Sync	②	DNF formulas
Chain relays [4–6, 56]	Sync	③	PoW, PoS	Sync	③	Arbitrary logic
XCLAIM [47], XCC [57]	TTP	①	PoW, PoS	Sync	③	Arbitrary logic
Gravity Bridge [40]	TTP	③	PoS, BFT	TTP	③	Arbitrary logic

Table 4: State-of-the-art CCC protocols w.r.t.: (i) the assumption they make (Trusted Third Party (TTP) or Synchrony), (ii) their scripting requirements (SR), (iii) the consensus they operate on, and the expressiveness they achieve.

scripting languages: ① comprises hash locks, time locks, and signature locks, ② includes the operations in ① along with the following functionalities for string concatenation and hash comparison, and ③ finally represents any quasi-Turing complete language.

Lock Contract Limitations. Existing cross-chain communication solutions not relying on a TTP fall into two main categories: lock contracts and chain relays. Lock contracts are an umbrella term for non-custodial locking mechanisms (e.g., Hashed-Timelocked-Contracts⁹, adaptor signatures) that achieve security and atomicity from the hardness of some cryptographic assumptions. Hash locks and adaptor signatures are, for instance, lock contract schemes broadly used to encode blockchain applications such as atomic swaps, payment channels [30, 58], multi-hop payments [33, 59], virtual channels [32, 34, 60, 61], and discreet log contracts [28]. Lock contracts use a statement S that ties the authorization of a transaction T_{x_2} to the leakage of a secret witness s of some hard relation (usually leaked within a transaction T_{x_1} posted on-chain). Lock contracts can encode a class of *asymmetric problems*: *The party posting transaction T_{x_1} cannot be the same posting transaction T_{x_2} .* Intuitively, the party who posts transaction T_{x_2} has to gain knowledge of s only after transaction T_{x_1} has been posted. Lock contracts are cheap and lightweight, and since they require minimal scripting capabilities, they can be leveraged on all existing chains. On the other hand, they enable a very limited number of (asymmetric) applications: They cannot be used, e.g., for Proofs-of-Burn, wrapping and unwrapping of tokens, etc.

B Verifying DNF Formulas with Glimpse

As introduced in Section 3.4, Glimpse can efficiently verify *DNF formulas over descriptions* (Figure 2), allowing to synchronize *any logical combination of transactions on \mathcal{L}_S* with corresponding transactions on \mathcal{L}_D . DNF formulas express truth tables in terms of disjunctions (OR) of conjunctions (AND) of one or more descriptions.

We explain how Glimpse achieves this degree of expressiveness by showing a concrete example. Let us consider two

⁹HTLCs are contracts storing a pair (h, t) and ensuring that if the contract receives the secret s such that $h = \mathcal{H}(s)$ before time t , then the ownership of the asset locked in the contract is transferred to the counter party.

oracles, i.e., O_1 and O_2 , operating on \mathcal{L}_S (they can also operate on different chains) and regularly posting information about a real-world event. On \mathcal{L}_D , prover P and verifier V lock $\frac{\alpha}{2}$ coins each in a Glimpse contract (Tx_G) and, e.g., they bet on a specific outcome for the event: if at least one oracle attests the desired outcome for the event (1-out-of-2 threshold), P can claim the α coins by proving to Tx_G that at least one oracle has published a transaction attesting the established outcome for the event; if the coins are still unspent after time T , V can claim the coins.

We recall that the to-be-verified outcome for the event is specified in the description Desc of each transaction and is hardcoded within Tx_G . We let Desc_i be the description for the transaction published by the oracle O_i , and we let

$$\mathcal{F}_S = (\text{Desc}_1 \wedge \neg \text{Desc}_2) \vee (\neg \text{Desc}_1 \wedge \text{Desc}_2) \vee (\text{Desc}_1 \wedge \text{Desc}_2)$$

be the DNF formula Glimpse has to verify. Glimpse proceeds as follows:

Setup. θ_P and θ_V are unspent outputs on \mathcal{L}_D holding $\frac{\alpha}{2}$ coins each and controlled by P and V , respectively. We denote with $\alpha := \theta_P.\text{coins} + \theta_V.\text{coins}$ the value locked in Glimpse and with ζ_P and ζ_V inputs spending θ_P and θ_V , respectively. The parties construct $[\text{Tx}_G] := (2, [\zeta_P, \zeta_V], 3, [\theta_\alpha, \theta_{\epsilon_1}, \theta_{\epsilon_2}])$, such that $\theta_\alpha := (\alpha, (\text{MuSig}(\text{pk}_P, \text{pk}_V)) \vee (\text{MuSig}(\text{pk}_P, \text{pk}_V) \wedge T_3))$, $\theta_{\epsilon_1} := (\epsilon_1, (\text{scriptG}(\text{Desc}_1, T_1, \mathcal{T}_S, n_1, P)))$, and $\theta_{\epsilon_2} := (\epsilon_2, (\text{scriptG}(\text{Desc}_2, T_1, \mathcal{T}_S, n_2, P)))$. We denote with ϵ_i the smallest possible amount of cash. Its value does not matter, since it is just used to enable the construction.

This means that when Glimpse verifies DNF a formula, Tx_G has to have as many outputs (θ_{ϵ_i}) as the number of transactions in the formula (in this case, two: $\theta_{\epsilon_1}, \theta_{\epsilon_2}$), plus an additional one holding the Glimpse value (θ_α). We now require *both* P and V to sample a random string and embed it in the transaction descriptions. Beside Tx_G , the parties create a set of transactions $(\text{Tx}_T, \text{Tx}_F, \text{Tx}_P)_i$ for each disjunctive term in the formula. In our example, \mathcal{F}_S has three terms: $(\text{Desc}_1 \wedge \neg \text{Desc}_2)$, $(\neg \text{Desc}_1 \wedge \text{Desc}_2)$, and $(\text{Desc}_1 \wedge \text{Desc}_2)$, therefore we will have three sets of transactions $(\text{Tx}_T, \text{Tx}_F, \text{Tx}_P)$. Figure B.1 shows an example of transaction set $(\text{Tx}_G, (\text{Tx}_T, \text{Tx}_F, \text{Tx}_P)_i, \text{Tx}_V)$ for the term $(\text{Desc}_1 \wedge \neg \text{Desc}_2)$ of \mathcal{F}_S . In general, $(\text{Tx}_T, \text{Tx}_F, \text{Tx}_P)_i$ is constructed as follows:

- Tx_T allows P to prove the inclusion of the oracles' transactions attesting the desired outcome for the event. Tx_T spends the outputs θ_{ϵ_i} for the Desc_i in the term (but not for the $\neg \text{Desc}_i$), and it has a single output $\theta_T := (\epsilon, \text{OneSig}(\text{pk}_P))$.
- Tx_F allows V to submit a proof for a transaction being posted on \mathcal{L}_S , as a reaction to a malicious P falsely claiming the transaction was not posted. Tx_F spends the outputs θ_{ϵ_i} for the $\neg \text{Desc}_i$ in the term (but not for the Desc_i), and it has as many outputs as the number of its inputs, each one of the form $\theta_{F,i} := (\epsilon_i, (\text{scriptG}(\text{Desc}_i, T_2, \mathcal{T}_S, n_i, (V, P)))$. For some terms, Tx_F is not needed at all, e.g., $(\text{Desc}_1 \wedge \text{Desc}_2)$.

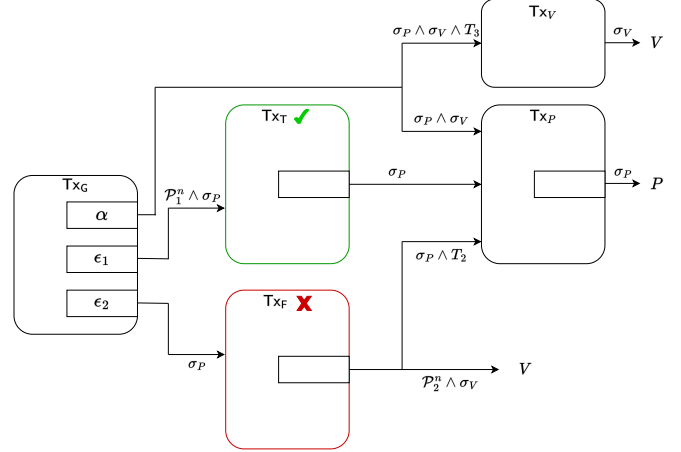


Figure B.1: Set $(\text{Tx}_G, (\text{Tx}_T, \text{Tx}_F, \text{Tx}_P)_i, \text{Tx}_V)$ of transactions for efficiently verifying the term $\text{Desc}_1 \wedge \neg \text{Desc}_2$ of the DNF formula in the example.

- Tx_P allows P to claim the α coins if all the outputs of Tx_F are still unspent. It spends θ_T , all the $\theta_{F,i}$, and θ_α . Finally, the parties create transaction Tx_V allowing V to spend the output θ_α after time T_3 . We have $T_1 < T_2 < T_3$.

At this point, P signs $[\text{Tx}_V]$ and sends to V the Glimpse specifics $(\text{Desc}_1, \text{Desc}_2, T_1, T_2, \mathcal{T}_S, n_1, n_2, \alpha, \text{scriptG}, [\text{Tx}_G], ([\text{Tx}_T], [\text{Tx}_F], [\text{Tx}_P])_{i,\forall i}, [\text{Tx}_V], \sigma_P([\text{Tx}_V]))$. V checks if the Glimpse specifics are well-formed, and verifies the validity of P 's signature. If everything is correct, V signs $[\text{Tx}_G]$ and $[\text{Tx}_P]_{i,\forall i}$, and sends the signatures to P . P checks if V 's signatures are valid and, if so, posts Tx_G on \mathcal{L}_D .

Commit on \mathcal{L}_S . The oracles publish transactions attesting the outcome of the real-world event, e.g., O_1 publishes Tx_1 s.t. $[\text{Tx}_1] \leftrightarrow \text{Desc}_1$ and O_2 publishes Tx_2 s.t. $[\text{Tx}_2] \leftrightarrow \text{Desc} \neq \text{Desc}_2$. This is, e.g., the case depicted in Figure B.1.

Verify & Commit on \mathcal{L}_D . P and V monitor \mathcal{L}_S (or query a relayer R) checking for the inclusion of transactions matching descriptions Desc_1 and Desc_2 . P constructs the corresponding proofs, and claims the coins by posting the corresponding set $(\text{Tx}_T, \text{Tx}_F, \text{Tx}_P)_i$.

V checks whether the set of transactions published by P corresponds to the correct term of \mathcal{F}_S realized by the oracles. If P does not spend θ_α , V can publish Tx_V and get the Glimpse coins after T_3 . If P misbehaves, V can react within T_2 and spend one of Tx_F 's outputs, thereby invalidating Tx_P_i and claiming the money via Tx_V .

Remarks. We observe that when verifying DNF formulas with Glimpse, V needs to be able to query the relayer R (or run a full node), as he needs to construct and submit proofs in case P cheats. Although increasing the off-chain communication overhead, this construction results in up to three on-chain transactions in the optimistic case, no matter the complexity of the formula to verify.