

# Astrape: Anonymous Payment Channels with Boring Cryptography\*

Yuhao Dong, Ian Goldberg, Sergey Gorbunov, and Raouf Boutaba

University of Waterloo

**Abstract.** The increasing use of blockchain-based cryptocurrencies like Bitcoin has run into inherent scalability limitations of blockchains. Payment channel networks, or PCNs, promise to greatly increase scalability by conducting the vast majority of transactions outside the blockchain while leveraging it as a final settlement protocol. Unfortunately, first-generation PCNs have significant privacy flaws. In particular, even though transactions are conducted off-chain, anonymity guarantees are very weak.

In this work, we present Astrape, a novel PCN construction that achieves strong security and anonymity guarantees with simple, black-box cryptography, given a blockchain with flexible scripting. Existing anonymous PCN constructions often integrate with specific, often custom-designed, cryptographic constructions. But at a slight cost to asymptotic performance, Astrape can use any generic public-key signature scheme and any secure hash function, modeled as a random oracle, to achieve strong anonymity, by using a unique construction reminiscent of onion routing. This allows Astrape to achieve provable security that is “generic” over the computational hardness assumptions of the underlying primitives. Astrape’s simple cryptography also lends itself to more straightforward security proofs compared to existing systems.

Furthermore, we evaluate Astrape’s performance, including that of a concrete implementation on the Bitcoin Cash blockchain. We show that despite worse theoretical time complexity compared to state-of-the-art systems that use custom cryptography, Astrape operations on average have a very competitive performance of less than 10 milliseconds of computation and 1 KB of communication on commodity hardware. Astrape explores a new avenue to secure and anonymous PCNs that achieves similar or better performance compared to existing solutions.

## 1 Introduction

### 1.1 Payment channel networks

Blockchain cryptocurrencies are gaining in popularity and becoming a significant alternative to traditional government-issued money. For instance, over 300,000 Bitcoin transactions alone [2] are processed every day. Unfortunately, such high demand inevitably leads to well-known scalability barriers [8]. Bitcoin, for instance, processes less than 10 transactions every second [20], far less than a reasonable global payment system.

---

\* An extended version of this paper, and its accompanying source code, is available. [12]

*Payment channels* [11] are a common technique to scale cryptocurrency transactions. In a nutshell, Alice and Bob open a payment channel by submitting a single transaction to the blockchain, locking up a sum of cryptocurrency from both of the parties. They can then pay each other by simply mutually signing a division of the locked money. Additional blockchain transactions are required only when the channel is closed by submitting an up-to-date signed division, unlocking the latest balances of Alice and Bob. This allows most activity to remain off-chain, while retaining the blockchain for final settlement: as long as the blockchain is secure, nobody can steal funds. More importantly, payment channels can be organized into *payment-channel networks* (PCNs) [20], where users without any open channels between them can pay each other through intermediaries.

## 1.2 Anonymity in PCNs

Unfortunately, “first-generation” PCNs based on the HTLC (hash time-locked contract), such as Lightning Network [11], have a significant problem — poor anonymity [18]. In the worst case, HTLC payments are as transparently linkable as blockchain payments [18], threatening the improved privacy that is often cited [23, 24] as a benefit of PCNs. Furthermore, naive implementations fall victim to subtle fee-stealing attacks, like the “wormhole attack” [19], that threaten economic viability.

A sizable body of existing work on fixing PCN security and privacy exists. On one hand, specialized constructions achieve strong anonymity in specific settings, such as Bolt [14] for hub-based PCNs on the Zcash blockchain, providing for indistinguishability of two concurrent transactions even when all intermediaries are malicious. On the other hand, general solutions for all PCN topologies, like Fulgor [18] and the AMHL (Anonymous Multi-hop Locks) family [19], achieve a somewhat weaker, topology-dependent notion of anonymity: *relationship anonymity* [6, 17]. This property, common to onion-routing and other anonymous communication protocols, means that two concurrent transactions cannot be distinguished as long as they share at least one honest intermediary.

## 1.3 Why boring cryptography?

Unfortunately, there remains a shortcoming common to all existing anonymous PCN constructions — custom, often number-theoretic and sometimes complex cryptographic primitives. No existing anonymous PCN construction limits itself to the bare-bones cryptographic primitives used in HTLC — black-box access to a generic signature scheme and hash function. For example, AMHL uses either homomorphic one-way functions or special constructions that exploit the mathematical structure of ECDSA or Schnorr signatures and Tumblebit uses a custom cryptosystem based on the RSA assumption. Blitz [5], though relying on an ostensibly black-box signature scheme, requires it to have a property<sup>1</sup> that rules out many post-quantum signature schemes.

<sup>1</sup> In particular, the ability for any party, given any public key, to generate new public keys that correspond to the same private key yet are unlinkable to the previous public

However, it is unclear that relationship anonymity requires sophisticated techniques. Relationship anonymity appears to be relatively “easy” elsewhere. Well-understood anonymous constructs like onion routing and mix networks exist for communication with no more than standard primitives used in secure communication (symmetric and asymmetric encryption). Of course, communication is probably easier — indeed some go beyond relationship anonymity with only simple cryptography — but it seems plausible that PCNs can use similarly elementary primitives to achieve anonymity.

Furthermore, “boring” cryptography has practical advantages. For one, non-standard cryptography poses significant barriers to adoption. Reliable and performant implementations of novel cryptographic functions are difficult to obtain, and tight coupling between a PCN protocol and a particular cryptographic construction makes swapping out primitives impossible. With use of black-box cryptography, a system is *generic over cryptographic hardness assumptions* — instead of assuming that, say, the RSA or discrete-log problems are hard, we only need to assume that there exists, for example, *some* secure signature scheme and *some* secure hash function.

Thus, we believe that efficient yet privacy-preserving PCNs that only use well-understood and easily replaced black-box cryptographic primitives are crucial to usable PCNs. In fact, AMHL’s authors already proposed that “an interesting question related to [anonymous PCN constructions] is under which class of hard problems such a primitive exists” [19] that they conjecturally answered with linear homomorphic one-way functions.

## 1.4 Our contributions

In this paper, we present Astrape,<sup>2</sup> a PCN protocol that limits itself to “boring”, generic cryptography already used in HTLC, yet achieves strong relationship anonymity. Despite achieving comparable security, privacy, and performance to other anonymous PCN constructions, Astrape does not introduce any cryptographic constructs other than those used in HTLC. This is accomplished using a novel construct reminiscent of onion routing that avoids the use of any form of zero-knowledge verification.

## 2 Background and related work

### 2.1 First-generation PCNs with HTLC

An extremely useful property of payment channels is that they can be used to construct *payment channel networks* (PCNs) [20, 10, 8], allowing users without channels directly between each other to pay each other via intermediaries. At the heart of any PCN is a *secure multi-hop transaction* mechanism — some way of

---

key. This is crucial to the “stealth addresses” that Blitz’s pseudonymous privacy rests upon.

<sup>2</sup> Greek for “lightning”, pronounced “As-trah-pee”.

Alice paying Bob to pay Carol without any trust in Bob. Most PCNs implement this using a smart contract known as the *Hash Time-Lock Contract* (HTLC). An HTLC is parameterized over a *sender* Alice, the *recipient* Bob, a deadline  $t$ , and a *puzzle*  $s$ . It locks up a certain amount of money, unlocking it according to the following rules:

- The money goes to Bob if he produces  $\pi$  where  $H(\pi) = s$  before time  $t$ , where  $H$  is a secure hash function.
- Otherwise, the money goes to Alice.

We can use HTLC to construct secure multi-hop transactions. Consider a sender  $U_0$  wishing to send money to a recipient  $U_n$  through untrusted intermediaries  $U_1, \dots, U_{n-1}$ . At first,  $U_0$  will generate a random  $\pi$  and  $s = H(\pi)$ , while sending the pair  $(\pi, s)$  to  $U_n$  over a secure channel.  $U_0$  can then lock money in a HTLC parameterized over  $U_0, U_1, s, t_1$ , notifying  $U_1$ .  $U_1$  would send an HTLC over  $U_1, U_2, s, t_2$ , notifying  $U_2$ , and so on. The deadline must become earlier at each step —  $t_1 > t_2 > \dots > t_n$  — this ensures that in case of an uncooperative or malicious intermediary, funds always revert to the sender.

The payment eventually will be routed to  $U_n$ , who will receive an HTLC over  $U_{n-1}, U_n, s, t_n$ . The recipient will claim the money by providing  $\pi$ ; this allows  $U_{n-1}$  to claim money from  $U_{n-2}$  using the same  $\pi$ , and so on, until all outstanding HTLC contracts are fulfilled.  $U_0$  has successfully sent money to  $U_n$ , while the preimage resistance of  $H$  prevents any intermediary from stealing the funds.

## 2.2 Hub-based anonymous payment channels

Unfortunately, HTLC has an inherent privacy problem — a common identifier  $s = H(\pi)$  visible to all nodes in the payment path [14, 18, 19]. This motivates *anonymous* PCN design. *Hub-based* approaches form the earliest kind of anonymous PCN design. Here, the shape of the network is limited to a star topology with users communicating with a centralized hub. Some solutions are highly specialized, such as Green and Miers’ Bolt [14], which relies on the Zcash blockchain’s zero-knowledge cryptography. Other solutions, such as Tumblebit [15] and the more recent A2L [22], provide more general solutions that work on a wide variety of blockchains.

Hub-based PCN constructions tackle the difficult problem of providing unlinkability between transactions despite the existence of only a single untrusted intermediary. It is therefore unsurprising that specialized cryptography is needed to protect anonymity. On the other hand, observations of real-world PCNs like the Lightning Network, as well as economic analysis [13], show that actual PCNs often have intricate topologies without dominating hubs. General, topology-agnostic solutions are thus more important to deploying private PCNs in practice.

## 2.3 Relationship-anonymous payment channels

Unlike hub-based approaches, where no intermediaries are trusted, general private PCN constructions target *relationship anonymity*. This concept, shared

with onion routing and other anonymous communication protocols, assumes at least one honest intermediary. Thus intermediaries are in fact crucial to relationship-anonymous PCNs’ privacy properties. Like most hub-based approaches, relationship-anonymous payment channels do not by themselves deal with information leaked by side channels such as timing and value.

The earliest solution to PCN privacy in this family was probably Fulgor and Rayo [18], a closely related pair of constructions that can be ported to almost all HTLC-based PCNs. Fulgor/Rayo combines a “multi-hop HTLC” contract with out-of-band ZKPs to remove the common identifier across payment hops.

In a later work, Malavolta et al. [19] introduced *anonymous multi-hop locks* (AMHL), a rigorous theoretical framework for analyzing private PCN contracts. The AMHL paper provided a concrete instantiation using linear homomorphic one-way functions (hOWFs), as well as a conjecture that hOWFs are necessary for implementing anonymous PCNs. They also presented a variant that uses a clever encoding of homomorphic encryption in ECDSA to be used in ECDSA-based cryptocurrencies like Bitcoin. The latter “scriptless” variant was generalized in later work to a notion of adaptor signatures [4], where a signature scheme like ECDSA is “mangled” in such a way that a correct signature reveals a secret based on a cryptographic condition. The authors of AMHL also discovered “wormhole attacks” on HTLC-based PCNs. These attacks exploit a fundamental flaw in the HTLC construction to allow malicious intermediaries to steal transaction fees from honest ones, a problem that AMHL’s anonymity techniques also solve.

More recently, Blitz [5] introduced *one-phase* payment channels that support multi-hop payments without a two-phase separation of coin creation and spending, improving performance and reliability. Blitz also achieves stronger anonymity than HTLC, but its notion of anonymity is strictly weaker than the relationship anonymity of AMHL and Fulgor/Rayo. Other relationship-anonymous systems consider powerful adversaries that control most nodes and achieve indistinguishability of concurrent transactions, but Blitz considers local adversaries controlling a single intermediary and limits itself to hiding the rest of the path from this intermediary.

### 3 Our approach

As we argued in Section 1.3, all of these existing solutions share an undesirable reliance on either custom cryptographic constructions or primitives with special properties, like Blitz’s stealth-address signature schemes. This causes inflexibility, difficult implementation, and an inability to respond to cryptanalytic breakthroughs like practical quantum computing.

Astrape is our solution to this problem. We show with a novel design that avoids the zero-knowledge verification paradigm, anonymous and atomic multi-hop transactions can be constructed with nothing but the two building blocks of HTLCs — hashing and signatures. Unlike existing work, no specific assumptions about the structure of the hash function or signature scheme are made, allowing Astrape to be easily ported to different concrete cryptographic primitives and

its security properties to “fall out” from those of the primitives. This also allows Astrape to achieve high performance on commodity hardware using standard cryptographic libraries.

### 3.1 Generalized multi-hop locks

In our discussion of Astrape, we avoid describing the concrete details of a specific payment channel network and cryptocurrency. Instead, we introduce an abstract model — generalized multi-hop locks. This model readily generalizes to different families of payment channel networks.

We model a *sender*,  $U_0$ , sending money to a *receiver*,  $U_n$ , through intermediaries  $U_1, \dots, U_{n-1}$ . We assume a “source routing” model, where the graph of all valid payment paths in the network is publicly known and the sender can choose any valid path to the recipient. After an *initialization* phase where the sender may securely communicate parameters to each hop, each user  $U_i$  where  $i < n$  creates a coin and notifies  $U_{i+1}$ . This coin is simply a contract  $\ell_{i+1}$  known as a *lock script*, that essentially releases money to  $U_{i+1}$  given a certain key  $k_{i+1}$ . We call this lock the *right lock* of  $U_i$  and the *left lock* of  $U_{i+1}$ .

Finally, the payment completes once all coins created in the protocol have been unlocked and spent by fulfilling their lock scripts. Typically, this happens through a chain reaction where the recipient’s left lock  $\ell_n$  is unlocked, allowing  $U_{n-1}$  to unlock its left lock, etc.

Formally, we model a GMHL over a set of participants  $U_i$  as a tuple of four PPT algorithms  $\mathbb{L} = (\text{Init}, \text{Create}, \text{Unlock}, \text{Vf})$ , defined as follows:

**Definition 1.** A GMHL  $\mathbb{L} = (\text{Init}, \text{Create}, \text{Unlock}, \text{Vf})$  consists of the following polynomial-time protocols:

1.  $\langle s_0^I, \dots, (s_n^I, k_n) \rangle \Leftarrow \langle \text{Init}_{U_0}(1^\lambda, U_1, \dots, U_n), \text{Init}_{U_1}, \dots, \text{Init}_{U_n} \rangle$ : the initialization protocol, started by the sender  $U_0$ , that takes in a security parameter  $1^\lambda$  and the identities of all hops  $U_i$  and returns an initial state  $s_i^I$  to all users  $U_i$ . Additionally, the recipient receives a key  $k_n$ .
2.  $\langle (\ell_i, s_{i-1}^R), (\ell_i, s_i^L) \rangle \Leftarrow \langle \text{Create}_{U_{i-1}}(s_{i-1}^I), \text{Create}_{U_i}(s_i^I) \rangle$ : the coin-creating protocol run between two adjacent hops  $U_{i-1}$  and  $U_i$ , creating the “coin sent from  $U_{i-1}$  to  $U_i$ ”. This includes a lock representation  $\ell_i$  as well as additional state on both ends — unlocking the lock represented by  $\ell_i$  releases the money.
3.  $k_i \Leftarrow \text{Unlock}_{U_i}(\ell_{i+1}, (s_i^I, s_i^L, s_i^R), k_{i+1})$ : the coin-spending protocol, run by each intermediary  $U_i$  where  $i < n$ , obtains a valid unlocking key  $k_i$  for the “left lock”  $\ell_i$  given its “right lock”  $\ell_{i+1}$ , its unlocking key  $k_{i+1}$  (already verified by  $\text{Vf}$  below), and  $U_i$ ’s internal state.
4.  $\{0, 1\} \Leftarrow \text{Vf}(\ell, k)$ : given a lock representation  $\ell$  and an unlocking key  $k$ , return 1 iff the  $k$  is a valid solution to the lock  $\ell$

As an example, a formalization of HTLC in the GMHL model can be found in the extended version of this paper [12, App. A].

*Generalizability to non-PCN systems.* We note here that GMHL makes no mention of typical PCN components such as channels, the blockchain, etc. This is because GMHL is actually agnostic of *how* exactly the locks are evaluated and enforced. In a typical PCN, these locks will be executed within bilateral payment channels, falling back to a public blockchain for final settlement.

However, other enforcement mechanisms can be used. Notably, all the locks could simply be contracts directly executing on a blockchain. In this way, any anonymous PCN formulated in the GMHL model is equivalent to a specification for a provably anonymous *on-chain, multi-hop coin tumbling service* that can anonymize entirely on-chain payments by routing them through multiple intermediaries.

*Comparison to existing work.* GMHL is an extension of *anonymous multi-hop locks*, the model used in the eponymous paper by Malavolta et al. [19]. In particular, AMHL defines an anonymous PCN construction in terms of the operations KGen, Setup, Lock, Rel, Vf, four of which correspond to GMHL functions.

Although AMHL’s model is useful, we could not use it verbatim. This is largely because AMHL’s original definition [19] also included its security and privacy properties, while we wish to be able to use the same framework in a purely *syntactic* fashion to discuss PCNs with other security and anonymity goals.

Nevertheless, GMHL can be considered as AMHL, reworded and used in a more general context. As we will soon see, Astrape’s desired security and privacy properties are actually very similar to those of AMHL, though we will consider other systems formulated in the GMHL framework along the way. Astrape can be considered an alternative implementation of the same “anonymous multi-hop locks” [19] construct.

### 3.2 Security and execution model

Now that we have a model to discuss PCN constructions, we can discuss our security model, as well as a model of the GMHL execution environment in which Astrape will execute.

*Active adversary.* We use a similar adversary model to that of AMHL [19]. That is, we model an adversary  $\mathcal{A}$  with access to a functionality  $\text{corrupt}(U_i)$  that takes in the identifier of any user  $U_i$  and provides the attacker with the complete internal state of  $U_i$ . The adversary will also see all incoming and outgoing communication of  $U_i$ .  $\text{corrupt}(U_i)$  will also give the adversary active control of  $U_i$ , allowing it to impersonate  $U_i$  when communicating with other participants.

*Anonymous communication.* We assume there is a secure and anonymous message transmission functionality  $\mathcal{F}_{\text{anon}}$  that allows any participant to send messages to any other participant. Messages sent by an honest (non-corrupted) user with  $\mathcal{F}_{\text{anon}}$  hide the identity of the sender and cannot be read by the adversary, although the adversary may arbitrarily delay messages.

There are many ways of implementing  $\mathcal{F}_{\text{anon}}$ , the exact choice of which is outside the scope of this paper. One solution recommended by existing work [18, 19] is an onion-routing circuit constructed over the same set of users  $U_i$ , constructed with a provably private protocol like Sphinx [9]. Public networks such as Tor may also be used to implement  $\mathcal{F}_{\text{anon}}$ .

*Exposed lock activity.* In contrast to communication, *lock activity* — the content of all locks being created, as well as the unlocking keys during unlocking — is not secure. This is because in practice, lock activity often happens on public media like blockchains. We pessimistically assume that the adversary can see all lock activity, while a non-adversary only sees lock activity concerning locks that it sends and receives.

*Liveness and timeouts.* We assume that every coin lock  $\ell_i$  comes with an appropriate timeout that will return money to  $U_{i-1}$  (i.e., able to be unlocked by a signature from  $U_{i-1}$  after the timeout) if  $U_i$  does not take action. We also assume that each left lock  $\ell_i$ 's deadline is at least  $\delta$  later than that of the right lock  $\ell_{i+1}$ , where  $\delta$  is an upper bound on network latency between honest parties, even under disruption by the adversary. In the most common setting of a PCN consisting of bilateral payment channels backed by a blockchain, this is essentially a blockchain censorship-resistance assumption. With a liveness assumption, we can then omit timeout handling from the description of the protocol, in line with related work (such as AMHL [19]).

*Infallible lock execution.* We formulate Astrape in the GMHL model, and assume the existence of a mechanism that will guarantee that cryptocurrency locks are always correctly executed in the face of arbitrary adversarial activity. In practice, both bilateral payment channels falling back to a general-purpose public blockchain (like Ethereum) and direct use of this blockchain are good approximations of this mechanism.

*Lock functionality.* We assume that inside our on-chain contracts we are able to use at least the following operations:

- *Concatenation*, producing a bitstring  $x||y$  of length  $|n+m|$  from two bitstrings  $x, y$ , where  $x$  has length  $n$  and  $y$  has length  $m$ .
- *Bitwise XOR*, producing a bitstring  $x \oplus y$  from two bitstrings  $x, y$

as well as the cryptographic hash function  $H$  defined below. An implication of this assumption is that PCNs on blockchains with highly restricted scripting languages, like Bitcoin, cannot use Astrape.

*Cryptographic assumptions.* One of Astrape's main goals is to make minimal cryptographic assumptions. We assume only:



- *Generic cryptographic hash function.* We assume a hash function  $H$ , modeled as a random oracle for the purpose of security proofs, producing  $\lambda$  bits of output, where  $1^\lambda$  is the security parameter. We use the random oracle both as a pseudorandom function and as a commitment scheme, which is well known [7] to be secure.
- *Generic signature scheme.* We assume a secure signature scheme that allows for authenticated communication between any two users  $U_i$  and  $U_j$ .

### 3.3 Security and privacy goals

Against the adversary we described above, we want to achieve the following security and privacy objectives:

*Relationship anonymity.* Given two simultaneous payments between different senders  $S_{\{0,1\}}$  and receivers  $R_{\{0,1\}}$  with payment paths of the same length intersecting at the same position at at least one honest intermediate user, an adversary corrupting all of the other intermediate users cannot determine, with probability non-negligibly better than  $1/2$  (guessing), whether  $S_0$  paid  $R_0$  and  $S_1$  paid  $R_1$ , or  $S_0$  paid  $R_1$  and  $S_1$  paid  $R_0$ . This is an established standard for anonymity in payment channels [18, 19] and is analogous to similar definitions for anonymous communication [21, 6]. It is important to note that the adversary is not allowed to corrupt the sender — senders always know who they are sending money to.

*Balance security.* For an honest user  $U_i$ , if its right lock  $\ell_i$  is unlocked,  $U_i$  must always be able to unlock its left lock  $\ell_{i-1}$  even if all other users are corrupt. Combined with the timeouts mentioned in our security model, this guarantees that no intermediary node can lose money even if everybody else conspires against it.

*Wormhole resistance.* We need to be immune to the *wormhole attack* on PCNs, where malicious intermediaries steal fees from other intermediaries. The reason why is rather subtle [19], but for our purposes this means that given an honest sender and an honest intermediary  $U_{i+1}$ ,  $\ell_i$  cannot be spent by  $U_i$  until  $U_{i+1}$  spends  $\ell_{i+1}$ . Intuitively, this prevents honest intermediaries from being “left out”.

## 4 Construction

### 4.1 Core idea: balance security + honest-sender anonymity

Unlike existing systems that utilize the mathematical properties of some cryptographic construction to build a secure and anonymous primitive, Astrape is constructed out of two separate *broken* constructions, both of which use boring cryptography and are straightforward to describe:

- **XorCake**, which has relationship anonymity but lacks balance security if the sender  $U_0$  is malicious

- **HashOnion** which has balance security, but *loses relationship anonymity* in the Unlock phase. That is, an adversary limited only to observing Init and Create cannot break relationship anonymity, but an adversary observing Unlock can.

The key insight here is that if we can combine XorCake and HashOnion in such a way to ensure that HashOnion’s Unlock phase *can only reveal information when the sender is malicious*, we obtain a system, Astrape, that has both relationship anonymity and balance security. This is because the definition of relationship anonymity assumes an honest sender: if the sender is compromised, it can always simply tell the adversary the identity of its counterparty, breaking anonymity trivially. It is important to note that such a composition *does not in any way weaken anonymity* compared to existing “up-front anonymity” systems like AMHL, even in the most pessimistic case.<sup>3</sup>

We now describe XorCake and HashOnion, and their composition into Astrape.

## 4.2 XorCake: anonymous but insecure against malicious senders

Let us first describe XorCake’s construction. XorCake is an extremely simple construction borrowed from “multi-hop HTLC”, a building block of Fulgor [18]. It has relationship anonymity, but not balance security against malicious senders.

Recall that in GMHL, the sender ( $U_0$ ) wishes to send a sum of money to the recipient ( $U_n$ ) through  $U_1, \dots, U_{n-1}$ . At the beginning of the transaction, the sender samples  $n$  independent  $\lambda$ -bit random strings  $(r_1, \dots, r_n)$ . Then, for all  $i \in 1, \dots, n$ , she sets  $n$  values  $s_i = H(r_i \oplus r_{i+1} \oplus \dots \oplus r_n)$ , where  $H$  is a secure hash function. That is,  $s_i$  is simply the hash of the XOR of all the values  $r_j$  for  $j \geq i$ .  $U_0$  then uses the anonymous channel  $\mathcal{F}_{\text{anon}}$  to provide  $U_n$  the values  $(r_n, s_n)$  and all the other  $U_i$  with  $(r_i, s_i, s_{i+1})$ .

Then, for each pair of neighboring nodes  $(U_i, U_{i+1})$ ,  $U_i$  sends  $U_{i+1}$  a coin encumbered by a regular HTLC  $\ell_{i+1}$  asking for the preimage of  $s_{i+1}$ .  $U_n$  knows how to unlock  $\ell_n$ , and the solution would let  $U_{n-1}$  unlock  $\ell_{n-1}$ , and so on. That is, each lock  $\ell_i$  is simply an HTLC contract asking for the preimage of  $s_i$ .

In the extended version [12, App. A], we give the formal definition of XorCake in the GMHL framework.

XorCake by itself satisfies relationship anonymity. A full proof is available in the Fulgor paper from which XorCake was borrowed [19], but intuitively this is because  $r_i$  will be randomly distributed over the space of possible strings because  $H$  behaves like a random oracle. This means that unlike in HTLC, no two nodes  $U_i$  and  $U_j$  can deduce that they are part of the same payment path unless they are adjacent.

<sup>3</sup> In a sense then, Astrape has “pseudo-optimistic” anonymity. Its design superficially suggests an optimistic construction with an anonymous “happy path” and a non-anonymous “unhappy path”, but the latter non-anonymity is illusory — the sender can always prevent the “unhappy” path from deanonymizing the transaction even if all other parties are malicious.

*State-mismatch attack.* Unfortunately, *XorCake* does not have balance security. Consider a malicious sender who follows the protocol correctly, except for sending an incorrect  $r_i$  to  $U_i$ . (Note that  $U_i$  cannot detect that  $r_i$  is incorrect given a secure hash function.) Then, when  $\ell_{i+1}$  is unlocked,  $\ell_i$  cannot be spent! In an actual PCN such as the Lightning Network, all coins “left” of  $U_i$  will time out, letting the money go back to  $U_0$ .  $U_0$  paid  $U_n$  with  $U_i$ ’s money instead of her own. We call this the “state-mismatch attack”, and because of it, *XorCake* is not a viable PCN construction on its own. In Fulgor, *XorCake* was combined with out-of-band zero-knowledge proofs of the correctness of  $r_i$ , but as we will see shortly, *Astrape* can dispense with them.

### 4.3 HashOnion: secure but eventually non-anonymous

We now present HashOnion, a PCN construction that has balance security but not relationship anonymity. Note that unlike HTLC, HashOnion’s non-anonymity stems entirely from information leaked in the Unlock phase, a property we will leverage to build a fully anonymous construction combining HashOnion and *XorCake*.

At the beginning of the transaction,  $U_0$  generates random values  $s_i$  for  $i \in \{1, \dots, n\}$ , then “onion-like” values  $x_i$ , recursively defined as  $x_i = H(s_i || x_{i+1})$ ,  $x_n = H(s_n || 0^\lambda)$ .

Essentially,  $x_i$  is a value that commits to all  $s_j$  where  $j \geq i$ . An onion-like commitment is used rather than a “flat” commitment (say, a hash of all  $s_j$  where  $j \geq i$ ) as it is crucial for balance security, as we will soon see.

For all intermediate nodes  $0 < i < n$ , the sender sends  $(x_{i+1}, s_i)$  to  $U_i$ , while for the destination, the sender sends  $s_n$ . Then, each intermediary  $U_{i-1}$  sends to its successor  $U_i$  a lock  $\ell_i$ , which can be only be unlocked by some  $k_i = (s_i, \dots, s_n)$  where  $H(s_i || H(s_{i+1} || H(\dots H(s_n || 0^\lambda)))) = x_i$ .  $U_{i-1}$  constructs this lock from the  $x_i$  it received from the sender. Finally, during the unlock phase, the recipient  $U_n$  solves  $\ell_n$  with  $k_n = (s_n)$ . This allows each  $U_i$  to spend  $\ell_i$ , completing the transaction.

For balance security, we need to show that with a solution  $k_{i+1} = (s_{i+1}, \dots, s_n)$  to  $\ell_{i+1}$ , and  $s_i$ , we can always construct a solution to  $\ell_i$ . This is obvious: we just add  $s_i$  to the solution:  $k_i = (s_i, s_{i+1}, \dots, s_n)$ .

One subtle problem is that  $U_i$  needs to make sure that its left lock is actually the correct  $\ell_i$  and not some bad  $\ell'_i$  parameterized over some  $x'_i \neq H(s_i || x_{i+1})$ . Otherwise, its right lock might get unlocked with a solution that does not let it unlock its left lock. Fortunately, this is easy: given  $s_i, x_{i+1}$  from the sender,  $U_i$  can just check that its left lock, parameterized over some  $x_i$ , matches  $x_i = H(s_i || x_{i+1})$  before sending out  $\ell_{i+1}$  (parameterized with  $x_{i+1}$ ) to the next hop. Thus, every user can make sure that if its right lock is unlocked, so can its left lock, so balance security holds.

We also see that although the unlocking procedure breaks relationship anonymity by revealing all the  $s_i$ , before the unlock happens, HashOnion does have relationship anonymity. This is because the adversary cannot connect the different  $x_i$  as long as one  $s_i$  remains secret — that of the one honest intermediary.

#### 4.4 Securing XorCake+HashOnion

We now move on to composing XorCake and HashOnion. We do so by creating a variant of HashOnion that embeds XorCake and recognizes an *inconsistency witness*. That is, this variant of HashOnion will unlock only when given a combination of values that proves an attempt by the sender to execute a state-mismatch attack for XorCake.

To construct such a lock, after generating the XorCake parameters,  $U_0$  creates  $n$   $\lambda$ -bit values  $x_i$  recursively:

$$x_n = o_n, \quad x_i = H(\overbrace{r_i || s_i || s_{i+1}}^{\text{XorCake parameters}} || o_i || x_{i+1})$$

where  $o_i$  is a random nonce sampled uniformly from all possible  $\lambda$ -bit values.<sup>4</sup> The intuition here is that  $x_i$  *commits* to all the information  $U_0$  would give to all hops  $U_j$  where  $j \geq i$ .

Afterwards, the sender then uses  $\mathcal{F}_{\text{anon}}$  to send  $(o_i, x_i, x_{i+1})$ , in addition to the XorCake parameters  $(r_i, s_i, s_{i+1})$ , to every hop  $i$ . Every hop  $U_i$  checks that all the parameters are consistent with each other.

We next consider what will happen if the sender attempts to fool an intermediate hop  $U_i$  with a state-mismatch attack.  $U_{i+1}$  would unlock its left lock  $\ell_{i+1}$  by giving  $k_{i+1}$  where  $H(k_{i+1}) = s_{i+1}$  but  $H(r_i \oplus k_{i+1}) \neq s_i$ . This then causes  $U_i$  to fail to unlock its left lock.

But this attempt allows  $U_i$  to generate a cryptographic witness verifiable to anybody knowing  $x_i$ :  $\lambda$ -bit values  $k_{i+1}, r_i, s_i, s_{i+1}, o_i, x_{i+1}$  where:

$$H(k_{i+1}) = s_{i+1}, H(r_i \oplus k_{i+1}) \neq s_i, H(r_i || s_i || s_{i+1} || o_i || x_{i+1}) = x_i$$

This inconsistency witness proves that the preimage of  $s_{i+1}$  XOR-ed with  $r_i$  does not equal the preimage of  $s_i$ , demonstrating that the values given to  $U_i$  are inconsistent and that  $U_0$  is corrupt. Since  $U_{i-1}$  knows  $x_i$ ,  $U_i$  can therefore prove that it was a victim of a state-mismatch attack to  $U_{i-1}$ .

Since  $x_i$  commits to *all* XorCake initialization states “rightwards” of  $U_i$ ,  $U_i$ , in cooperation with  $U_{i-1}$ , can also produce a witness that  $U_{i-2}$  can verify using  $x_{i-1}$ . This is simply a set of  $\lambda$ -bit values  $k_{i+1}, r_{i-1}, s_{i-1}, r_i, s_i, s_{i+1}, x_i, o_i, o_{i-1}, x_{i+1}$  where:

$$H(k_{i+1}) = s_{i+1}, H(r_i \oplus k_{i+1}) \neq s_i, \\ H(r_i || s_i || s_{i+1} || o_i || x_{i+1}) = x_i, H(r_{i-1} || s_{i-1} || s_i || o_{i-1} || x_i) = x_{i-1}$$

We can clearly extend this idea all the way back to  $U_1$  — given a witness demonstrating a state-mismatch attack against  $U_i$ ,  $U_{i-1}$  can verify the witness and generate a similar one verifiable by  $U_{i-2}$ , and so on. This forms the core construction that Astrape uses to fix XorCake’s lack of balance security.

#### 4.5 Complete construction

We now present the complete construction of Astrape, as formalized in Figure 1 within the GMHL framework. Note that we use the notation  $\text{Tag}[x_1, \dots, x_n]$  to

<sup>4</sup>  $||$  denotes concatenation. In our case, it is possible to unambiguously separate concatenated values, since we only ever concatenate  $\lambda$ -bit values.

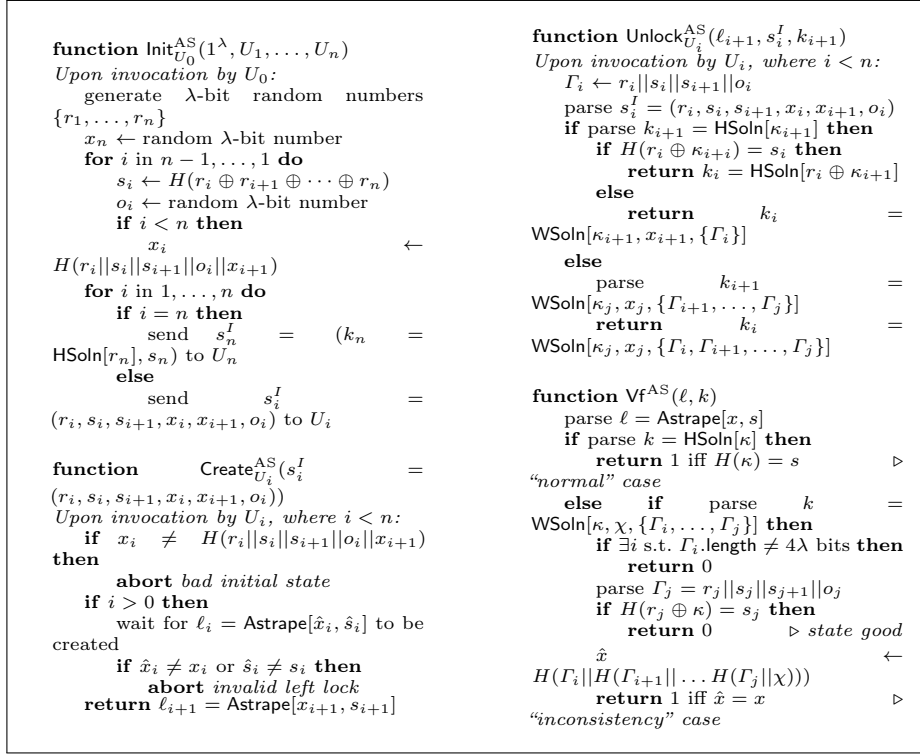


Fig. 1: Astrape as a GMHL protocol

represent a  $n$ -tuple of values with an arbitrary “tag” that identifies the type of value.

*Initialization.* In the first phase, represented as `Init` in GMHL, the sender  $U_0$  first establishes communication to the  $n$  hops  $U_1, \dots, U_n$ , the last one of which is the receiver. When talking to intermediaries and the recipient,  $U_0$  uses our abstract functionality  $\mathcal{F}_{\text{anon}}$ .

The sender then generates random  $\lambda$ -bit strings  $(r_1, \dots, r_n)$  and  $(o_1, \dots, o_n)$ , deriving  $s_i = H(r_i \oplus r_{i+1} \oplus \dots \oplus r_n)$  and  $x_i = H(r_i || s_i || s_{i+1} || o_i || x_{i+1})$ ;  $x_n = H(o_n)$ . She then sends to each intermediate hop  $U_i$  the tuple  $s_i^I = (r_i, s_i, s_{i+1}, x_i, x_{i+1}, o_i)$  and gives the last hop  $U_n$  the initial state  $s_n^I = (r_n, s_n)$  and the unlocking key  $k_n = \text{HSoln}[r_n]$ .

*Creating the coins.* We now move on to `Create`, where all the coins are initially locked.  $U_0$  then sends  $U_1$  a coin encumbered with a lock represented as  $\ell_1 = \text{Astrape}[x_1, s_1]$ . When each hop  $U_i$  receives a correctly formatted coin from its previous hop  $U_{i-1}$ , it sends the next hop  $U_{i+1}$  a coin with a lock  $\ell_{i+1} = \text{Astrape}[x_{i+1}, s_{i+1}]$ . Note that  $U_i$  checks whether its left lock is consistent with

the parameters it received from  $U_0$ ; this ensures that when  $U_i$ 's right lock unlocks later,  $U_i$  can always construct a solution for its left lock. If the checks fail, `Create` aborts, and all of the locks will eventually time out (see Section 3.2), returning money to the sender.

As specified in `Vf`, each of these locks  $\ell_i$  can be spent either through solving a XorCake-type puzzle to find the preimage of  $s_i$  (the “normal” case) or by presenting an inconsistency witness with a HashOnion-type witness demonstrating  $x_i$ 's commitment to inconsistent data (the “inconsistency” case). After all the transactions with `Astrape`-encumbered coins are sent,  $U_n$  can finally claim its money, triggering the next phase of the protocol.

*Unlocking the coins.* The last step is `Unlock`. After receiving the final coin from  $U_{n-1}$ , the recipient unlocks its lock  $\ell_n$  by providing to `Vf` the preimage of the HTLC puzzle:  $k_n = \text{HSoln}[r_n]$  — this is the only way an honest recipient can claim the money in a payment originating from an honest sender. Each intermediate node  $U_i$  reacts when its right lock  $\ell_{i+1}$  is unlocked with key  $k_{i+1}$ :

- If  $U_{i+1}$  solved the HTLC puzzle with  $k_{i+1} = \text{HSoln}[\kappa_{i+1}]$ , construct  $\kappa_i = r_i \oplus \kappa_{i+1}$ 
  - If  $H(\kappa_i) = s_i$ , this means that there is no state-mismatch attack happening. We unlock our left lock with  $k_i = \text{HSoln}[\kappa_i]$ .
  - Otherwise, there must be an attack happening. We construct a witness and create a key that embeds the witness verifiable with  $x_i$ . This gives us  $k_i = \text{WSoln}[\kappa_{i+1}, x_{i+1}, \{\Gamma_i\}]$ , where  $\Gamma_i = r_i || s_i || s_{i+1} || o_i$ .
- Otherwise,  $U_{i+1}$  demonstrated that the sender attempted to defraud some  $U_j$ , where  $j > i$  unlocked  $\ell_{i+1}$  by presenting an inconsistency witness  $k_{i+1} = \text{WSoln}[\kappa_j, x_j, \{\Gamma_{i+1}, \Gamma_{i+2}, \dots, \Gamma_j\}]$ .
  - We can simply construct  $k_i = \text{WSoln}[\kappa_j, x_j, \{\Gamma_i, \Gamma_{i+1}, \dots, \Gamma_j\}]$  where  $\Gamma_i = r_i || s_i || s_{i+1} || o_i$ . This transforms the witness verifiable with  $x_{i+1}$  to a witness verifiable with  $x_i$ .

Note that both cases are covered by `Vf` — it accepts and verifies both “normal” unlocks with `HSoln`-tagged tuples, and “inconsistency” unlocks with `WSoln`. Thus, even though `Astrape` is a composition of XorCake and HashOnion, the final construction fully “inlines” the two into the same flow of initialization, coin creation, and unlocking, with no separate procedure to process inconsistency witnesses. Unlocking continues backwards towards the sender until all the locks created in the previous step are unlocked. We have balance security — node  $U_i$  can unlock its left lock  $\ell_i$  if and only if node  $U_{i+1}$  has unlocked  $\ell_{i+1}$ , so no intermediaries can lose any money.

Security proofs, as well as a discussion on side-channel and griefing attacks, can be found in the extended version [12, §5].

## 5 Blockchain implementation

`Astrape` is easy to implement on blockchains with Turing-complete scripting languages, like Ethereum, as well as layer-2 PCNs such as Raiden built on these

blockchains, but blockchains without Turing-complete scripts involve two main challenges.

First, these blockchains typically do not allow recursion or loops in lock scripts. This means that we cannot directly implement the  $\forall f$  function. Instead, we must “unroll”  $\forall f$  to explicitly check for witnesses to inconsistencies in the parameters given to  $U_i$ ,  $U_{i+1}$ , etc. So for an  $n$ -hop payment the size of every lock script grows to  $\Theta(n)$ . In practice, the mean path length in the Lightning Network is currently around 5 (see our measurements in Section 6.4), and privacy-focused onion routing systems such as Tor or I2P typically use 3 to 5 hops. We believe linear-length script sizes are not a significant concern for Astrape deployment. An Astrape deployment can simply pick an arbitrary maximum for the number of hops supported and achieve reasonable worst-case performance.

The second issue is more serious: some blockchains have so little scripting that Astrape cannot be implemented. Astrape requires an “append-like” operation  $\parallel$  that can take in two bytestrings and combine them in a collision-resistant manner. Unfortunately, the biggest blockchain Bitcoin has disabled all string-manipulation opcodes. Whether an implementation based solely on the 32-bit integer arithmetic that Bitcoin uses is possible is an interesting open question.

## 6 Comparison with existing work

In this section, we compare Astrape with existing PCN constructions. First, we compare Astrape’s design choices and features with that of other systems, showing that it explores a novel design space. Then, we evaluate Astrape’s concrete performance. We compare Astrape’s performance with that of other PCN constructions, both anonymous and non-anonymous. Finally, we explore Astrape’s performance on a real-world network graph from the Lightning Network.

### 6.1 Design comparison

In Table 1, we compare Astrape’s properties with those of existing payment channel networks. We see that except for HTLC, which does not achieve anonymity, all previous PCN networks use cryptographic constructions specialized

Table 1: Comparison of different PCNs

	Topology	Anon <sup>a</sup>	Efficient <sup>b</sup>	Crypto
HTLC	Mesh	No	Yes	Sig. + hash
Tumblebit	Hub	Yes	No	Custom RSA
Bolt	Hub	Yes	Yes	NIZKP
Teechain	Hub	Yes	Yes	Trusted comp.
Fulgor/Rayo	Mesh	Yes	No	ZKP
AMHL <sub>van</sub> <sup>c</sup>	Mesh	Yes	Yes	Homom. OWF
AMHL <sub>ecd</sub>	Mesh	Yes	Yes	ECDSA,
AMHL <sub>sch</sub>	Mesh	Yes	Yes	Homom. enc. Schnorr sigs
Blitz	Mesh	Weak <sup>d</sup>	Yes	SA <sup>e</sup> sig. + hash
Astrape	Mesh	Yes	Yes	Sig. + hash

<sup>a</sup> Relationship anonymity

<sup>b</sup> Roughly comparable performance to HTLC. For example, ZKPs requiring many orders of magnitude more computation time than HTLC are not considered “efficient”.

<sup>c</sup> AMHL is a family of three closely related constructions. We denote by *van*, *ecd*, *sch* the “vanilla”, ECDSA, and Schnorr implementations respectively.

<sup>d</sup> See discussion in Section 2.3.

<sup>e</sup> A “stealth-address” signature scheme; i.e., a signature scheme where any party knowing a public key can generate unlinkably different public keys that correspond to the same private key.

Table 2: Resource usage of different PCN systems ( $n$  hops,  $c$ -byte HTLC contract,  $d$ -byte Astrape contract); AHML variants van,ecd,sch as in Table 1

	<i>Plain HTLC</i>	Fulgor/ Rayo	AMHL	Astrape (Bitcoin Cash)	Astrape
Comput. time (ms)	$< 0.001$	$\approx 200n$	$\approx n$ (van) $\approx 3n$ (ecd) $\approx 3n$ (sch)	$\approx 0.7n$	$\approx 0.25n$
Comm. size (bytes)	$32n$	$1650000n$	$32 + 96n$ (van) $416 + 128n$ (ecd) $256 + 128n$ (sch)	$192n$	$192n$
Lock (bytes)	$32 + c$	$32 + c$	$32 + c$	$108 + 39 \cdot n$	$64 + d$
Unlock, normal case (bytes)	32	32	$32$ (van) / 64	32	32
Unlock, worst case (bytes)	32	32	$32$ (van) / 64	$64 + 128 \cdot n$	$64 + 128 \cdot n$

for their use case. Furthermore, only more recent constructions achieve efficiency comparable to HTLC. It is thus clear that Astrape is the first and only PCN construction that works on all PCN topologies, achieves strong anonymity, and performs at high efficiency, while using the same simple cryptography as HTLC.

## 6.2 Implementation and benchmark setup

To demonstrate the feasibility and performance of our construction, we developed a prototype implementation in the Go programming language. We implemented all the cryptographic constructions of Astrape inside a simulated GMHL model. We used the libsodium library’s implementation of the ed25519 [16] signature scheme and blake2b [3] hash function. In addition, we generated script locks in Bitcoin Cash’s scripting language to illustrate script sizes for scripting languages with no loops. The Bitcoin Cash scripts, written in the higher-level CashScript language, can be found in the extended version [12, App. B].

All tests were done on a machine with a 3.2 GHz Intel Core i7 and 16 GB RAM. Network latency is not simulated, as this is highly application dependent. These conditions are designed to be maximally similar to those under which Fulgor [18] and AMHL [19] were evaluated, allowing us to compare the results directly.

## 6.3 Resource usage

Our first set of tests compares Astrape’s resource usage to that of other PCN constructions. We compare both a simulation of Astrape and a concrete implementation using Bitcoin Cash’s scripting language to traditional HTLC, Fulgor/Rayo, and all three variants of AMHL.

We summarize the results in Table 2, where  $n$  refers to the number of hops,  $c$  to the size of an HTLC contract, and  $d$  to the size of an Astrape contract. We copy results for Fulgor/Rayo [18] and AMHL (ECDSA) [19] from their original sources, which use an essentially identical setup.



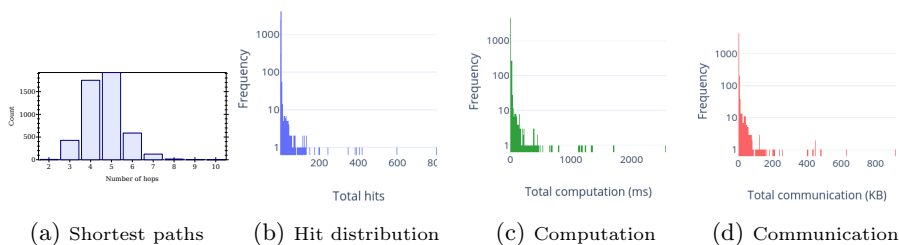


Fig. 2: Overhead distribution

*Computation time.* We measure computation time, with communication and other overhead ignored. The time measured is the sum of the CPU time taken by each hop, for all steps of the algorithm. We note that by eschewing non-standard cryptographic primitives, Astrape achieves lower computation times across the board compared to Fulgor/Rayo and AMHL.

*Communication overhead.* We also measure the communication overhead of each system. This is defined as all the data that needs to be communicated *other than the locks and their opening solutions*. For example, in Astrape, this includes all the setup information sent from  $U_0$ , while in AMHL this includes everything exchanged during the Setup, Lock, and Rel [19] phases. We see that Astrape has by far the least communication overhead of all the anonymous PCN constructions. Note especially the extreme overhead of the zero-knowledge proofs used in Fulgor/Rayo.

*Lock overhead.* The last measure is *per-coin* lock overhead — the size of each lock script (the “lock size”) and that of the information required to unlock it (the “unlock size”). This is a very important component of a system’s resource usage, since lock and unlock sizes directly translate into transaction fees in blockchain cryptocurrencies. Astrape’s performance differs in two important ways.

First of all, Astrape’s Vf function is expressed in a recursive manner. In blockchains like Bitcoin Cash that support neither recursion nor loops in their scripting language, we must “unroll” the Vf implementation. This causes lock sizes to be linear in the number of hops. In blockchains with general-purpose scripting languages, though, lock size is generally constant. Second, the worst-case unlock size is larger for Astrape. When the sender is malicious and all coins have to be spent by invoking HashOnion, we need  $n$  parameters  $(I_1, \dots, I_n)$  to unlock each coin for an  $n$ -hop payment. However, despite this asymptotic disadvantage, we believe that Astrape nevertheless offers competitive lock performance. This is because payment routes are quite short in practice, as we will shortly see.

#### 6.4 Statistical simulation

Finally, we simulate the performance of Astrape on the network graph of the Lightning Network (LN).

*Setup.* We set up a mainnet Lightning Network node using the `lnd` [1] reference implementation. We then used the `lncli describegraph` command to capture the network topology of Lightning Network in February 2021. This gives us a graph of 9566 nodes and 72164 edges. Finally, we randomly sample 5,000 pairs of nodes in the network and calculate the shortest paths between them. This gives us a randomly sampled set of real-life payment paths.

*Path statistics.* As we have previously shown, paths more than 10 hops long still have fairly small overhead even with non-recursive lock scripts, but much longer paths will cause rather large unlock sizes. We examine whether the graph topology will force payments to grow too long; Figure 2a illustrates the distribution of lengths for our 5,000 randomly selected payment paths. On average, a payment path was 5.12 hops long, though the Lightning Network specification allows up to 20 hops. This indicates that shortest payment paths long enough to pose seriously ballooning worst-case overhead are practically nonexistent.

*Total scalability.* One important attribute we wish to explore with the LN topology is the total scalability of the network — how fast can a PCN process transactions as a whole.

To do so, we keep track of how many times each node appears, or is “hit”, in our 5,000 randomly selected payment routes. On average, this is 2.99, but the vast majority of nodes are hit only once, while a few nodes are hit hundreds of times. The distribution of hits is plotted in Figure 2b as a log-linear histogram. We then look at the *distribution of overhead* in the network for both computation and communication. This is by calculating the total computation and communication cost for each node “hit” by the 1,000 random payments, using values from the Bitcoin Cash implementation.

Computation cost is plotted in Figure 2c. We see that the most heavily loaded node in the entire network did around 2,000 ms of computation to process 1,500 transactions. This indicates that the largest hubs in a PCN with the current Lightning Network topology will be able to process around 750 transactions a second per CPU core. Such a throughput is orders of magnitude higher than that of typical blockchains and is within reach of many traditional payment systems. We note that this is only the maximum throughput of a *single CPU core* — in practice hubs likely have multicore machines, and with many hubs the total LN throughput will be many times this number.

Communication cost is plotted in Figure 2d. We pessimistically assume that all payments are settled through HashOnion. Even so, the total network load averages to only about 3.58 KB per node. The largest hubs’ total load still do not exceed 1 MB. This illustrates that the bottleneck is actually computation, not communication.

In summary, we see that Astrape’s worst-case asymptotic performance poses no barriers to the total throughput of an Astrape-powered payment channel network. PCNs can enjoy the superb scalability associated with them just as easily with Astrape-powered privacy and security.

## 7 Conclusion

First-generation payment channel networks and other trust-minimizing intermediarized cryptocurrency payment systems lack strong privacy and security guarantees. Existing research, although solving the privacy and security problems, tend to rely on custom cryptographic primitives that cannot be easily swapped with alternatives based on different computational hardness assumptions.

We presented Astrape, a novel PCN construction that breaks this conundrum. Astrape is the first PCN that achieves relationship anonymity and balance security with only black-box access to generic conventional cryptography. It relies on a general idea of using non-anonymous post-hoc inconsistency witnesses to achieve balance security, while avoiding any information leaks when senders are not corrupt. This allows Astrape to avoid dealing with the zero-knowledge verification used to achieve balance security in existing relationship-anonymous PCNs without sacrificing any anonymity or security properties.

Furthermore, we demonstrate that Astrape is practical to deploy in the real world. Performance is superior on average to existing private PCNs, even on blockchains that are unsuitable for free-form smart contracts. We also showed that Astrape achieves high scalability on a real-world payment channel network graph.

## Acknowledgements

We thank the reviewers and Sherman Chow for helping to improve this paper. We thank NSERC for Discovery Grant RGPIN-07014 and Create 498002-2017. This research was undertaken, in part, thanks to funding from the Canada Research Chairs program.

## References

1. Lightning Network Daemon. <https://github.com/lightningnetwork/lnd> (2019)
2. Blockchain Charts. <https://www.blockchain.com/charts> (2021), accessed on 2022-04-01
3. Aumasson, J.P., Neves, S., Wilcox-O’Hearn, Z., Winnerlein, C.: BLAKE2: simpler, smaller, fast as MD5. In: International Conference on Applied Cryptography and Network Security. pp. 119–135. Springer (2013)
4. Aumayr, L., Maffei, M., Ersoy, O., Erwig, A., Faust, S., Riahi, S., Hostáková, K., Moreno-Sanchez, P.: Bitcoin-compatible virtual channels. In: 2021 IEEE Symposium on Security and Privacy (SP). pp. 901–918. IEEE (2021)
5. Aumayr, L., Monero-Sanchez, P., Maffei, M.: Blitz: Secure Multi-Hop Payments Without Two-Phase Commits. In: 30th USENIX Security Symposium (2021)
6. Backes, M., Kate, A., Manoharan, P., Meiser, S., Mohammadi, E.: Anoa: A framework for analyzing anonymous communication protocols. In: 2013 IEEE 26th Computer Security Foundations Symposium. pp. 163–178. IEEE (2013)
7. Camenisch, J., Drijvers, M., Gagliardoni, T., Lehmann, A., Neven, G.: The wonderful world of global random oracles. In: IACR EuroCrypt. pp. 280–312. Springer (2018)

8. Croman, K., Decker, C., Eyal, I., Gencer, A.E., Juels, A., Kosba, A., Miller, A., Saxena, P., Shi, E., Sirer, E.G., et al.: On scaling decentralized blockchains. In: International Conference on Financial Cryptography and Data Security. pp. 106–125. Springer (2016)
9. Danezis, G., Goldberg, I.: Sphinx: A compact and provably secure mix format. In: 30th IEEE Symposium on Security and Privacy. pp. 269–282. IEEE (2009)
10. Decker, C., Russell, R., Osuntokun, O.: eltoo: A Simple Layer2 Protocol for Bitcoin. <https://blockstream.com/eltoo.pdf> (2018)
11. Decker, C., Wattenhofer, R.: A fast and scalable payment network with bitcoin duplex micropayment channels. In: Symposium on Self-Stabilizing Systems. pp. 3–18. Springer (2015)
12. Dong, Y., Goldberg, I., Gorbunov, S., Boutaba, R.: Astrape: Anonymous Payment Channels with Boring Cryptography. <https://github.com/nullchinchilla/astrape-paper/> (2022)
13. Engelmann, F., Kopp, H., Kargl, F., Glaser, F., Weinhardt, C.: Towards an economic analysis of routing in payment channel networks. In: Proceedings of the 1st Workshop on Scalable and Resilient Infrastructures for Distributed Ledgers. pp. 1–6 (2017)
14. Green, M., Miers, I.: Bolt: Anonymous payment channels for decentralized currencies. In: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security. pp. 473–489. ACM (2017)
15. Heilman, E., Alshenibr, L., Baldimtsi, F., Scafuro, A., Goldberg, S.: Tumblebit: An untrusted bitcoin-compatible anonymous payment hub. In: Network and Distributed System Security Symposium (2017)
16. Josefsson, S., Liusvaara, I.: Edwards-Curve Digital Signature Algorithm (EdDSA). RFC 8032 (Jan 2017). <https://doi.org/10.17487/RFC8032>, <https://rfc-editor.org/rfc/rfc8032.txt>
17. Lai, R.W.F., Cheung, H.K.F., Chow, S.S.M., So, A.M.C.: Another look at anonymous communication. In: Phan, R.C.W., Yung, M. (eds.) Paradigms in Cryptology – Mycrypt 2016. Malicious and Exploratory Cryptology. pp. 56–82. Springer International Publishing, Cham (2017)
18. Malavolta, G., Moreno-Sanchez, P., Kate, A., Maffei, M., Ravi, S.: Concurrency and privacy with payment-channel networks. In: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security. pp. 455–471. ACM (2017)
19. Malavolta, G., Moreno-Sanchez, P., Schneidewind, C., Kate, A., Maffei, M.: Anonymous multi-hop locks for blockchain scalability and interoperability. In: NDSS (2019)
20. McCorry, P., Möser, M., Shahandasti, S.F., Hao, F.: Towards bitcoin payment networks. In: Australasian Conference on Information Security and Privacy. pp. 57–76. Springer (2016)
21. Pfitzmann, A., Hansen, M.: A terminology for talking about privacy by data minimization: Anonymity, unlinkability, undetectability, unobservability, pseudonymity, and identity management (2010)
22. Tairi, E., Moreno-Sanchez, P., Maffei, M.: A2L: Anonymous atomic locks for scalability and interoperability in payment channel hubs. In: 42nd IEEE Symposium on Security and Privacy (2021)
23. Van Wirdum, A.: How the Lightning Network Layers Privacy on Top of Bitcoin. <https://bitcoinmagazine.com/articles/how-the-lightning-network-layers-privacy-on-top-of-bitcoin-1482183775> (2016), accessed on 2022-04-01
24. Yousaf, H., Kappos, G., Piotrowska, A., Kanjalkar, S., Delgado-Segura, S., Miller, A., Meiklejohn, S.: An empirical analysis of privacy in the lightning network (2021)