Atomic Cross-Chain Swaps

Maurice Herlihy
Computer Science Department
Brown University
Providence, Rhode Island
maurice.herlihy@gmail.com

ABSTRACT

An *atomic cross-chain swap* is a distributed coordination task where multiple parties exchange assets across multiple blockchains, for example, trading bitcoin for ether.

An *atomic swap* protocol guarantees (1) if all parties conform to the protocol, then all swaps take place, (2) if some coalition deviates from the protocol, then no conforming party ends up worse off, and (3) no coalition has an incentive to deviate from the protocol.

A cross-chain swap is modeled as a directed graph $\mathcal{D}$, whose vertexes are parties and whose arcs are proposed asset transfers. For any pair $(\mathcal{D}, L)$, where $\mathcal{D} = (V, A)$ is a strongly-connected directed graph and $L \subset V$ a feedback vertex set for $\mathcal{D}$, we give an atomic cross-chain swap protocol for $\mathcal{D}$, using a form of hashed timelock contracts, where the vertexes in L generate the hashlocked secrets. We show that no such protocol is possible if $\mathcal{D}$ is not strongly connected, or if $\mathcal{D}$ is strongly connected but L is not a feedback vertex set. The protocol has time complexity $O(\text{diam}(\mathcal{D}))$ and space complexity (bits stored on all blockchains) $O(|A|^2)$.

ACM Reference Format:

Maurice Herlihy. 2018. Atomic Cross-Chain Swaps. In *PODC '18: ACM Symposium on Principles of Distributed Computing, July 23-27, 2018, Egham, United Kingdom*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3212734.3212736>

1 MOTIVATION

Carol wants to sell her Cadillac for bitcoins. Alice is willing to buy Carol's Cadillac, but she wants to pay in an "alt-coin" cryptocurrency. Fortunately, Bob is willing to trade alt-coins for bitcoins. Alice, Bob, and Carol need to arrange a three-way swap: Alice will transfer her alt-coins to Bob, Bob will transfer his bitcoins to Carol, and Carol will transfer title of her Cadillac to Alice¹. Of course, no one trusts anyone else. How can we devise a protocol that ensures that if all parties behave rationally, in his or her own self-interest, then all assets are exchanged, but if some parties behave irrationally, then no rational party will end up worse off?

In many blockchains, assets are transferred under the control of so-called *smart contracts* (or just *contracts*), scripts published on the blockchain that establish and enforce conditions necessary to transfer an asset from one party to another. For example, let

¹Naturally, they live in a state that records automobile titles in a blockchain.

This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2018 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-5795-1/18/07.

<https://doi.org/10.1145/3212734.3212736>

$H(\cdot)$ be a cryptographic hash function. Alice might place her alt-coins in escrow by publishing on the alt-coin blockchain a smart contract with *hashlock* h and *timelock* t . Hashlock h means that if Bob sends the contract a value s , called a *secret*, such that $h = H(s)$, then the contract irrevocably transfers ownership of those alt-coins from Alice to Bob. Timelock t means that if Bob fails to produce that secret before time t elapses, then the escrowed alt-coins are refunded to Alice.

Here is a simple protocol for Alice, Bob, and Carol's three-way swap, illustrated in Figures 1 and 2. Let Δ be enough time for one party to publish a smart contract on any of the blockchains, or to change a contract's state, and for the other party to detect the change.

- Alice creates a secret s , $h = H(s)$, and publishes a contract on the alt-coin blockchain with hashlock h and timelock 6Δ in the future, to transfer her alt-coins to Bob.
- When Bob confirms that Alice's contract has been published on the alt-coin blockchain, he publishes a contract on the Bitcoin blockchain with the same hashlock h but with timelock 5Δ in the future, to transfer his bitcoins to Carol.
- When Carol confirms that Bob's contract has been published on the Bitcoin blockchain, she publishes a contract on the automobile title blockchain with the same hashlock h , but with timeout 4Δ in the future, to transfer her Cadillac's title to Alice.
- When Alice confirms that Carol's contract has been published on the title blockchain, she sends s to Carol's contract, acquiring the title and revealing s to Carol.
- Carol then sends s to Bob's contract, acquiring the bitcoins and revealing s to Bob.
- Bob sends s to Alice's contract, acquiring the alt-coins and completing the swap.

What could go wrong? If any party halts while contracts are being deployed, then all contracts eventually time out and trigger refunds. If any party halts while contracts are being triggered, then only that party ends up worse off. For example, if Carol halts without triggering her contract, then Alice gets the Cadillac and Bob gets a refund, so Carol's misbehavior harms only herself.

The order in which contracts are deployed matters. If Carol were to post her contract with Alice before Bob posts his contract with Carol, then Alice could take ownership of the Cadillac without paying Carol.

Timelock values matter. If Carol's contract with Bob were to expire at the same time as Bob's contract with Alice, then Carol could reveal s to collect Bob's bitcoins at the very last moment, leaving Bob no time to collect his alt-coins from Alice.

Figure 1: Atomic cross-chain swap: deploying contracts

Figure 2: Atomic cross-chain swap: triggering arcs

What if parties behave irrationally? If Alice (irrationally) reveals s before the first phase completes, then Bob can take Alice's alt-coins, and perhaps Carol can take Bob's bitcoins, but Alice will not get her Cadillac, so only she is worse off.

A *atomic swap protocol* guarantees (1) if all parties conform to the protocol, then all swaps take place, (2) if some parties deviate from the protocol, then no conforming party ends up worse off², and (3) no coalition has an incentive to deviate from the protocol. Alice, Bob, and Carol's swapping adventure suggests broader questions: when are atomic cross-chain swaps possible, how can we implement them, and what do they cost?

While swapping digital assets is the immediate motivation for this study, atomic cross-chain swap protocols have other possible applications. *Sharding* [7] splits one blockchain into many for better load-balancing and scalability. Most of the time, activities on different shards proceed independently. When they cannot, an atomic swap protocol can coordinate needed cross-chain updates. In a decentralized distributed system, *upgrades* from one software version to another, or from one data schema to another, could benefit from atomic cross-chain swaps. An atomic swap protocol can be thought of as a trust-free, Byzantine-hardened form of *distributed commitment* [26]. An atomic cross-chain swap is a special case of a *distributed atomic transaction* [25], although not all atomic transactions can be expressed as cross-chain swaps.

²Other than the inconvenience of having assets temporarily locked up.

Cross-chain swaps are well-known to the blockchain community [4, 6, 9, 20, 21, 27], but to our knowledge, this is the first systematic analysis of the theory underlying such protocols. We make the following contributions. A cross-chain swap is modeled as a directed graph (digraph) $\mathcal{D}$, whose vertexes are parties and whose arcs are proposed asset transfers. For any pair $(\mathcal{D}, L)$, where $\mathcal{D} = (V, A)$ is a *strongly-connected* digraph and $L \subset V$ a *feedback vertex set* for $\mathcal{D}$, we give an atomic cross-chain swap protocol using a form of hashed timelock contracts, where the vertexes in L , called *leaders*, generate the hashlocked secrets. (Vertexes that are not leaders are *followers*.) We also show that no such protocol is possible if $\mathcal{D}$ is not strongly connected, or if $\mathcal{D}$ is strongly connected but the set of leaders L is not a feedback vertex set. The protocol has time complexity $O(\text{diam}(\mathcal{D}))$ and communication complexity (bits published on blockchains) $O(|A| \cdot |L|)$.

2 MODEL

2.1 Digraphs

A *directed graph* (or *digraph*) $\mathcal{D}$ is a pair (V, A) , where V is a finite set of *vertexes*, and A is a finite set of ordered pairs of distinct vertexes called *arcs*. We use $V(\mathcal{D})$ for $\mathcal{D}$'s set of vertexes, and $A(\mathcal{D})$ for its set of arcs. An arc (u, v) has *head* u and *tail* v . An arc *leaves* its head and *enters* its tail. An arc (u, v) *enters* a set of vertexes $W \subseteq V$ if $u \notin W$ and $v \in W$, and similarly for leaving.

A digraph $\mathcal{C}$ is a *subdigraph* of $\mathcal{D}$ if $V(\mathcal{C}) \subseteq V(\mathcal{D})$, $A(\mathcal{C}) \subseteq A(\mathcal{D})$ and every arc in $A(\mathcal{C})$ has both its head and tail in $V(\mathcal{C})$.

A *path* p in $\mathcal{D}$ is a sequence of vertexes $(u_0, \dots, u_\ell)$ such that $u_0, \dots, u_{\ell-1}$ are distinct. Path p has *length* ℓ , denoted by $|p|$. If v is a vertex, and $(u_0, \dots, u_\ell)$ a path that does not include the arc (v, u_0) , then $v + p$ denotes the path $(v, u_0, \dots, u_\ell)$. For vertexes u, v , $D(u, v)$ is the length of the longest path from u to v in $\mathcal{D}$.

A path $(u_0, \dots, u_\ell)$ is a *cycle* if $u_0 = u_\ell$. A digraph is *acyclic* if it has no cycles. Vertex v is *reachable* from vertex u if there is a path from u to v . $\mathcal{D}$'s *diameter* $\text{diam}(\mathcal{D})$ is the length of the longest path from any vertex to any other. $\mathcal{D}$ is *connected* if its underlying graph is connected, and *strongly connected* if, for every pair u, v of distinct vertexes in $\mathcal{D}$, u is reachable from v and v is reachable from u . A *feedback vertex set* is a subset of V whose deletion leaves $\mathcal{D}$ acyclic.

The *transpose* $\mathcal{D}^T$ is the digraph obtained from $\mathcal{D}$ by reversing all arcs. If $\mathcal{D}$ is strongly connected, so is $\mathcal{D}^T$, and any feedback vertex set for $\mathcal{D}$ is also a feedback vertex set for $\mathcal{D}^T$.

2.2 Blockchains and Smart Contracts

For our purposes, a *blockchain* is a distributed service that allows clients to publish transactions to a publicly-readable, tamper-proof distributed ledger. Our analysis is independent of the particular blockchain algorithm. We assume a timing model where there is a known duration Δ long enough for one party to publish a contract to a blockchain, and for a second party to confirm that the contract has been published.

The owner of an asset (such as a unit of cryptocurrency or an automobile title) can create a smart contract to transfer ownership of that asset to a *counterparty* if specified conditions are met. A contract is *published* when its creator places it on a blockchain ledger. Once a contract is published, it is irrevocable: neither the

contract's creator nor any other party can remove the contract nor tamper with its terms.

A *rational* party acts in its own self-interest, deviating from a protocol only if it is profitable to do so. Rational parties can collude with one another to disadvantage other parties. An *irrational* party may deviate from a protocol even if it is not profitable to do so. Parties may behave irrationally out of spite, because they were hacked, or because they profit in ways not foreseen by the protocol designers.

Blockchain protocols typically require parties to have public and private keys. We use $\text{sig}(x, v)$ to denote the result of v using its private key to sign x .

3 SWAP DIGRAPHS AND GAMES

A cross-chain swap is given by a digraph $\mathcal{D} = (V, A)$, where each vertex in V represents a party, and each arc in A represents a proposed asset transfer from the arc's head to its tail via a shared blockchain. (We assume without loss of generality that $\mathcal{D}$ is connected, because a disconnected digraph can be treated as multiple swaps.) Henceforth, we use *party* and *vertex*, *blockchain* and *arc*, interchangeably, depending on whether we emphasize roles or digraph structure.

In the terminology of game theory, a swap $\mathcal{D}$ is a *cooperative game*, organized so that if all parties follow the protocol, each transfer on each arc happens. Each possible outcome is given by a subdigraph $\mathcal{E} = (V, A')$ of $\mathcal{D}$. If a proposed transfer $(u, v) \in A$ is also in $(u, v) \in A'$, then that transfer happened. For short, we say arc (u, v) was *triggered*.

A protocol is a *strategy* for playing a game: a set of rules that determines which step a party takes at any stage of a game. To model real-world situations where multiple parties are secretly controlled by a single adversary, the swap game is *cooperative*: parties can form *coalitions* where coalition members commit to a common strategy.

Figure 3: Partial order of protocol outcomes

Here are the outcomes for a party v , organized into classes. For brevity, each class has a shorthand name.

- The party acquires assets without paying: at least one arc entering v is triggered, but no arc leaving v is triggered (FREE RIDE).

- The party acquires assets while paying less than expected: all arcs entering v are triggered, but at least one arc leaving v is not triggered (DISCOUNT).
- The party swaps assets as expected: all arcs entering and leaving v are triggered (DEAL).
- No assets change hands: no arc entering or leaving v is triggered (NoDEAL).
- The party pays without acquiring all expected assets: at least one arc entering v is not triggered, and at least one arc leaving v is triggered (UNDERWATER).

Payoffs for a coalition $C \subset V$ are defined by replacing v with C in the definitions above.

The protocol design incorporates certain conservative assumptions about parties' preferences. The protocol's preferred outcome is for all conforming parties to end with outcome DEAL. In the presence of failures or deviation, however, it is acceptable for conforming parties to end with outcome NoDEAL, the *status quo*, rendering them no worse off. Furthermore, each party is assumed to prefer DEAL to NoDEAL, because otherwise it would not have agreed to the swap in the first place. It follows that each party prefers any FREERIDE outcome to NoDEAL, because it acquires additional assets "for free", without relinquishing any assets of its own. Similarly, each party prefers any DISCOUNT outcome to DEAL, since that party acquires the same assets in both outcomes, but relinquishes fewer in DISCOUNT outcomes. For these reasons, DEAL, NoDEAL, DISCOUNT, and FREERIDE are all considered acceptable outcomes for conforming parties if the protocol execution is unable to complete because of failures or adversarial behavior.

We consider the remaining UNDERWATER outcomes to be unacceptable to conforming parties. It is possible that in some idiosyncratic cases, a party may actually prefer particular UNDERWATER outcomes to NoDEAL. For example, a party with three entering arcs and one leaving arc may be willing to relinquish its asset in return for acquiring only two out of three of the entering arcs' assets. We leave the design of protocols that make such fine distinctions to future work.

Definition 3.1. A swap protocol $\mathbb{P}$ is *uniform* if it satisfies:

- If all parties follow $\mathbb{P}$, they all finish with payoff DEAL.
- If any coalition cooperatively deviates from $\mathbb{P}$, no conforming party finishes with payoff UNDERWATER.

A uniform protocol is not useful if rational parties will not follow it. A swap protocol is a *strong Nash equilibrium strategy* if no coalition improves its payoff when its members cooperatively deviate from that protocol.

Definition 3.2. A swap protocol $\mathbb{P}$ is *atomic* if it is both uniform and a strong Nash equilibrium strategy.

This definition formalizes the notion that if all parties are rational, all swaps happen, but if some parties are irrational, the rational parties will never end up worse off. Recall that a *conforming* party follows the protocol, while a *deviating* party does not.

LEMMA 3.3. *If $\mathcal{D}$ is strongly connected, then any uniform swap protocol $\mathbb{P}$ is atomic.*

PROOF. If a deviating coalition $C \subset V$ achieves a better payoff than DEAL, then that payoff is either in FREERIDE or DISCOUNT.

It follows that some arc that enters C is triggered, and some arc that leaves C is untriggered. Moreover, if any arc that enters C is untriggered, then all arcs that leave C are untriggered.

A conforming party $v \notin C$ cannot end up UNDERWATER, so if an arc entering v is untriggered, then every arc leaving v must be untriggered, and if an arc leaving v is triggered, then every arc entering v must be triggered.

Let (c, v) be an untriggered arc leaving C . Since v is conforming, every arc leaving v is untriggered. Because $\mathcal{D}$ is strongly connected, there is a path $(v, v_0), (v_0, v_1), \dots, (v_k, c_0)$ where each $v_i \notin C$, and $c_0 \in C$. By a simple inductive argument, each arc in this path is untriggered, so the arc (v_k, c_0) that enters C is untriggered, so every arc leaving C must be untriggered, and some arc entering C must be triggered.

Let (v, c) be a triggered arc entering C . Since v is conforming, every arc entering v is triggered. Because $\mathcal{D}$ is strongly connected, there is a path $(c_1, v_0), (v_0, v_1), \dots, (v_k, c)$ where each $v_i \notin C$, and $c_1 \in C$. By a simple inductive argument, each arc in this path is triggered, so the arc (c_1, v_0) leaving C is triggered, contradicting the fact that every arc leaving C is untriggered. $\square$

LEMMA 3.4. *If $\mathcal{D}$ is not strongly connected, then no uniform swap protocol is atomic.*

PROOF. Because $\mathcal{D}$ is not strongly connected, it contains vertexes x, y such that y is reachable from x , but not vice-versa. Let Y be the set of vertexes reachable from y , and X the rest: $X = V \setminus Y$. X is non-empty because it contains x . Because y is reachable from x , there is at least one arc from X to Y , but no arcs from Y to X .

Coalition X can improve its payoff by triggering all arcs between vertexes in X , but no arcs from X to Y , yielding payoff FREERIDE for X , since it triggers strictly fewer arcs leaving X , without affecting any arcs entering X . In fact, the payoff for each individual vertex in X is either the same or better than DEAL. $\square$

We have just proved:

THEOREM 3.5. *A uniform swap protocol for $\mathcal{D}$ is atomic if and only if $\mathcal{D}$ is strongly connected.*

Informally, if $\mathcal{D}$ is not strongly connected, then rational parties will deviate from any uniform protocol. In practice, such a swap would never be proposed, because the parties in X would never agree to a swap with the free riders in Y . Henceforth, $\mathcal{D}$ is assumed strongly connected.

We remark that Theorem 3.5 relies on the implicit technical assumption that all value transfers are explicitly represented on some blockchain. This theorem would be falsified, for example, if Carol responds to learning Alice's secret by sending a large drone to drop her Cadillac in the middle of Alice's driveway, without ever recording that transfer in a shared blockchain. We will assume that if swaps have off-chain consequences, as they typically do, that those consequences are explicitly recorded in the form of blockchain updates.

```

1  contract Swap {
2      Asset asset;           /* asset to be transferred or refunded */
3      Digraph digraph;       /* swap digraph */
4      address[] leaders;     /* leaders */
5      address party;         /* transfer asset from */
6      address counterparty;  /* transfer asset to */
7      uint[] timelock;       /* vector of timelocks */
8      uint[] hashlock;       /* vector of hashlocks */
9      bool[] unlocked;       /* which hashlocks unlocked? */
10     uint start;            /* protocol starting time */
11     /* constructor */
12     function Swap (Asset _asset; /* asset to be transferred or refunded */
13                   Digraph _digraph; /* swap digraph */
14                   address[] _leaders; /* leaders */
15                   address _party; /* transfer asset from */
16                   address _counterparty; /* transfer asset to */
17                   uint[] _timelock; /* vector of timelocks */
18                   uint[] _hashlock; /* vector of hashlocks */
19                   uint _start /* protocol starting time */
20                   ) {
21         asset = _asset; /* copy */
22         party = _party; counterparty = _counterparty; /* copy */
23         timelock = _timelock; hashlock = _hashlock; /* copy */
24         unlocked = [false, ..., false]; /* all unlocked */
25     }

```

Figure 4: Swap contract (part one)

4 AN ATOMIC SWAP PROTOCOL

4.1 Hashlocks and Hashkeys

In a simple two-party swap, each party publishes a contract that assumes temporary control of that party's asset. This *hashed time-lock contract* [5] stores a pair (h, t) , and ensures that if the contract receives the matching secret s , $h = H(s)$, before time t has elapsed, then the contract is *triggered*, irrevocably transferring ownership of the asset to the counterparty. If the contract does not receive the matching secret before time t has elapsed, then the asset is *refunded* to the original owner. For multi-party cross-chain swaps, we will need to extend these notions in several ways.

In the three-way swap recounted earlier, each arc had a single hashlock and a single timeout. Timeouts were assigned so that the timeout on each arc entering a follower v was later by at least Δ than the timeout on each arc leaving v . This gap ensures that if any arc leaving v is triggered, then v has time to trigger every entering arc.

If a swap digraph has only one leader, $\hat{v}$, then the subdigraph of its followers is acyclic. As in our three-way swap example, the hashlock on arc (u, v) can be given timeout $(\text{diam}(\mathcal{D}) + D(v, \hat{v}) + 1) \cdot \Delta$, where $D(v, \hat{v})$ is the length of the longest path from v to the unique leader $\hat{v}$. (See left-hand side of Figure 6.)

This formula does not work if a swap digraph has more than one leader, because the subdigraph of any leader's followers has a cycle,

and it is not possible to assign timeouts across a cycle in a way that guarantees a gap of at least Δ between entering and leaving arcs. (See right-hand side of Figure 6.)

Instead, for general digraphs, we must replace timed hashlocks with a more general mechanism, one that assigns different timeouts to different paths. Pick a set $L = \{v_0, \dots, v_\ell\}$ of vertexes, called *leaders*, forming a feedback vertex set for $\mathcal{D}$. Each leader v_i generates a secret s_i and hashlock value $h_i = H(s_i)$, yielding a *hashlock vector* $(h_0, \dots, h_\ell)$, which is assigned to every arc.

A *hashkey* for h on arc (u, v) is a triple (s, p, σ) , where s is the secret $h = H(s)$, p is a path $(u_0, \dots, u_k)$ in $\mathcal{D}$ where $u_0 = v$ and u_k is the leader who generated s , and

$$\sigma = \text{sig}(\dots \text{sig}(s, u_k), \dots, u_0),$$

the result of having each party in the path sign s . A hashkey (s, p, σ) *times out* at time $(\text{diam}(\mathcal{D}) + |p|) \cdot \Delta$ after the start of the protocol. That hashkey *unlocks* h on (u, v) if it is presented before it times out. An arc is *triggered* when all of its hashlocks are unlocked. A hashlock has *timed out* on an arc when all of its hashkeys on that arc have timed out. Figure 7 shows partial hashkeys for a two-leader swap digraph.

```

26 function unlock (int i, uint s, Path path, Sig sig) {
27     require (msg.sender == counterpart); /* only from counterpart */
28     if (now < start + (diam(digraph) + |path|) * Δ /* hashkey still valid? */
29         && hashlock[i] == H(s) /* secret correct? */
30         && isPath(path, digraph, leader[i], counterpart) /* path valid? */
31         && verifySigs(sig, s, path) { /* signatures valid? */
32         unlocked[i] = true;
33     }
34 }
35 function refund () {
36     require (msg.sender == party); /* only from party */
37     if (any hashlock unlocked and timed out) {
38         transfer asset to party;
39         halt;
40     }
41 }
42 function claim () {
43     require (msg.sender == counterpart); /* only from counterpart */
44     if (every hashlock unlocked) {
45         transfer asset to counterpart;
46         halt;
47     }
48 }
49 }

```

Figure 5: Swap contract (part two)

Figure 6: *A* is the leader, *B* and *C* followers. Timeouts can be assigned when the follower subdigraph is acyclic (left) but not when it is cyclic (right)

4.2 Market Clearing

For simplicity, assume the swap digraph is constructed by a (possibly centralized) market-clearing service, which perhaps communicates with the parties through its own blockchain. The clearing service is not a trusted party, because the parties can check the consistency of the clearing service's responses.

Each party creates a secret s and matching hashlock $h = H(s)$. It sends the clearing service its hashlock, along with an offer characterizing the swaps it is willing to make. The service combines these offers, and publishes a swap digraph $\mathcal{D} = (V, A)$, a vector $L \subset V$ of *leaders* forming a feedback vertex set, a vector of those leaders' hashlocks $h_0, \dots, h_\ell$, and a *starting time* T , at least Δ in the future.

Figure 7: Hashkey paths for arcs of two-leader digraph

Figure 8: Concurrent contract propagation for two-leader digraph

If all parties conform to the protocol, all contracts will be triggered before $T + 2 \cdot \text{diam}(\mathcal{D}) \cdot \Delta$, but if some parties deviate, the conforming parties' assets will be refunded by then.

4.3 Contracts

Figures 4 and 5 show pseudocode³ for a hashed timelock swap contract. A smart contract resembles an object in an object-oriented programming language, providing *long-lived state* (Lines 2-9), a *constructor* to initialize that state (Lines 12-25), and one or more *functions* to manage that state (Lines 26-48).

The contract's long-lived state records the asset to be transferred or refunded (Line 2), the digraph $\mathcal{D}$ (Line 3), the digraph's set of leaders (Line 4), the party transferring the asset (Line 5), the counterparty receiving the asset (Line 6), a vector of timelocks

(Line 7), a vector of hashlocks (Line 8), and a Boolean unlocked vector marking which hashlocks have been unlocked (Line 9).

When the contract is initialized, its constructor copies the fields provided into the contract's long-lived state (Lines 21-23) and sets each entry in unlocked to *false* (Line 24).

The `unlock()` function (Line 26), callable only by the counterparty (Line 27), takes an index i , a secret s_i , a path p , and the signature sig . The hashlock h_i is unlocked if

- the current time is less than $T + (\text{diam}(\mathcal{D}) + |p|) \cdot \Delta$ (Line 28),
- $h_i = H(s_i)$ (Line 29),
- p is a path in $\mathcal{D}$ from the counterparty to the leader who generated s_i (Line 30), and
- the signature is the result of signing s_i by the parties in p (Line 31)

The `refund()` function (Line 35), callable only by the party, refunds the asset to the party if any unlocked hashlock has timed out. The `claim()` function (Line 42), callable only by the counterparty

³ This pseudocode is based loosely on the popular Solidity programming language for smart contracts [24].

(Line 36), transfers the asset to the counterparty if all hashlocks have been unlocked.

4.4 Pebble Games

We analyze the protocol using two variations on a simple pebble game. We are given a strongly-connected digraph $\mathcal{D} = (V, A)$, and a vertex feedback set $L \subset V$ of *leaders*.

In the *lazy* pebble game, start by placing pebbles on the arcs leaving each leader. Place new pebbles on the arcs leaving vertex v when there is a pebble on *every* arc entering v .

In the *eager* pebble game, start by placing a single pebble on one vertex z . Place new pebbles on the arcs leaving v when there is a pebble on *any* arc entering v . Both games continue until no more pebbles can be placed.

LEMMA 4.1. *In the lazy game, every arc in $\mathcal{D}$ eventually has a pebble.*

PROOF. Suppose by way of contradiction, the game stops in a state where an arc (u, v) has no pebble. There must be a pebble-free arc (u', u) entering u , because otherwise the game would have placed a pebble on (u, v) . Continuing in this way, build a longer and longer pebble-free path until it becomes a pebble-free cycle. But leaders form a feedback vertex set, so every cycle in $\mathcal{D}$ includes a leader, and the arcs leaving that leader have pebbles placed in the first step. $\square$

LEMMA 4.2. *In the eager game, every arc in $\mathcal{D}$ eventually has a pebble.*

PROOF. Suppose by way of contradiction, the game stops in a state where an arc (u, v) has no pebble. Because $\mathcal{G}$ is strongly connected, there is a path from z to v . Since z has a pebble and v does not, there is an arc (w, w') on that path where w has a pebble but w' does not, so w' will get a pebble in the next step, contradicting the hypothesis that the game has stopped. $\square$

Suppose there is a worst-case delay Δ between when the last pebble is placed on any arc entering v , and when the last pebble is placed on any arc leaving v .

LEMMA 4.3. *In both pebble games, every arc will have a pebble in time at most $\text{diam}(\mathcal{D}) \cdot \Delta$ from when the game started.*

PROOF. For the lazy game, it is enough to show that in each interval of time Δ , the longest pebble-free path shrinks by one. At any time after the first step, let $a_0, \dots, a_k$ be a pebble-free path of maximal length. That path cannot be a cycle, because then it would include a leader, who would have placed a pebble on a_0 in the first step. It follows that every arc entering the head of a_0 must have a pebble, because otherwise we could construct a longer pebble-free path. By hypothesis, within time Δ , a_0 will have a pebble, and the path will have shrunk by one.

For the eager game, it is enough to observe that in each interval of time Δ , for every vertex v , the number of unpebbled vertexes in every path from z to w shrinks by one. Because $\mathcal{D}$ is strongly connected, such a path always exists. $\square$

COROLLARY 4.4. *Under the stated timing assumptions, for both games, every arc has a pebble within time $\text{diam}(\mathcal{D}) \cdot \Delta$.*

4.5 The Protocol

There are two phases. In Phase One, instances of the Swap contract (Figures 4 and 5) are propagated through $\mathcal{D}$, starting at the leaders. Each time a party observes that a contract has been published on an entering arc, it verifies that contract is a correct swap contract, and abandons the protocol otherwise.

Here is the Phase-One protocol for leaders:

- (1) Publish a contract on every arc leaving the leader, then
- (2) wait until contracts have been published on all arcs entering the leader.

Here is the protocol for followers:

- (1) wait until correct contracts have been published on all arcs entering the vertex, then
- (2) publish a contract on every arc leaving the vertex.

Figure 7 shows how contracts are propagated in a swap digraph with two leaders.

In Phase Two, the parties disseminate secrets via hashkeys. While contracts propagate in the direction of the arcs, from party to counterparty, hashkeys propagate in the opposite direction, from counterparty to party. Informally, each party is motivated to trigger the contracts on entering arcs to acquire the assets controlled by those contracts.

We now trace how the secret s_i generated by leader v_i is propagated. At the start of the phase, v_i calls $\text{unlock}(s_i, v_i, \text{sig}(s_i, v_i))$ at each entering arc's contract (here, the function's arguments are the hashkey, and v_i is a degenerate path). The first time any other party v observes that hashlock h_i on a leaving arc's contract has been unlocked by a call to $\text{unlock}(s_i, p, \sigma)$, it calls $\text{unlock}(s_i, v + p, \text{sig}(\sigma, v))$ at each entering arc's contract. The propagation of s_i is complete when h_i has either timed out or has been unlocked on all arcs.

LEMMA 4.5. *If all parties conform to the protocol, then every arc has a contract within time $\text{diam}(\mathcal{D}) \cdot \Delta$ of when the protocol started.*

PROOF. Phase One is an instance of the lazy pebble game on $\mathcal{D}$, so the claim follows from Lemmas 4.1 and 4.3. $\square$

LEMMA 4.6. *If all parties conform to the protocol, then every arc's contract is triggered within time $2 \cdot \text{diam}(\mathcal{D}) \cdot \Delta$ of when the protocol started.*

PROOF. Each secret's dissemination is an instance of the eager pebble games on $\mathcal{D}^T$, the transpose digraph. The secrets are disseminated in parallel. $\square$

THEOREM 4.7. *If all parties conform to the protocol, then every contract is triggered within time $2 \cdot \text{diam}(\mathcal{D}) \cdot \Delta$ of when the protocol started.*

The deadline $2 \cdot \text{diam}(\mathcal{D}) \cdot \Delta$ bounds the time assets can be held in escrow when things go wrong. In practice, one would expect actual running times to be shorter.

There is a simple optimization that ensures that Phase Two completes in constant time when all parties conform to the protocol. We use a shared blockchain, perhaps that of the market-clearing service, as a broadcast medium. Each leader v_i publishes its secret s_i on the shared blockchain, and each follower monitors that blockchain, triggering its entering arcs when it learns the secret. (Logically, we

create an arc from each follower directly to that leader.) Unfortunately, while this broadcasting blockchain can “short-circuit” the Phase Two protocol, it cannot replace it, because a deviating leader might refrain from publishing the secret on that blockchain, but publish it on others. (Miller *et al.* [18] propose a similar optimization for the Lightning network.)

LEMMA 4.8. *If hashlock h times out on any arc entering a conforming v , then h must have timed out on every arc leaving v .*

PROOF. Suppose h was triggered on (v, w) by hashkey (s, p, σ) . If v does not appear in p , then $v + p$ is a path from v to the leader, and v can immediately trigger h on (u, v) using the hashkey $(s, v + p, \text{sig}(\sigma, v))$, which has not timed out. If v appears in p , then v has already received (and signed) a hashkey that triggers h on (u, v) . $\square$

THEOREM 4.9. *No conforming party ends up UNDERWATER.*

PROOF. Assume by way of contradiction that some conforming party v ends up UNDERWATER: a leaving arc (v, w) has a triggered contract, but an entering arc (u, v) does not and will not.

First, arc (u, v) must have a contract. Suppose v is a leader. Since (v, w) has been triggered, v has revealed its secret through a hashkey. But a leader issues hashkeys in Phase Two only after a contract has been published on every entering arc during Phase One. Suppose instead v is a follower. Since (v, w) has been triggered, one of the arcs leaving v has a contract. But in Phase One, a follower publishes a contract on a leaving arc only after contracts have been published on all of its incoming arcs.

Since (u, v) has a contract, one of that arc’s hashlocks must have timed out. By Lemma 4.8, the arc (v, w) must also have timed out, a contradiction. $\square$

THEOREM 4.10. *For $\mathcal{D} = (V, A)$ with leaders $L \subset V$, the space complexity, measured as the number of bits stored on all blockchains, is $O(|A|^2)$.*

PROOF. There are $|A|$ contracts, one on each arc, each with a copy of the digraph $\mathcal{D}$, which requires $O(|A|)$ storage. $\square$

Finally, any atomic cross-chain swap protocol using hashed time-locks must assign secrets to a feedback vertex set.

LEMMA 4.11. *In any uniform hashed timelock swap protocol, no follower v can publish a contract on an arc leaving v before contracts have been published on all arcs entering v .*

PROOF. If follower v has a contract on arc (v, w) but no contract on arc (u, v) , then the parties other than v could collude to trigger the contract on (v, w) , while refusing to publish a contract on (u, v) , leaving v UNDERWATER. $\square$

THEOREM 4.12. *In any uniform swap protocol based on hashed time-locks, the set L of leaders is a feedback vertex set in $\mathcal{D}$.*

PROOF. Suppose, instead, there is a uniform swap protocol where the leaders do not form a vertex feedback set.

At any step in the protocol, the *waits-for* digraph W is the subdigraph of $\mathcal{D}^T$ where (v, u) is an arc of W if (u, v) has no published contract. Informally, Lemma 4.11 implies that v must be waiting for u to publish a contract on (u, v) before u can publish any contracts on its own outgoing arcs. In the initial state, if $\mathcal{D} \setminus L$ contains a

cycle, so does W . At each protocol step, a follower v can publish a contract on a leaving arc only if v has indegree zero in the current *waits-for* digraph. But no vertex on a cycle in the *waits-for* digraph will ever have indegree zero, a contradiction. $\square$

4.6 Single-Leader Digraphs

As noted, in the common special case where a swap digraph needs only one leader, we can replace hashkeys with simple timeouts, reducing message sizes and eliminating the need for digital signatures. In the following, let $\mathcal{D}$ be a swap digraph with a single leader $\hat{v}$ with hashlock h .

LEMMA 4.13. *If each arc (u, v) has timeout $(\text{diam}(\mathcal{D}) + D(v, \hat{v}) + 1) \cdot \Delta$, then for every conforming $v \neq \hat{v}$, the timeout on each arc (u, v) is later by at least Δ than the timeout on each arc (v, w) .*

PROOF. Let p be the longest path from w to the leader $\hat{v}$. Because the subdigraph of followers is acyclic, $v + p$ is a path of length $D(w, \hat{v}) + 1$ from v to $\hat{v}$, so $D(v, \hat{v}) \geq D(w, \hat{v}) + 1$. $\square$

LEMMA 4.14. *For a single-leader digraph using timeouts, if hashlock h times out on any arc entering a conforming v , then h must have timed out on every arc leaving v .*

PROOF. By Lemma 4.14, once h is triggered on (v, w) , v has time at least Δ to trigger (u, v) . $\square$

From this point on, the bounds on running time and proofs of safety for the single-leader-using-timeouts protocol are essentially the same as for the general protocol.

5 REMARKS

We have seen that single-leader swap digraphs do not require hashkeys and digital signatures, only timeouts. Is there a way to reduce the use of digital signatures in the general case?

Finding a minimal feedback vertex set for $\mathcal{D}$ is NP-complete [15], although there exists an efficient 2-approximation [3].

The protocol is easily extended to a model where there may be more than one arc from one vertex to another, so-called *directed multi-graphs* [2], reflecting the situation where Alice wants to transfer assets on distinct blockchains to Bob.

The swap protocol is still vulnerable to a weak denial-of-service attack where an adversarial party repeatedly proposes an attractive swap, and then fails to complete the protocol, triggering refunds, but temporarily rendering assets inaccessible. We leave for future work the question whether one could require parties to post bonds, and following a failed swap, examine the blockchains to determine who was at fault (by failing to execute an enabled transition).

An interesting open problem is the extent to which this swap protocol can be modified to provide better privacy, analogous to the way the Bolt network [12] improves on Lightning.

As noted, some parties may be willing to accept certain UNDERWATER outcomes rejected by the swap protocol presented here. Future work might investigate protocols where parties are endowed with customized objective functions to provide finer-grained control which outcomes are acceptable.

The swap protocol can be made recurrent by having the leaders distribute the next round’s hashlocks in Phase Two of the previous round. If swaps are recurrent, then it would be useful to conduct

swaps *off-chain* as much as possible, similar to the way that Lightning [22] and Raiden [19] networks support off-chain transactions for bitcoin and ERC20 tokens.

One limitation of the swap protocol presented here is the assumption that the swap digraph, its leaders, and their hashlocks are common knowledge among the participants. Future work might address constructing and propagating this information dynamically.

6 RELATED WORK

The use of hashed timelock contracts for two-party cross-chain swaps is believed to have emerged from an on-line discussion forum in 2016 [4, 20]. There is open-source code [6, 9, 21] for two-party cross-chain swap protocols between selected currencies, and proposals for applications using swaps [27].

Off-chain payment networks [8, 12, 19, 22] circumvent the scalability limits of existing blockchains by conducting multiple transactions off the blockchain, eventually resolving final balances through a single on-chain transaction. The Revive network [16] rebalances off-chain networks in a way that ensures that compliant parties do not end up worse off. These algorithms also use hashed timelock contracts, but they address a different set of problems.

Multi-party swaps arise when matching kidney donors and recipients. A transplant recipient with an incompatible donor can swap donors to ensure that each recipient obtains a compatible organ. A number of algorithms [1, 10, 13] have been proposed for matching donors and recipients. Shapley and Scarf [23] consider the circumstances under which certain kinds of swap markets have strong equilibriums. Kaplan [14] describes a polynomial-time algorithm that given a set of proposed swaps, constructs a swap digraph if one exists. These papers and many others focus on “the clearing problem”, roughly analogous to constructing a swap digraph, but not on how to execute those swaps on blockchains.

The *fair exchange* problem [11, 17] is a precursor to the atomic cross-chain swap problem. Alice has a digital asset Bob wants, and vice-versa, and at the end of the protocol, either Alice and Bob have exchanged assets, or they both keep their assets. In the absence of blockchains, trusted, or semi-trusted third parties are required, but roles of those trusted parties can be minimized in clever ways.

A *atomic cross-chain transaction* is a distributed task where a *sequence* of exchanges occurs at each blockchain. An atomic cross-chain swap is an atomic cross-chain transaction, but not vice-versa, because not all transactions can be expressed as swaps. In our original example, Alice could not borrow bitcoins from Bob to pay Carol, because then Alice would have to execute two steps in sequence (borrow, then spend) instead of executing a single swap. A better understanding of atomic cross-chain transactions is the subject of future work.

REFERENCES

- [1] David J. Abraham, Avrim Blum, and Tuomas Sandholm. 2007. Clearing Algorithms for Barter Exchange Markets: Enabling Nationwide Kidney Exchanges. In *Proceedings of the 8th ACM Conference on Electronic Commerce (EC '07)*. ACM, New York, NY, USA, 295–304. <https://doi.org/10.1145/1250910.1250954>
- [2] J. Bang-Jensen and G. Gutin. 2001. *Digraphs: Theory, Algorithms, and Applications*. Springer. <https://books.google.com/books?id=d3gpAQAAAJ>
- [3] Ann Becker and Dan Geiger. 1996. Optimization of Pearl's method of conditioning and greedy-like approximation algorithms for the vertex feedback set problem. *Artificial Intelligence* 83, 1 (1996), 167–188. [https://doi.org/10.1016/0004-3702\(95\)00004-6](https://doi.org/10.1016/0004-3702(95)00004-6)
- [4] bitcoinwiki. [n. d.]. Atomic cross-chain trading. https://en.bitcoin.it/wiki/Atomic_cross-chain_trading. As of 9 January 2018.
- [5] bitcoinwiki. [n. d.]. Hashed Timelock Contracts. Retrieved 8 January 2018 from https://en.bitcoin.it/wiki/Hashed_Timelock_Contracts
- [6] Sean Bowe and Daira Hopwood. [n. d.]. Hashed Time-Locked Contract transactions. Retrieved 9 January 2018 from <https://github.com/bitcoin/bips/blob/master/bip-0199.mediawiki>
- [7] Vitalik Buterin. [n. d.]. On sharding blockchains. Retrieved 8 January 2018 from <https://github.com/ethereum/wiki/wiki/Sharding-FAQ>
- [8] Christian Decker and Roger Wattenhofer. 2015. A Fast and Scalable Payment Network with Bitcoin Duplex Micropayment Channels. In *Stabilization, Safety, and Security of Distributed Systems*, Andrzej Pelc and Alexander A. Schwarzmann (Eds.). Springer International Publishing, Cham, 3–18.
- [9] DeCred. [n. d.]. Decred cross-chain atomic swapping. Retrieved 8 January 2018 from <https://github.com/decred/atomicswap>
- [10] John P. Dickerson, David F. Manlove, Benjamin Plaut, Tuomas Sandholm, and James Trimble. 2016. Position-Indexed Formulations for Kidney Exchange. *CoRR abs/1606.01623* (2016). arXiv:1606.01623 <http://arxiv.org/abs/1606.01623>
- [11] Matthew K. Franklin and Gene Tsudik. 1998. Secure Group Barter: Multi-party Fair Exchange with Semi-Trusted Neutral Parties. In *Financial Cryptography*.
- [12] Matthew Green and Ian Miers. 2016. Bolt: Anonymous Payment Channels for Decentralized Currencies. *Cryptology ePrint Archive*, Report 2016/701. <https://eprint.iacr.org/2016/701>.
- [13] Zhipeng Jia, Pingzhong Tang, Ruosong Wang, and Hanrui Zhang. 2017. Efficient Near-optimal Algorithms for Barter Exchange. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems (AAMAS '17)*, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 362–370. <http://dl.acm.org/citation.cfm?id=3091125.3091181>
- [14] Randy M. Kaplan. 2011. An improved algorithm for multi-way trading for exchange and barter. *Electronic Commerce Research and Applications* 10, 1 (2011), 67–74. <https://doi.org/10.1016/j.elerap.2010.08.001> Special Section: Service Innovation in E-Commerce.
- [15] Richard M. Karp. 1972. Reducibility Among Combinatorial Problems. In *Proceedings of a symposium on the Complexity of Computer Computations, held March 20-22, 1972, at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York*. 85–103. <http://www.cs.berkeley.edu/~luca/cs172/karp.pdf>
- [16] Rami Khalil and Arthur Gervais. 2017. Revive: Rebalancing Off-Blockchain Payment Networks. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (CCS '17)*. ACM, New York, NY, USA, 439–453. <https://doi.org/10.1145/3133956.3134033>
- [17] Silvio Micali. 2003. Simple and Fast Optimistic Protocols for Fair Electronic Exchange. In *Proceedings of the Twenty-second Annual Symposium on Principles of Distributed Computing (PODC '03)*. ACM, New York, NY, USA, 12–19. <https://doi.org/10.1145/872035.872038>
- [18] Andrew Miller, Iddo Bentov, Ranjit Kumaresan, and Patrick McCorry. 2017. Sprites: Payment Channels that Go Faster than Lightning. *CoRR abs/1702.05812* (2017). arXiv:1702.05812 <http://arxiv.org/abs/1702.05812>
- [19] Raiden Network. [n. d.]. What is the Raiden Network? Retrieved 26 January 2018 from <https://raiden.network/101.html>
- [20] T. Nolan. [n. d.]. Atomic swaps using cut and choose. Retrieved 9 January 2018 from <https://bitcointalk.org/index.php?topic=1364951>
- [21] The Komodo Organization. [n. d.]. The BarterDEX Whitepaper: A Decentralized, Open-Source Cryptocurrency Exchange, Powered by Atomic-Swap Technology. Retrieved 9 January 2018 from <https://supernet.org/en/technology/whitepapers/BarterDEX-Whitepaper-v0.4.pdf>
- [22] J. Poon and T. Dryja. 2016. The Bitcoin Lightning Network: Scalable Off-Chain Instant Payments. <https://lightning.network/lightning-network-paper.pdf> As of 29 December 2017.
- [23] Lloyd Shapley and Herbert Scarf. 1974. On cores and indivisibility. *Journal of Mathematical Economics* 1, 1 (1974), 23–37. <https://EconPapers.repec.org/RePEc:eee:mateco:v:1:y:1974:i:1:p:23-37>
- [24] Solidity documentation. [n. d.]. <http://solidity.readthedocs.io/en/latest/index.html>.
- [25] Gerhard Weikum and Gottfried Vossen. 2001. *Transactional Information Systems: Theory, Algorithms, and the Practice of Concurrency Control and Recovery*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- [26] Wikipedia. [n. d.]. Two-phase commit protocol. Retrieved 18 May 2018 from https://en.wikipedia.org/wiki/Two-phase_commit_protocol
- [27] G. Zyskind, C. Kisagun, and C. Fromknecht. [n. d.]. Enigma Catalyst: a machine-based investing platform and infrastructure for crypto-assets. https://www.enigma.co/enigma_catalyst.pdf. As of 25 January 2018.