



HAL
open science

Contingent payments on a public ledger: models and reductions for automated verification

Sergiu Bursuc, Steve Kremer

► **To cite this version:**

Sergiu Bursuc, Steve Kremer. Contingent payments on a public ledger: models and reductions for automated verification. ESORICS 2019 - The 24th European Symposium on Research in Computer Security, Sep 2019, Luxembourg, Luxembourg. hal-02269063

HAL Id: hal-02269063

<https://hal.science/hal-02269063v1>

Submitted on 22 Aug 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Contingent payments on a public ledger: models and reductions for automated verification

Sergiu Bursuc and Steve Kremer

Inria Nancy-Grand'Est & LORIA, France

Abstract. We study protocols that rely on a public ledger infrastructure, concentrating on protocols for zero-knowledge contingent payment, whose security properties combine diverse notions of fairness and privacy. We argue that rigorous models are required for capturing the ledger semantics, the protocol-ledger interaction, the cryptographic primitives and, ultimately, the security properties one would like to achieve.

Our focus is on a particular level of abstraction, where network messages are represented by a term algebra, protocol execution by state transition systems (e.g. multiset rewrite rules) and where the properties of interest can be analyzed with automated verification tools. We propose models for: (1) the rules guiding the ledger execution, taking the coin functionality of public ledgers such as Bitcoin as an example; (2) the security properties expected from ledger-based zero-knowledge contingent payment protocols; (3) two different security protocols that aim at achieving these properties relying on different ledger infrastructures; (4) reductions that allow simpler term algebras for homomorphic cryptographic schemes.

Altogether, these models allow us to derive a first automated verification for ledger-based zero-knowledge contingent payment using the Tamarin prover. Furthermore, our models help in clarifying certain underlying assumptions, security and efficiency tradeoffs that should be taken into account when deploying protocols on the blockchain.

1 Introduction

The blockchain and its associated public ledger promise a practical solution to a basic need for security protocols: a system that operates as stated, providing reliable outcome to all agents. Both deployed [1–4] and abstract [5, 6] ledgers are ordered sequences of states - *state transition systems* respecting operational constraints. The goal of the underlying distributed protocols is to ensure that the ledger is indeed public, unique, alive and consistent. Protocols can then be based on transaction and smart contract semantics - i.e. rules that guide the state transition system - to implement functionality that would otherwise be inefficient or require trusted parties. Take *fair exchange*: two parties want to swap assets according to a contract that ensures fairness : any information or value transfer is reciprocated as planned [7]. The problem can be solved with optimistic assumptions, calling a trusted third party only when needed [8–10], or with digital (counter)cheques and transactions inside multi-party computations [11–13].

A public ledger provides an alternative solution to the problem, specified as a *zero-knowledge contingent payment* (ZKCP) for a seller and buyer. We suppose that the

information of interest can be expressed as data (a *witness*) satisfying functional constraints (a desired *result*), e.g. a sudoku solution respects additive constraints, a prime factor decomposition satisfies multiplicative constraints, etc. ZKCP goals are: for the **Seller** - a delivered witness will be paid for; for the **Buyer** - a paid for witness will be delivered. Classically, these properties require trust and coordination with third parties. On public ledgers, reliable semantics and dedicated cryptographic protocols can minimize trust and interaction [14–18].

Challenges. Protocol actions occur at distinct levels: from local cryptographic objects, to network transactions, to ledger confirmation. Their respective semantics is useful in protocol design, where parties can agree on desired ledger actions beforehand, yet the concurrent environment opens up new challenges:

- *Multiple sessions, concurrent ledger access.* Asynchronicity leads to ambiguity about what it means to be paid. For example, a seller should ensure it will not be *paid* the same coin for two witnesses. If multiple sessions run in parallel, some with colluding parties, protocol messages may be mixed up and exploited. Valid transaction requests do not necessarily result in confirmed ledger transactions : if the adversary obtains private keys by exploiting the protocol, a race ensues between honest and adversarial messages claiming a coin. Protocols should ensure this does not happen - this is not usually an explicit goal.
- *Transaction finality.* In fact, it is commonly advised to wait for transactions to be finalized on the ledger to ensure payment. Yet, we show that ZKCP protocols (have to) provide a stronger property: as early as a transaction request is being sent over the network, one should ensure that the corresponding coin cannot be spent in any other way, because specific fields from the transaction may help the adversary in revealing secrets - so we cannot afford the transaction to fail.
- *Cryptographic interaction.* Ledger-based protocols produce complex cryptographic objects that engage ledger transitions at the same time as private data transfer, e.g. [15] relies on homomorphic encryption to produce a (secret) ECDSA signature that will perform a ledger transaction; this signature is committed in a zero-knowledge proof ensuring the corresponding ledger transition will furthermore reveal the witness. Such interaction between cryptography and the ledger extends the scope of crypto primitives to new protocols - dedicated, fine-grained security models are needed to evaluate them.
- *Security foundations.* Compounding all of above: ledger-based protocols are network cryptographic protocols executed in an adversarial environment. There is history of attacks and foundations for such protocols - see e.g. [19–23] for recent examples - showing the importance of rigorous security specification and automated verification. Furthermore, we need generic models that allow a clear separation between security properties, ledger infrastructure and cryptographic protocols.

Our contributions address these challenges by formal models connecting the ledger, the ledger-based protocols, the cryptographic primitives and the desired security properties in a specification that can be used as input for automated verification tools. We use the Tamarin prover [24] for verification: it provides an expressive language to specify (cryptographic) state transition systems and to restrict their traces by logical formulas.

- *Public ledger.* We show that the model of the blockchain as a structured computational resource has a natural formal (or symbolic) counterpart combining multiset rewriting,

term algebras and first order logic [24–26]. We identify minimal restrictions on multiset rewriting rules that make them function as a blockchain transition system, i.e. a smart contract. We also show how protocol rules can operate in order to exploit the ledger semantics. We specify the electronic coin functionality provided in e.g. Bitcoin [1] as an example (section 3).

- *ZKCP on public ledgers.* We consider two ZKCP protocols [14, 15] and perform their formal verification in a unified, generic model that captures their different features (sections 4 and 5). The specification tackles a strong attacker that can run multiple sessions, corrupt parties, control the network (in particular drop, reorder, replace the messages to the ledger) and exploit the cryptographic properties of messages. The formal security properties clearly circumscribe the expected ZKCP guarantees, both in their positive and in their negative aspects: e.g. a buyer will learn the witness or otherwise it can obtain a refund; a seller will obtain payment, unless there is a delivery delay to the ledger; etc. The security properties are parametric, so that different protocols can accordingly instantiate the notions of payment, time delay, witness extraction, etc.

- *Advanced cryptography.* The protocol we consider in section 5 aims at a basic version of Bitcoin, with a minimal scripting language for signature verification; this calls for complex cryptography, intertwining homomorphic encryption, randomized signatures, diffie-hellman exponentiation and specialized zero-knowledge proofs. The corresponding formal specification as a message theory is out of the scope for any current automated verification tools. We provide a theoretical framework and a reduction result showing that it is sound to consider a simplified theory as input (section 6). We start from a general theory where some of the function symbols are homomorphic: from $f(u, \bar{w})$ and v , one can derive $f(u * v, \bar{w})$, where $*$ is the product in an abelian group. In the reduced theory: 1) the homomorphic properties are restricted as follows: the adversary can derive $f(u * v, \bar{w})$ from $f(u, \bar{w})$ only if u is a product of messages created by honest parties; 2) the abelian group is degenerated: the adversary can derive the factors u_1, \dots, u_k of any product $u_1 * \dots * u_k$, without being required to know any inverse.

2 Preliminaries: computation model

Term algebra [27]. \mathcal{F} denotes the set of function symbols and $\mathcal{F}^{(n)}$ those of arity n . The set of terms built from \mathcal{F} , a set of names and a set of variables is \mathcal{T} . Tuples of terms are denoted by an overline, e.g. $\bar{u} = (u_1, \dots, u_n)$. We let $st(t)$ be the subterms of a term t , and $top(t)$ be its top symbol. \mathcal{F} is endowed with a rewrite system: a set of rewrite rules \mathcal{R} , that we denote by $l \rightarrow r$, modulo a set of equations \mathcal{E} , that we denote by $l \approx r$. \mathcal{R} or \mathcal{E} can be empty. For a term t , $t \downarrow_{\mathcal{R}}$ is its normal form, obtained after applying all possible rewrite steps (modulo \mathcal{E}) from \mathcal{R} . Implicitly, terms are normalized and term equalities interpreted modulo $(\mathcal{R}, \mathcal{E})$.

Example 1. For the theory of randomized signatures, as instantiated e.g. by (EC)DSA [28], we let $\mathcal{F}_{\text{sig}} = \{\text{sign}, \text{ver}, \text{ok}, g\}$ and \mathcal{R}_{sig} be the signature verification rule: $\text{ver}(\text{sign}(x, y, z), x, g(y)) \rightarrow \text{ok}$. Here $g(y)$ represents the public key corresponding to a secret key y , i.e. the group element that corresponds to raising a group generator g to a scalar power y . The third argument of sign takes the role of the randomness: $\text{sign}(m, k, r_1)$ and $\text{sign}(m, k, r_2)$ are two distinct signatures of m with key k .

The theory of an abelian group (AG), e.g. \mathbb{Z}_q , is modeled by the signature $\mathcal{F}_* = \{*, i\}$ and the set of equations $\text{AG} = \{x * i(x) \approx 1, x * 1 \approx x\} \cup \text{AC}$ where $\text{AC} = \{x * y \approx y * x, (x * y) * z \approx x * (y * z)\}$ models associativity and commutativity.

Multiset rewriting and state transitions [24, 26]. The signature is extended with fact symbols to represent adversarial knowledge, protocol state, freshness information, etc. A fact is represented by $F(t_1, \dots, t_k)$, where F is a fact symbol and t_1, \dots, t_k are terms. There are the following special fact symbols: K - for attacker knowledge; Fr - for fresh data; In and Out - for protocol inputs and outputs. Other symbols may be added as required by the protocol, e.g. for representing the state. These symbols can be persistent (the corresponding facts cannot disappear), or linear (the corresponding facts are consumed by rules and protocol rules can update them). Persistent fact symbols are prefixed by !, e.g. !F. A multiset can contain multiple copies of the same linear fact.

A multiset rewriting (msr) rule is defined by $[L] \dashv [M] \mapsto [N]$, where L, M, N are multisets of facts called respectively premisses, actions and conclusions. We denote such a rule by $[L] \Rightarrow [N]$ when M is empty. To ease protocol specification, we extend the syntax of multiset rules with variable assignments and equality constraints, i.e. we can write rules of the form $[L] \dashv [\Phi, M] \mapsto [N]$ where L may contain expressions $x = t$ to define local variables and Φ is a set of equations of the form $u \approx v$. Equations are not directly supported in Tamarin, but can be easily encoded with restrictions as we show in Example 3. For two multisets of facts M_0, M_1 and rule $P = [L] \dashv [\Phi, M] \mapsto [N]$ we say that M_1 can be obtained from M_0 by applying the rule P , instantiated with θ if: (1) every equality in $\Phi\theta$ is true; (2) every fact in $L\theta$ is included in M_0 (counting multiplicities for linear facts); (3) M_1 is obtained from M_0 by removing linear facts included in $L\theta$ and adding all facts from $N\theta$.

A special set of *message deduction rules* defines how the attacker can derive new knowledge and make use of existing knowledge to interact with the protocol. Within this set, we distinguish *network deduction rules* and *intruder deduction rules*. Network deduction rules are fixed: they define outputs, inputs, public and fresh data.

$$[\text{Out}(x)] \Rightarrow [\text{K}(x)]; [\text{K}(x)] \Rightarrow [\text{In}(x)]; \Rightarrow [\text{K}(y)]; \Rightarrow [\text{Fr}(z)]; [\text{Fr}(x)] \Rightarrow [\text{K}(x)]$$

The semantics ensures that y and z above are instantiated to public, resp. fresh names.

Intruder deduction rules are of the form $[\text{K}(u_1), \dots, \text{K}(u_k)] \Rightarrow [\text{K}(v)]$ - defining operations on messages. These are typically $[\text{K}(x_1), \dots, \text{K}(x_k)] \Rightarrow [\text{K}(f(x_1, \dots, x_k))]$ for all $f \in \mathcal{F}^{(k)}$. We also allow more general deduction rules, as in Example 2 and Figure 4. Such rules can wlog replace rewrite rules $f(l_1, \dots, l_k) \rightarrow r$ for symbols f with no other occurrence in the rewrite system and whose occurrence in protocol rules is not under a term context. An *intruder theory*, that we denote by \mathcal{I} , is thus given by a set of intruder deduction rules plus $(\mathcal{R}, \mathcal{E})$. For a set of terms $\{t_1, \dots, t_n, t\}$ we let $\{t_1, \dots, t_n\} \vdash t$ if $\text{K}(t)$ can be obtained from $\text{K}(t_1), \dots, \text{K}(t_n)$ using intruder deduction rules. *Protocol rules* model the execution of the protocol by honest parties. There are basic restrictions ensuring that protocol rules are a sound model of protocol executions [26]; we will follow them implicitly in our models and examples.

Example 2. Exponentiation in a Diffie-Hellman group can be represented by the rewrite rule $\text{exp}(g(x), y) \rightarrow g(x * y)$ together with the deduction rule $[\text{K}(x_1), \text{K}(x_2)] \Rightarrow$

$[K(exp(x_1, x_2))]$. Alternatively, the deduction rule $[K(g(x)), K(y)] \Rightarrow [K(g(x * y))]$ allows to model the corresponding operation performed by the attacker (without requiring explicit application of exp). Similarly, a protocol rule can directly perform exponentiation without explicit use of the symbol exp , e.g. $[In(g(x)), Fr(y)] \Rightarrow [Out(g(x * y))]$.

For a rule P , we let $facts(P)$, $in(P)$, $out(P)$, $lhs(P)$, $rhs(P)$, $act(P)$ be respectively the set of all facts, of input facts (e.g. $In(u)$), of output facts (e.g. $Out(u)$), of left-hand side facts (i.e. premisses), of right-hand side facts (i.e. conclusions) and of action facts. For a set of facts \mathbf{F} , we let $msg(\mathbf{F})$ be the set of messages that are arguments of facts in \mathbf{F} . We let $io(P) = msg(in(P) \cup out(P))$.

Traces and properties. A *trace* τ is a sequence of applications of $n \geq 1$ msr rules, interleaving applications of protocol, intruder and network deduction rules. For every $i \in \{1, \dots, n\}$, we let P_i be the rule applied at step i and θ_i be the corresponding substitution. We define:

- $facts(\tau, i) = act(P_i)\theta_i\downarrow$ if P_i is a protocol or network deduction rule;
- $facts(\tau, i) = \{K(v\theta_i\downarrow)\}$ if P_i is an intruder deduction rule with $rhs(P_i) = \{K(v)\}$

For a set of rules Q , we denote by $traces(Q)$ the set of all valid traces that can be derived from elements in Q . Consider a set of timepoint variables, denoted by i, j, l, \dots , which will be interpreted over rational numbers. A *trace atom* is either \perp , or a term equality $t_1 \approx t_2$, or a timepoint ordering $i < j$, or a timepoint equality $i = j$, or an action fact $\mathbf{F}@i$ for a fact \mathbf{F} and timepoint i . A *trace formula* is a first-order logic formula obtained from trace atoms by applying the usual quantification and logical connectives. Given a trace τ and trace formula ϕ , whose variables are all bound, the satisfaction relation $\tau \models \phi$, is defined recursively as expected, in particular $\tau \models \mathbf{F}@i$ iff $\mathbf{F} \in facts(\tau, i)$.

For a set of rules Q and trace formulas Ψ, Φ , we let $Q \models \Phi$ iff $\forall \tau \in traces(Q). \tau \models \Phi$ and $Q; \Psi \models \Phi$ iff $\forall \tau \in traces(Q). \tau \models \Psi \Rightarrow \Phi$. For verification, $(Q; \Psi)$ will be a system specification and Φ a property to verify; Q defines local transition rules, while Ψ defines additional, global restrictions on the set of traces for the specified system.

Example 3. Consider the binary fact symbol Eq and the formula

$$\Psi_{eq} : \forall x, y, i. Eq(x, y) @ i \Rightarrow x \approx y.$$

An $Eq(u, v)$ action in a rule allows then to test that $u \approx_\varepsilon v$ before proceeding. Take $P = [In(u), In(v), Fr(s)] \dashv [Eq(u, v)] \dashv [Out(s)]$. Then $K(a), K(a), Eq(a, a), K(s)$ is a trace of P satisfying Ψ_{eq} , while $K(a), K(f(a)), Eq(a, f(a)), K(s)$ does not.

Consider the unary symbol $Fresh$ and the restriction

$$\Psi_{fresh} : \forall x, i, j. Fresh(x) @ i \wedge Fresh(x) @ j \Rightarrow i = j.$$

It ensures that every occurrence of $Fresh(t)$ is with a different t . Assume we add $Fresh(\langle u, v \rangle)$ as an action in P . Then, among $traces(P)$, $\dots Eq(a, a), \dots, Eq(a, a)$ does not satisfy Ψ_{fresh} , while $\dots Eq(a, a), \dots, Eq(b, b)$ does.

Example 4. Consider the set of rules Q_{keys} :

- $[Fr(k)] \text{---} [!Key(k)] \text{---} [Pk(g(k)), !Key(k), Out(g(k))]$
- $[!Key(x)] \text{---} [Corrupt(g(x))] \text{---} [Out(x)]$

It models a basic key infrastructure. The formula $\Phi : !Key(x) @ i \Rightarrow \neg \exists j. K(x) @ j$ says that keys are secret. Then $Q_{keys} \not\models \forall x, i. \Phi$, since the second rule in Q_{keys} allows the attacker to corrupt keys. Now consider the protocol rule

$$Q_{sign} : [Fr(a), !Key(x)] \text{---} [Honest(g(x)), Sign(x)] \text{---} [Out(sign(a, k, \rho_r))]$$

the formula $\Phi' : Sign(x) @ j \Rightarrow \neg \exists j. K(x) @ j$ - saying that keys used in Q_{sign} are secret - and the restriction: $\Psi_{hon} : \forall x, i. Honest(x) @ i \Rightarrow \neg \exists j. Corrupt(x) @ j$. Then we have $Q_{keys}, Q_{sign}; \Psi_{hon} \models \forall x, i. \Phi'$ because we have added the restrictions that keys in Q_{sign} are honest and that honest keys cannot be corrupted.

Public data. Tamarin allows the use of variables that can be instantiated only with messages of a public sort. They are denoted by $\$x$, and can occur anywhere in a protocol msr rule. As in Example 4, we will use annotations of ρ for such data, e.g. ρ_r for a public nonce, ρ_{sn} for a serial number, etc.

Protocol state. Specifications rely on sequences of protocols rules (P_0, \dots, P_k) , where each rule P_i should be executed before P_{i+1} and can pass on, via facts, state data to P_{i+1} . To avoid clutter, we use a symbol $state_i$ to represent this transmission, and we allow P_{i+1} to reference any variables from P_i that should be formally passed via state facts. We denote by $state_i[x = u]$ the pattern matching of state variable x by a term u .

3 Public ledgers: facts, rules, coins

Coin ledger. The protocols we consider are based on coin contracts of e.g. Bitcoin [1]: a *coin* is represented by an object $(sn, g(k))$ on the ledger, where sn is a serial number, and $g(k)$ is the public key of the coin owner. Serial numbers are computed as the hash of the transaction that created the coin; for simplicity, we assume they are fresh public numbers. To spend a coin, i.e. transfer it to a new owner, the ledger expects a transaction request, attested by a signature from the current owner, containing the sn of the coin to be spent, the public key $g(k')$ of the new owner and (implicitly) the serial number sn' of the new coin. If the signature is valid, the coin $(sn, g(k))$ is marked as spent, and a new coin $(sn', g(k'))$ is created for the new owner. We call *basecoins* these coins.

We will also make use of *hashcoins: hashed timelock contracts* [29] used to establish trust relationships outside the ledger [30, 31]. They perform a transaction by which one of the two parties, say A, obtains the preimage of a hash - which can e.g. be a key encrypting some data of interest - while the other party, say B, provides the hash preimage and obtains a basecoin in return. A performs a ledger transaction pledging one of A's coins into a hashcoin, providing the desired hash image and the public key of B. B can then claim the coin using a (signed) inverse of the image. A timeout mechanism ensures the coin can be returned to A if there was no action from B in due time. A hashcoin can be represented by a tuple $(sn, g(k), h(x), g(k'))$ here $g(k)$ represents the coin creator, who can obtain it after timeout, $h(x)$ is the desired hash image, and $g(k')$ is the party that can claim sn by supplying x .

Formal model. We consider two special sets of disjoint fact symbols: one for *ledger facts*, denoted by $\mathbf{F}_{\mathcal{L}}$, and one for *check facts*, denoted by $\mathbf{F}_{\mathcal{C}}$. Ledger facts will be used to represent the state of the ledger. For example, they can record who is the owner of an asset, what are the elements of a given transaction, etc. Ledger facts are assumed persistent because the ledger history cannot change. Check facts, on the other hand, will be used by protocols to restrict their executions with respect to the (current or past) states of the ledger. For example, they can be used to ensure that a coin, whose existence is recorded by a ledger fact, has not yet been spent.

Example 5. Let $\mathcal{F}_{\mathcal{L}}^{\text{coin}} = \{\text{!Coin}, \text{!HCoin}, \text{!Spend}, \text{!Time}\}$ and $\mathcal{F}_{\mathcal{C}}^{\text{coin}} = \{\text{Unspent}\}$. The corresponding facts represent: $\text{!Coin}(\text{sn}, g(k)) @ i$ - a coin sn created at timepoint i belonging to the public key $g(k)$; $\text{!HCoin}(\text{sn}, \langle g(k_1), g(k_2), h(t) \rangle) @ i$ - a hashcoin sn that can be claimed for $g(k_2)$ by supplying t and a signature, or for $g(k_1)$ after timeout by supplying a signature; $\text{!Spend}(\text{sn}, u, w, v) @ i$ - the transfer of a coin (sn, u) to a new owner v at timepoint i , relying on supporting data w : w is a signature when sn is a basecoin, plus possibly a hash preimage when sn is a hashcoin; $\text{!Time}(\text{sn}) @ i$ marks the fact that the hashcoin sn was reclaimed after a timeout at timepoint i ; $\text{Unspent}(\text{sn}) @ i$ checks the ledger to ensure the coin sn is unspent at i .

The semantics of the ledger is defined by msr rules that can only be triggered by ledger facts and public inputs, and can only produce ledger facts and public outputs. Ledger restrictions ensure additional constraints for the states produced by the ledger. These rules and constraints define the ledger state transition system and make it available for external protocols, which may be executed by honest or adversarial parties.

Definition 1. A msr rule P is a ledger rule if: (1) $\text{facts}(P) \subseteq \text{in}(P) \cup \text{out}(P) \cup \mathbf{F}_{\mathcal{L}}$; (2) $\text{rhs}(P) \subseteq \text{act}(P)$. P is ledger-respecting if $(\text{act}(P) \cup \text{rhs}(P)) \cap \mathbf{F}_{\mathcal{L}} = \emptyset$. A ledger restriction is a trace formula with facts in $\mathbf{F}_{\mathcal{L}} \cup \mathbf{F}_{\mathcal{C}}$.

Properties of ledger rules in Definition 1 ensure that: (1) the ledger transition system depends only on ledger facts and public inputs; (2) all produced ledger facts are recorded as actions in the trace. In this paper we consider public ledgers, e.g. [1–4], so the ledger rules will also satisfy (3) $\text{msg}(\text{rhs}(P)) \subseteq \text{msg}(\text{out}(P))$. This is not an inherent restriction of the model, and partially public ledgers, e.g. [32], may be considered in the scope of Definition 1. Bearing in mind the properties (2) and (3) of our considered ledger rules, in order to simplify the presentation of our examples in the paper, we will avoid duplication, writing $[F_0] \text{---} [\Phi] \text{---} [F_1]$ instead of $[F_0] \text{---} [\Phi, F_1] \text{---} [F_1, \text{Out}(\text{msg}(F_1))]$ as expected. All protocol rules will be ledger-respecting as in Definition 1, so the only way to produce ledger facts is by passing through ledger rules; on the other hand, protocol rules can freely access ledger facts to check the state of the ledger, so we can have $\text{lhs}(P) \cap \mathbf{F}_{\mathcal{L}} \neq \emptyset$.

In Fig. 1, the rule R_{new} abstracts the coin mining process; the other rules model formally the coin transactions as described above: spending coins to coins, to hashcoins, and back to coins. The rule R_{h2cr} produces a ledger fact $\text{!Time}(x_{\text{sn}})$ to record that the corresponding coin was reclaimed after a timeout. The rules $S_{\text{c2h}}, S_{\text{h2c}}$ assume Hash and Inv to be defined by their context as a hash image of interest and a hash preimage.

Fig. 1. Ledger coin rules: $\mathcal{L}_{\text{base}} = \{\text{R}_{\text{new}}, \text{R}_{\text{c2c}}\}$; $\mathcal{L}_{\text{hash}} = \mathcal{L}_{\text{base}} \uplus \{\text{R}_{\text{c2h}}, \text{R}_{\text{h2c}}, \text{R}_{\text{h2cr}}\}$

$\text{R}_{\text{new}} : [\text{!Pk}(x_{\text{pk}}), \text{In}(\langle s, x_{\text{sn}} \rangle)] \text{---} [\text{ver}(s, x_{\text{sn}}, x_{\text{pk}}) \approx \text{ok}] \text{---} [\text{!Coin}(x_{\text{sn}}, x_{\text{pk}})]$
$\text{R}_{\text{c2c}} : [\text{!Coin}(x_{\text{sn}}, x_{\text{pk}}), \text{In}(u)] \text{---} [\Phi_{\text{c2c}}(x_{\text{sn}}, x_{\text{pk}}, u)] \text{---} [\text{!Spend}(x_{\text{sn}}, x_{\text{pk}}, v), \text{!Coin}(y_{\text{sn}}, y_{\text{pk}})]$
$\text{R}_{\text{c2h}} : [\text{!Coin}(x_{\text{sn}}, x_{\text{pk}}), \text{In}(u)] \text{---} [\Phi_{\text{c2h}}(x_{\text{sn}}, x_{\text{pk}}, u)] \text{---} [\text{!Spend}(x_{\text{sn}}, x_{\text{pk}}, s, y), \text{!HCoin}(y_{\text{sn}}, y)]$
$\text{R}_{\text{h2c}} : [\text{!HCoin}(x_{\text{sn}}, y), \text{In}(u)] \text{---} [\Phi_{\text{h2c}}(x_{\text{sn}}, y, u)] \text{---} [\text{!Spend}(x_{\text{sn}}, y, s, y_{\text{pk}}), \text{!Coin}(z_{\text{sn}}, y_{\text{pk}})]$
$\text{R}_{\text{h2cr}} : [\text{!HCoin}(x_{\text{sn}}, y), \text{In}(u)] \text{---} [\Phi_{\text{h2cr}}(x_{\text{sn}}, y, u)] \text{---} [\dots, \text{!Coin}(z_{\text{sn}}, x_{\text{pk}}), \text{!Time}(x_{\text{sn}})]$

where

$\text{R}_{\text{c2c}} : u = \langle s, y_{\text{sn}}, y_{\text{pk}} \rangle$; $\Phi_{\text{c2c}} = \text{ver}(s, \langle \text{c2c}, x_{\text{sn}}, y_{\text{sn}}, y_{\text{pk}} \rangle, x_{\text{pk}}) \approx \text{ok}$; $v = \langle s, y_{\text{pk}} \rangle$
$\text{R}_{\text{c2h}} : u = \langle s, y_{\text{sn}}, y_{\text{pk}}, y_h \rangle$; $\Phi_{\text{c2h}} = \text{ver}(s, \langle \text{c2h}, x_{\text{sn}}, y_{\text{sn}}, y_{\text{pk}}, y_h \rangle, x_{\text{pk}}) \approx \text{ok}$; $y = \langle x_{\text{pk}}, y_{\text{pk}}, y_h \rangle$
$\text{R}_{\text{h2c}} : y = \langle x_{\text{pk}}, y_{\text{pk}}, y_h \rangle$; $u = \langle s, y_{\text{sn}}, y_w \rangle$; $\Phi_{\text{h2c}} = \text{ver}(s, \langle \text{h2c}, x_{\text{sn}}, y_w \rangle, y_{\text{pk}}) \approx \text{ok} \wedge y_h \approx h(y_w)$ (similarly for R_{h2cr})

Ledger-based protocol rules (typical examples)

$\text{S}_{\text{c2c}} : [\text{!Key}(x_{\text{sk}}), \text{!Pk}(y_{\text{pk}}), \text{!Coin}(x_{\text{sn}}, g(x_{\text{sk}})), x_s = \text{sign}(\langle \text{c2c}, x_{\text{sn}}, \rho_{\text{sn}}, y_{\text{pk}} \rangle, x_{\text{sk}}, \rho_r)]$ $\text{---} [\text{Unspent}(x_{\text{sn}})] \text{---} [\text{Out}(\langle x_s, \rho_{\text{sn}}, y_{\text{pk}} \rangle)]$
$\text{S}_{\text{c2h}} : [\text{!Key}(x_{\text{sk}}), \text{!Pk}(y_{\text{pk}}), \text{!Coin}(x_{\text{sn}}, g(x_{\text{sk}})), \text{Hash}(y_h)]$ $\text{---} [\text{Unspent}(x_{\text{sn}})] \text{---} [\text{Out}(u_{\text{c2h}})]$
$\text{S}_{\text{h2c}} : [\text{!Key}(y_{\text{sk}}), \text{!HCoin}(x_{\text{sn}}, \langle x_{\text{pk}}, g(y_{\text{sk}}), h(x_w) \rangle), \text{Inv}(y_w)]$ $\text{---} [\text{Unspent}(x_{\text{sn}}), \text{Claim}(x_{\text{sn}}, g(y_{\text{sk}}))] \text{---} [\text{Out}(u_{\text{h2c}})]$

where $t_{\text{c2h}} = \langle \text{c2h}, x_{\text{sn}}, y_{\text{pk}}, y_h \rangle$; $u_{\text{c2h}} = \langle \text{sign}(t_{\text{c2h}}, x_{\text{sk}}, \rho_r), \rho_{\text{sn}}, y_{\text{pk}}, y_h \rangle$
 $t_{\text{h2c}} = \langle \text{h2c}, x_{\text{sn}}, x_w, \rho_{\text{sn}} \rangle$; $u_{\text{h2c}} = \langle \text{sign}(t_{\text{h2c}}, y_{\text{sk}}, \rho_r), \rho_{\text{sn}}, y_w \rangle$

Ledger restrictions define additional constraints that should be satisfied by the public ledger. If $\text{facts}(\Phi) \subseteq \mathbf{F}_{\mathcal{L}}$ then the restriction Φ is *inherent to the semantics of the ledger*, i.e. it is a check performed by the (distributed) trusted party that builds the ledger. On the other hand, if $\exists \mathbf{F} \in \text{facts}(\Phi) \cap \mathbf{F}_{\mathcal{L}}$, then Φ *restricts the execution of the protocols* with respect to the public ledger: a protocol rule P with a substitution θ such that $\mathbf{F}\theta \in \text{act}(P\theta)$ can perform a transition at timepoint i , only if $\mathbf{F}\theta @ i$ is consistent with $\Phi\theta$ and the previous ledger facts.

Example 6. The following formulas define ledger restrictions for coins on $\mathcal{L}_{\text{base}}, \mathcal{L}_{\text{hash}}$

$$\begin{aligned} \Psi_0 &: \forall x, \bar{y}, \bar{z}, i, j. \text{!Spend}(x, \bar{y}) @ i \wedge \text{!Spend}(x, \bar{z}) @ j \Rightarrow i = j \wedge \bar{y} = \bar{z} \\ \Psi_1 &: \forall x, y, z, i, j. \text{!F}_1(x, y) @ i \wedge \text{!F}_2(x, z) @ j \Rightarrow i = j \wedge y = z \\ &\quad (\forall F_1, F_2 \in \{\text{Coin}, \text{HCoin}\}) \\ \Psi_2 &: \forall x, \bar{y}, i, j. \text{Unspent}(x) @ i \wedge \text{!Spend}(x, \bar{y}) @ j \Rightarrow i < j \end{aligned}$$

They ensure that - no coin can be spent twice (Ψ_0); - every fresh coin has a fresh serial number (Ψ_1); - Unspent can hold at timepoint i only if the corresponding coin has not already been spent on the ledger (Ψ_2). Note that Ψ_0, Ψ_1 are inherent ledger restrictions, while Ψ_2 is a protocol ledger restriction. We let $\Psi_{\text{coin}} = \Psi_0 \wedge \Psi_1 \wedge \Psi_2$.

4 Zero knowledge contingent payments

We specify in a general framework the security guarantees that parties can expect from ZKCP protocols. We allow several parameters in definitions, that can be instantiated

differently by specific protocols and ledgers - we illustrate it on $\mathcal{L}_{\text{base}}$ and $\mathcal{L}_{\text{hash}}$. We are interested in generic ZKCP protocols, where any functionality can be obtained by instantiating the protocol with a specific function f . Security is independent of the actual function f , so we consider a generic f in the following.

For intuition, consider first a protocol on $\mathcal{L}_{\text{hash}}$ [14, 16]. It assumes a zero-knowledge proof system showing that a ciphertext provided by a party contains a witness for a desired result, where the symmetric encryption key is the preimage of a given hash value. We represent such a proof by $zk(w, v, u)$ where w is the witness, v is the hash preimage used as symmetric key, and u is the secret key of the party constructing the proof (for brevity, we omit public data that may be part of the proof). The following rewrite rules represent symmetric encryption and zk proof verification:

$\text{sdec}(\text{senc}(x, y), y) \rightarrow x \quad \text{ver}_{\text{zk}}(\text{zk}(x, y, z), \text{senc}(x, y), f(x), h(y), g(z)) \rightarrow \text{ok}$.
 These define $\mathcal{I}_{\text{hash}}$, where also $\forall f \in \mathcal{F}^{(k)}. [\mathbf{K}(x_1), \dots, \mathbf{K}(x_k)] \Rightarrow [\mathbf{K}(f(x_1, \dots, x_k))]$.
 Assume a seller with private key ks wants to sell w to a buyer with public key $g(kb)$.

Seller 1: generate a fresh key k ; output $\text{senc}(w, k), h(k), g(ks), zk(w, k, ks)$;

Buyer 1: receive above data from seller and, if the zk proof verifies, invoke R_{c2h} on $\mathcal{L}_{\text{hash}}$ to create a hashcoin for the given $h(k)$ and $g(ks)$: $!\text{HCoin}(\text{sn}, \langle g(kb), g(ks), h(k) \rangle)$;

Seller 2: inspect $\mathcal{L}_{\text{hash}}$ to see if the above coin was created; invoke R_{h2c} with k and ks to claim the coin; this reveals k and thus reveals the witness;

Buyer 2: inspect $\mathcal{L}_{\text{hash}}$ to see if R_{h2c} was invoked for the created hashcoin; if yes, the ledger will also contain the key k that allows the decryption of the ciphertext received at step 1; if not, the rule R_{h2cr} can be invoked after a time delay so that the coin is returned to the original owner.

Timeout. The fairness properties for the ZKCP protocols will be relative to the timely execution of certain operations. More precisely, if a certain action is not performed by a party in due time, then there is another action - grounded on the semantics of the ledger as in Example 7 or on cryptographic primitives as in Example 8 - that can be performed in order to compensate for the missing action.

Example 7 (Ledger timeout). Consider the rule R_{h2cr} from Figure 1 modeling the refund of a hashcoin after a timeout. The execution of this rule at timepoint i is accompanied on the ledger by the fact $!\text{Time}(x_{\text{sn}}) @ i$ to record that this coin was spent due to a timeout. This allows to specify the possible effects of invoking R_{h2c} on $\mathcal{L}_{\text{hash}}$: either the transaction completes as expected, or there was a timeout, i.e. R_{h2cr} was invoked. Consider the rule S_{h2c} from Figure 1; note the Claim action. Then $\mathcal{L}_{\text{hash}}$ ensures the following property:

$$\forall x, y, z, z_1, z_2, i, j. \quad \text{Claim}(x, y) @ i \wedge \quad \text{!Spend}(x, z_1, z_2, z) @ j \Rightarrow \quad z = y \vee \quad \text{!Time}(x) @ j$$

where $z = y$ happens in a normal execution, and $!\text{Time}(x) @ j$ if the timeout occurs.

Example 8 (Cryptographic timeout [33, 34]). Time commitment schemes allow to produce a commitment to a message that keeps it secret for a period of time. We represent a time commitment to u by $\text{tcom}(u)$ and consider the following rule $\text{Q}_{\text{tcom}} : [\text{In}(\text{tcom}(x))] \text{---} [\text{!Time}(x)] \text{---} [\text{Out}(x)]$. We express that fresh committed data is either secret, or it was released after a timeout. Let $P : [\text{Fr}(s)] \text{---} [\text{Tcom}(s)] \text{---} [\text{Out}(\text{tcom}(s))]$. Then $\text{Q}_{\text{tcom}}, P \models \forall x, i, j. \text{Tcom}(x) @ i \wedge \mathbf{K}(x) @ j \Rightarrow \exists k. k < j \wedge \text{!Time}(x) @ k$

Fig. 2. Formal ZKCP on $\mathcal{L}_{\text{hash}}$; Seller = (S_0, S_1, S_2) ; Buyer = $(B_0, B_1, B_2^{\text{go}}, B_2^{\text{ab}})$

S_0 : [!Key(x_{ks}), !Witn(x_{wtn})] — [Sell($g(x_{\text{ks}}), x_{\text{wtn}}$)] → [state ₀]
S_1 : [state ₀ , Fr(k), $x_{\text{ew}} = \text{senc}(x_{\text{wtn}}, k)$, $x_{\pi} = \text{zk}(x_{\text{wtn}}, k, x_{\text{ks}})$] ⇒ [Out($\langle x_{\pi}, x_{\text{ew}}, h(k) \rangle$), state ₁]
S_2 : [state ₁ , !HCoin($x_{\text{sn}}, \langle x_{\text{pkb}}, g(x_{\text{ks}}), h(k) \rangle$)] — [Unspent(x_{sn}), Claim($g(x_{\text{ks}}), x_{\text{wtn}}, x_{\text{sn}}, x_{\text{sn}}$)] → [Out($\langle \text{sign}(\langle \text{h2c}, x_{\text{sn}}, \rho_{\text{sn}}, k \rangle), x_{\text{ks}}, k, \rho_{\text{sn}} \rangle$)]
B_0 : [!Res(x_{res}), !Key(x_{kb}), !Pk(x_{pks}), !Coin($x_{\text{sn}}, g(x_{\text{kb}})$)] ⇒ [state ₀]
B_1 : [state ₀ , In($\langle x_{\pi}, x_{\text{ew}}, x_h \rangle$)] — [ver _{zk} ($x_{\pi}, x_{\text{ew}}, x_{\text{res}}, x_h, x_{\text{pks}}$) ≈ ok, Pay($g(x_{\text{kb}}), x_{\text{res}}, \rho_{\text{sn}}, \langle x_{\pi}, x_{\text{ew}}, x_h \rangle$)] → [Out($\langle \text{sign}(\langle \text{c2h}, x_{\text{sn}}, \rho_{\text{sn}}, x_{\text{pks}}, x_h \rangle), x_{\text{kb}}, \rho_{\text{sn}}, x_{\text{pks}}, x_h \rangle$), state ₁]
B_2^{go} : [state ₁ , !Spend($\rho_{\text{sn}}, z, \langle x_s, x_k \rangle, x_{\text{pks}}$), $x_{\text{wtn}} = \text{sdec}(x_{\text{ew}}, x_k)$] — [$h(x_k) \approx x_h, f(x_{\text{wtn}}) \approx x_{\text{res}}, \text{Witness}(x_{\text{res}})$] → []
B_2^{ab} : [state ₁ , !HCoin($x_{\text{sn}}, \langle g(x_{\text{kb}}), x_{\text{pks}}, x_h \rangle$)] — [Unspent(x_{sn})] → [Out($\langle \text{sign}(\langle \text{h2cr}, x_{\text{sn}}, \rho_{\text{sn}} \rangle), x_{\text{kb}}, \rho_{\text{sn}} \rangle$)]

Definition 2. Let Q be a set of (protocol and ledger) rules and Ψ be a set of restrictions. We say that (Q, Ψ) is a

- coin infrastructure if Q produces !Spend($u_{\text{coin}}, \bar{u}, u_{\text{pk}}$) ledger facts and $\Psi_{\text{coin}} \subseteq \Psi$ (see Figure 1 and Example 6);
- time infrastructure if Q produces !Time(u) actions (see Example 7 and Example 8);
- key infrastructure if $Q_{\text{keys}} \subseteq Q$ (see Example 4)
- function model if Q contains the rules Q_{func} :

$$[\text{Fr}(x_w)] \Rightarrow [\text{!Witn}(x_w), \text{Out}(f(x_w))] ; \quad [\text{Fr}(x_w)] \Rightarrow [\text{!Res}(f(x_w)), \text{Out}(x_w)]$$

If all of these are satisfied we say that (Q, Ψ) is a ZKCP-context.

The fact !Witn(x_w) from a function model is used by an honest seller to determine a witness, and the adversary (playing the role of the buyer) obtains a desired result $f(x_w)$. The fact !Res($f(x_w)$) is used by an honest buyer to determine a desired result, and the adversary (playing the role of the seller) obtains the corresponding witness x_w .

Definition 3. A ZKCP Seller specification is given by a set of protocol rules that contains two special rules:

$$\text{sell: } [\dots] \text{— [Sell}(t_{\text{pk}}, t_{\text{wtn}}) \text{] → } [\dots]$$

$$\text{claim: } [\dots] \text{— [Claim}(t_{\text{pk}}, t_{\text{wtn}}, t_{\text{time}}, t_{\text{sn}}) \text{] → } [\dots]$$

The *sell* rule models the start of a seller session, recording in Sell($t_{\text{pk}}, t_{\text{wtn}}$) the seller public key and the witness. The *claim* rule models the seller claiming a coin as payment, producing an action fact Claim($t_{\text{pk}}, t_{\text{wtn}}, t_{\text{time}}, t_{\text{sn}}$) where $t_{\text{pk}}, t_{\text{wtn}}$ are as above, t_{time} is timeout constrained data, and t_{sn} the claimed coin. In our case studies, t_{time} is either a sn as in Ex. 7 or a secret key share, cryptographically committed as in Ex. 8. See in Fig. 2 the formal Seller specification for the protocol above.

Definition 4. Let (Q, Ψ) be a ZKCP-context and S be a ZKCP Seller specification. We say that these ensure seller security if $Q, S; \Psi \models \Phi_S$, where Φ_S is defined in Figure 3.

Fig. 3. Security properties for ZKCP on a ledger

Seller security: witness reveal vs payment: $\Phi_S := \Phi_0 \wedge \Phi_1 \wedge \Phi_2$
$\Phi_0 : \forall x_{pk}, x_{wtn}, i, j. \text{Sell}(x_{pk}, x_{wtn}) @ i \wedge K(x_{wtn}) @ j \Rightarrow \exists k, y_{pk}, x_t, x_{coin}. \text{Claim}(y_{pk}, x_{wtn}, x_t, x_{coin}) @ k$
$\Phi_1 : \forall \bar{y}, \bar{z}, x. \text{Claim}(\bar{y}, x) @ i \wedge \text{Claim}(\bar{z}, x) @ j \Rightarrow i = j$
$\Phi_2 : \forall x_{pk}, x_{wtn}, x_t, x_{coin}, i, j. \text{Claim}(x_{pk}, x_{wtn}, x_t, x_{coin}) @ i \wedge \text{!Spend}(x_{coin}, z, y, z_{pk}) @ j$ $\Rightarrow z_{pk} = x_{pk} \vee \exists k. k \leq j \wedge \text{!Time}(x_t) @ k$
Buyer security: pay gives witness or refund: $\Phi_B := [\forall i, j, x_{pk}, x_{res}, x_{coin}, \bar{x}_{state}. (\Phi_0 \wedge \Phi_1)] \wedge \Phi_2$
$\Phi_0(\Psi_0) : \text{Pay}(x_{pk}, x_{res}, x_{coin}, \bar{x}_{state}) @ i \wedge \text{!Spend}(x_{coin}, z, y, z_{pk}) @ j \Rightarrow z_{pk} = x_{pk} \vee \Psi_0(y, \bar{x}_{state})$
$\Phi_1(\Psi_1) : \text{Pay}(x_{pk}, x_{res}, x_{coin}, \bar{x}_{state}) @ i \Rightarrow \Psi_1(x_{res}, \bar{x}_{state})$
$\Phi_2(\Psi_0, \Psi_1) : \forall x_{res}, y, \bar{x}_{state}. \Psi_0(y, \bar{x}_{state}) \wedge \Psi_1(x_{res}, \bar{x}_{state}) \Rightarrow \exists x_w. x_{res} = f(x_w) \wedge y, \bar{x}_{state} \vdash x_w$

Intuitively, the formula $\Phi_S = \Phi_0 \wedge \Phi_1 \wedge \Phi_2$ from Definition 4 ensures that:

- Φ_0 : if the other party learns the witness, then (one of) the seller(s) for the corresponding witness is able to claim the payment of a coin into seller's account;
- Φ_1 : the other party cannot lead the seller into accepting the same payment twice, e.g. for two different witnesses;
- Φ_2 : the payment claimed by the seller will succeed as such on the ledger, unless the corresponding timeout event happened.

Note that, in Φ_0 , the key y_{pk} into which payment is claimed is not necessarily equal to the key x_{pk} that engaged in selling the witness: the two keys can differ when there are two sellers for the same witness; then the adversary can learn the witness in one session without paying in the second one. Φ_1 requires care to ensure session specific payments; simply checking unspent conditions on the ledger is not sufficient in case of concurrent sessions. Φ_2 is important because the coin claimed by the seller is jointly constructed with the adversary, so we need to ensure that there is no other way to spend it. The following is proved automatically with Tamarin [35]:

Proposition 1. For Seller of Figure 2, $Q_{keys}, \mathcal{L}_{hash}, \mathcal{I}_{hash}, Q_{func}, \text{Seller}; \Psi_{coins} \models \Phi_S$

ZKCP Buyer. As we can see in the \mathcal{L}_{hash} -based protocol presented above, in order to ensure the witness delivery from a ZKCP protocol, the buyer should perform some verification actions on the data (e.g. zero-knowledge proofs) received during the protocol execution. We model these checks by a formula $\Psi_1(x, \bar{x}_{state})$, where x represents the desired result for the function of interest, and \bar{x}_{state} represents protocol data that is relevant for buyer's verification actions. Ψ_1 and \bar{x}_{state} are protocol specific and they are parameters of our definition.

In addition to data received during the protocol execution, the buyer can also rely on data that is published on the ledger, and on the associated constraints that are ensured by the ledger semantics. We model these by $\Psi_0(y, \bar{x}_{state})$ where y represents the relevant ledger data. For example, in the \mathcal{L}_{hash} -based protocol, the semantics of the ledger ensures that the data y associated to the transaction that spends the hashcoin must contain the preimage of a hash recorded in \bar{x}_{state} , if the coin was spent by any party other than the buyer. A part of our security definition will require that Ψ_0 in conjunction with Ψ_1

does indeed reveal the witness. A second part of the definition will require that, if the buyer performed a payment transaction, then the buyer and the ledger will reach a state where Ψ_0 and Ψ_1 hold, or otherwise the buyer can obtain a refund.

Definition 5. A ZKCP Buyer specification is given by a set of protocol rules that contains the special rule **pay**: $[\dots] \dashv\vdash \text{Pay}(t_{\text{pk}}, t_{\text{res}}, t_{\text{coin}}, \bar{u}_{\text{state}}) \dashv\vdash [\dots]$.

The *pay* rule models the invocation of a payment transaction for a witness, where t_{pk} is the public key of the buyer, t_{res} is the desired result, t_{coin} is the target coin where the buyer makes the payment, and \bar{u}_{state} is state information that is relevant for obtaining the witness. See Fig. 2 for the Buyer specification in the protocol described above.

Definition 6. Let (Q, Ψ) be a ZKCP-context and \mathcal{B} be a ZKCP Buyer specification. We say that these ensure buyer security if $Q, \mathcal{B}; \Psi \models \Phi_B$, where Φ_B is defined in Figure 3.

Intuitively, the formulas Φ_0, Φ_1, Φ_2 from Definition 6 ensures that:

- Φ_0 : if the buyer has paid for a witness into a coin, then spending that coin on the ledger will either lead to a refund, i.e. $z_{\text{pk}} = x_{\text{pk}}$, or else the data y associated to the spending transaction together with buyer state data satisfy the constraint Ψ_0 ;
- Φ_1 : before paying, the buyer performs checks entailing the constraint Ψ_1 for the desired result and the buyer state;
- Φ_2 : Ψ_0 and Ψ_1 allow to derive a witness for the desired result, by combining transaction data y with data \bar{x}_{state} gathered from the protocol execution.

Proposition 2. For Buyer from Figure 2 and $Q = (Q_{\text{keys}}, \mathcal{L}_{\text{hash}}, \mathcal{I}_{\text{hash}}, Q_{\text{func}})$, we have

$$Q, \text{Buyer}; \Psi_{\text{coins}} \models \Phi_B \left\{ \begin{array}{l} \bar{x}_{\text{state}} : (x_{\pi}, x_{\text{ew}}, x_h, x_{\text{pks}}) \\ \Psi_0(y, \bar{x}_{\text{state}}) : \exists y_s, y_h. y \approx \langle y_s, y_h \rangle \wedge x_h \approx h(y_h) \\ \Psi_1(x_{\text{res}}, \bar{x}_{\text{state}}) : \text{ver}_{\text{zk}}(x_{\pi}, x_{\text{ew}}, x_{\text{res}}, x_h, x_{\text{pks}}) \approx \text{ok} \end{array} \right.$$

We prove Φ_0 from Φ_B with Tamarin [35]. The properties Φ_1 and Φ_2 are simple local deduction properties that can be checked by hand (if the state of the buyer would be more complex, automated tools can also be used for that).

Observations: • the seller (\mathcal{S}) and buyer (\mathcal{B}) public keys are linked on the ledger, while this is not a necessary consequence of the security properties. \mathcal{S} does not need to know the public key of \mathcal{B} in advance, while \mathcal{B} does need the public key of \mathcal{S} .

• private ledger keys of \mathcal{S} and \mathcal{B} do not have to be secret for security to hold: our models allow corruption of any key by the adversary (\mathcal{A}). For \mathcal{S} , security follows from the fresh symmetric key created for each session and, for \mathcal{B} , from the trusted ledger. Note, however, that these keys allow \mathcal{A} to spend the coins of their owner, but this is independent from the ZKCP protocol. In fact, a basic property of *any* ledger-based protocol should be that it does not reveal secret keys, i.e. $\forall x, i, j. !\text{Key}(x) @ i \wedge \text{K}(x) @ j \Rightarrow \exists \ell. \ell < j \wedge \text{Corrupt}(g(x)) @ \ell$. We also prove this property in Tamarin for our models.

• \mathcal{S} cannot reuse the same symmetric key and zero-knowledge proof in two different sessions, even if those sessions are for selling the same witness; • our intruder deduction rules assume a perfect zero-knowledge construction, in particular \mathcal{A} cannot tweak the proof parameters in order to reveal the witness, as exploited by attacks of [16]. In the next section we show that intruder deduction rules can also model finer-grained

properties of cryptographic constructions if required, in particular conditions when the witness may be revealed; • security for \mathcal{S} depends on the timely delivery of transactions to the ledger, while this is not the case for \mathcal{B} , who could obtain both the witness and the money back if there was a time delay; • the proof x_π is not necessary for extracting the witness so it can be discarded after verification by \mathcal{B} ; • our models consider a strong \mathcal{A} and, as such, do not cover the case of weaker, multiple \mathcal{A} 's, e.g. for two different buyers that do not collude or do not control the network, but they can be extended to.

5 ZKCP protocol on the basecoin ledger

Managing hashcoins - e.g. applying the hashing algorithm - sets tradeoffs for the agents that maintain the ledger; they may give priority to standard coins, i.e. preferring $\mathcal{L}_{\text{base}}$ over $\mathcal{L}_{\text{hash}}$. Another constraint that needs to be taken into account - by parties engaging in ZKCP - is the complexity of constructing and verifying the zero-knowledge proofs. In this section, we formalize and analyze the protocol of [15], which aims to implement the ZKCP functionality on $\mathcal{L}_{\text{base}}$. Other works, e.g. [18], aim to minimize the zk burden by appealing to special contracts that will be executed only in case of dispute.

Cryptographic primitives. For ZKCP on $\mathcal{L}_{\text{base}}$, [15] adopts timed cryptographic commitments [33, 34], as presented in Example 8, in order to emulate the ledger time-out. To link ledger transitions and data release, [15] exploits algebraic properties of the ECDSA signature used in Bitcoin: relying on homomorphic encryption, e.g. Paillier, an encrypted signature can be constructed from an encryption of the signing key, which can be constructed by adding shares of the signing key on top of an initial encrypted share [36–39]. A Diffie-Hellman group is used to establish a shared key. A special type of zk proof is also needed: a prover can encode the witness and convince the verifier that it can be extracted as soon as some committed structured data - for ZKCP: an ECDSA signature - is revealed. We rely on $\mathcal{I}_{\text{base}}$ from Figure 4 to model these crypto primitives. A term $\text{esign}(m, k, r_1, g(r_1 * r_2), pk(z))$ represents an encrypted partial signature of a message m , with signing key k , randomness share r_1 , public randomness $g(r_1 * r_2)$, and encryption public key $pk(z)$. Combining it with the decryption key z and the complementary randomness share r_2 , one can compute $\text{sign}(m, k, r_1 * r_2)$. The rules for extract and ver_{zk} model the connection between a valid signature and witness extraction. Time commitments can be checked wrt the public part $g(x)$ of private data x .

Fig. 4. Intruder theory $\mathcal{I}_{\text{base}}$; and $\forall f \in \mathcal{F}^{(k)}. [\mathbf{K}(x_1), \dots, \mathbf{K}(x_k)] \Rightarrow [\mathbf{K}(f(x_1, \dots, x_k))]$

$\text{Hom}_{\{\text{g, enc}\}} : [\mathbf{K}(g(x)), \mathbf{K}(y)] \Rightarrow [\mathbf{K}(g(x * y))] \quad [\mathbf{K}(\text{enc}(x, z)), \mathbf{K}(y)] \Rightarrow [\mathbf{K}(\text{enc}(x * y, z))]$ $\text{AG} : x * i(x) = 1, x * 1 = x, x * y = y * x, (x * y) * z = x * (y * z)$ $\mathcal{R}_0 : \text{homs}(\text{enc}(k, y), m, r_1, r) \rightarrow \text{esign}(m, k, r_1, r, y) \quad \text{dec}(\text{enc}(x, pk(y)), y) \rightarrow x$ $\text{decs}(\text{esign}(m, k, r_1, g(r_1 * r_2), pk(z)), r_2, z) \rightarrow \text{sign}(m, k, r_1 * r_2)$ $\text{ver}(\text{sign}(x, y, z), x, g(y)) \rightarrow \text{ok} \quad \text{open}(\text{com}(x, r), r) \rightarrow x \quad \text{extract}(\text{zk}(x, y, z), z) \rightarrow x$ $\text{ver}_{\text{tc}}(\text{tcom}(x), g(x)) \rightarrow \text{ok} \quad \text{ver}_{\text{zk}}(\text{zk}(x, f(x), \text{sign}(y, z, w)), f(x), y, g(z)) \rightarrow \text{ok}$
--

Jointly signing a message. Assume two parties A_1 (holding k_1, r_1) and A_2 (holding k_2, r_2) want to create $\text{sign}(t, k_1 * k_2, r_1 * r_2)$ for some agreed upon t . Then, say, A_1 can generate a fresh key pair $k, \text{pk}(k)$ and send $\text{enc}(k_1, \text{pk}(k))$ to A_2 . Relying on Hom_{enc} , A_2 can obtain $\text{enc}(k_1 * k_2, \text{pk}(k))$, which with $t, r_2, g(r_1 * r_2)$ as arguments to homs gives $\text{esign}(t, k_1 * k_2, r_2, g(r_1 * r_2), \text{pk}(k))$. Sent back to A_1 , the joint signature is derived by applying decs to this term and r_1, k . Note that A_1 gets the signature and can decide when to show it to A_2 . On the other hand, both parties contribute to randomness in the signature; no party can force a particular value for the randomness. Both of these features will be needed to ensure the security properties for the ZKCP protocol:

1) Based on DH key-exchange and commitments, compute a public key $pk_{12} = g(k_1 * k_2)$ such that the private key $k_1 * k_2$ is secret-shared between the seller (\mathcal{S}), who holds $k_1, g(k_2)$, and the buyer (\mathcal{B}), who holds $k_2, g(k_1)$. Similarly, secret-shared randomness $r_1 * r_2$ is computed: $\# \text{Public} : pk_{12}, g(r_1 * r_2) \quad \text{Seller} : k_1, r_1 \quad \text{Buyer} : k_2, r_2 \#$

2) The key pk_{12} is used for an intermediate transfer from \mathcal{B} to \mathcal{S} . The two agree on the transaction that transfers a coin from pk_{12} to \mathcal{S} : $\# \text{Public} : t = \langle \text{c2c}, \rho_{\text{sn}}^1, \rho_{\text{sn}}^2, g(\text{ks}) \rangle \#$, where $\rho_{\text{sn}}^1, \rho_{\text{sn}}^2$ are fresh public serial numbers and $g(\text{ks})$ is the public key of \mathcal{S} . This transaction is not signed, so cannot yet lead to a transfer. Also, \mathcal{B} has not yet transferred coins into pk_{12} .

3) Based on crypto as shown above, \mathcal{S} (with \mathcal{B} 's help) obtains $s = \text{sign}(t, k_1 * k_2, r_1 * r_2)$. \mathcal{S} checks that s is valid by applying the signature verification algorithm. It then outputs the zero-knowledge proof $\pi = \text{zk}(w, f(w), s)$ and a time commitment to \mathcal{S} 's share of the joint secret key: $\# \text{Seller} : s \quad \text{Public} : \pi, \text{tcom}(k_1) \#$

4) \mathcal{B} verifies the proof and the time commitment, and transfers a coin to pk_{12} , leading to an update of the ledger: $\# \text{Ledger} : !\text{Coin}(\rho_{\text{sn}}^1, pk_{12}) \#$

5) The seller claims ρ_{sn}^1 by invoking R_{c2c} on the ledger, relying on the signature s obtained previously. The ledger will record a $!\text{Spend}$ fact with the corresponding transaction data, including the signature: $\# \text{Ledger} : !\text{Spend}(\rho_{\text{sn}}^1, pk_{12}, s, g(\text{ks})) \#$

6) The buyer obtains s from the ledger and extracts the witness from the zk proof: $w = \text{extract}(\pi, s)$. If the seller aborted, no one can redeem the coin ρ_{sn}^1 , until the time commitment reveals k_1 , so the buyer can reconstruct $k_1 * k_2$ and redeem the coin. The formal specification is in Fig. 5, with details of joint signing omitted.

Proposition 3. For Seller and Buyer from Figure 5 and Q_{tcom} from Example 8,

$$Q, \text{Seller}; \Psi_{\text{coins}} \models \Phi_S \quad Q, \text{Buyer}; \Psi_{\text{coins}} \models \Phi_B \quad Q = (\text{Q}_{\text{keys}}, \text{Q}_{\text{tcom}}, \mathcal{L}_{\text{base}}, \mathcal{I}_{\text{base}}, \text{Q}_{\text{func}})$$

where $\bar{x}_{\text{state}} : \langle x_\pi, x_{\text{tcom}}, x_{pk}^{12} \rangle, \Psi_0(y, \bar{x}_{\text{state}}) : \exists z, x. x_\pi \approx \text{zk}(z, x, x_s) \wedge y \approx x_s;$
 $\Psi_1(x_{\text{res}}, \bar{x}_{\text{state}}) : \text{ver}_{\text{zk}}(x_\pi, x_{\text{res}}, x_{\text{tcom}}, x_{pk}^{12}) \approx \text{ok}$

Tamarin verification: we prove Φ_S and Φ_0 for Φ_B automatically with Tamarin relying on the reduction that we present in the next section for termination within 1 minute. We prove two helper lemmas along the way: 1) if the adversary knows a time commitment, then it either knows the committed message at an earlier time, or the commitment is constructed by an honest party; 2) fresh randoms and keys stay secret - unless opened by a time commitment. The Tamarin code is available online [35].

Observations: • as for $\mathcal{L}_{\text{hash}}$, the \mathcal{S} and \mathcal{B} are linked on the ledger; the secret keys of any party can be corrupted, we prove however that the protocol does not itself reveal these keys; • the cryptographic constructions from [15] are a particular instance of $\mathcal{I}_{\text{base}}$;

Fig. 5. ZKCP on $\mathcal{L}_{\text{base}}$; Seller = (S_0, \dots, S_4) ; Buyer = $(B_0, \dots, B_3, B_4^{\text{go}}, B_4^{\text{ab}})$

$$\begin{array}{l}
\hline
S_0: [!\text{Key}(x_{ks}), !\text{Witn}(x_{\text{wtn}})] \dashv [\text{Sell}(g(x_{ks}), x_{\text{wtn}})] \mapsto [\text{state}_0] \\
S_1: [\text{state}_0, \text{Fr}(k_1), \text{Fr}(r_1), \text{Fr}(r)] \Rightarrow [\text{Out}(\text{com}(g(k_1), r)), \text{Out}(g(r_1)), \text{state}_1)] \\
S_2: [\text{state}_1, \text{In}(y_{k_2}), \text{Fr}(k_e)] \Rightarrow [\text{Out}(r), \text{Out}(\text{enc}(k_1, \text{pk}(k_e))), \text{state}_2] \\
S_3: [\text{state}_2 [y_{k_2} = g(x_{k_2})], x_{\text{pk}}^{12} = g(x_{k_2} * k_1), c_k = \text{tcom}(k_1), x_\pi = \text{zk}(x_{\text{wtn}}, f(x_{\text{wtn}}, s))] \\
(\text{JointSign} \mapsto t = \langle \text{c2c}, \rho_{\text{sn}}^1, \rho_{\text{sn}}^2, g(x_{ks}) \rangle, s = \text{sign}(t, \dots)) \Rightarrow [\text{Out}(\langle c_k, x_\pi \rangle), \text{state}_3] \\
S_4: [\text{state}_3, !\text{Coin}(\rho_{\text{sn}}^1, x_{\text{pk}}^{12})] \dashv [\text{Unspent}(\rho_{\text{sn}}^1), \text{Claim}(g(x_{ks}), x_{\text{wtn}}, k_1, \rho_{\text{sn}}^1)] \mapsto [\text{Out}(\langle s, \rho_{\text{sn}}^2, g(x_{ks}) \rangle)] \\
\hline
B_0: [!\text{Res}(x_{\text{res}}), !\text{Key}(x_{\text{kb}}), !\text{Pk}(x_{\text{pks}}), !\text{Coin}(x_{\text{sn}}^0, g(x_{\text{kb}}))] \Rightarrow [\text{state}_0] \\
B_1: [\text{state}_0, \text{In}(\langle x_{ck}, y_{r_1} \rangle), \text{Fr}(k_2), \text{Fr}(r_2)] \Rightarrow [\text{Out}(\langle g(k_2), g(r_2) \rangle), \text{state}_1] \\
B_2: [\text{state}_1 [x_{ck} = \text{com}(g(x_{k_1}), x_r), y_{r_1} = g(x_{r_1})], \text{In}(x_r), x_{\text{pk}}^{12} = g(x_{k_1} * k_2), x_r^{12} = g(x_{r_1} * r_2)] \\
(\text{JointSign} \mapsto t = \langle \text{c2c}, \rho_{\text{sn}}^1, \rho_{\text{sn}}^2, x_{\text{pks}} \rangle, s = \text{sign}(t, \dots)) \Rightarrow [\text{state}_2] \\
B_3: [\text{state}_2, \text{In}(\langle x_{\text{tcom}}, x_\pi \rangle), \text{Fr}(r)] \dashv [\text{ver}_{\text{zk}}(x_\pi, x_{\text{res}}, t, x_{\text{pk}}^{12}) \approx \text{ok}, \text{ver}_{\text{tc}}(x_{\text{tcom}}, g(x_{k_1})) \approx \text{ok}, \\
\text{Pay}(g(x_{\text{kb}}), x_{\text{res}}, \rho_{\text{sn}}^1, \langle x_\pi, x_{\text{tcom}}, x_{\text{pk}}^{12} \rangle)] \mapsto \\
[\text{Out}(\langle \text{sign}(\langle \text{c2c}, x_{\text{sn}}^0, \rho_{\text{sn}}^1, x_{\text{pk}}^{12} \rangle, x_{\text{kb}}, r), \rho_{\text{sn}}^1, x_{\text{pk}}^{12} \rangle), \text{state}_3] \\
B_4^{\text{go}}: [\text{state}_3, !\text{Spend}(\rho_{\text{sn}}^1, z, s, x_{\text{pks}}), x_{\text{wtn}} = \text{extract}(x_\pi, s)] \dashv [x_{\text{res}} \approx f(x_{\text{wtn}}), \text{Witness}(x_{\text{res}})] \mapsto [] \\
B_4^{\text{ab}}: [\text{state}_3, !\text{Coin}(\rho_{\text{sn}}^1, g(x_k^{12})), \text{In}(x_{k_1}), \text{Fr}(r), x_k^{12} = x_{k_1} * k_2, x_s = \text{sign}(\langle \rho_{\text{sn}}^1, \rho_{\text{sn}}^2, g(x_{\text{kb}}) \rangle, x_k^{12}, r)] \\
\dashv [x_{\text{tcom}} \approx \text{tcom}(x_{k_1}), \text{Unspent}(\rho_{\text{sn}}^1)] \mapsto [\text{Out}(\langle x_s, \rho_{\text{sn}}^2, g(x_{\text{kb}}) \rangle)] \\
\hline
\end{array}$$

it may admit more efficient instances, and our proofs could still be relied on for the security guarantees; • $\mathcal{I}_{\text{base}}$ does not cover the full algebra of homomorphic encryption, where we have $[\text{K}(\text{enc}(x, z)), \text{K}(\text{enc}(y, z))] \Rightarrow [\text{K}(\text{enc}(x * y, z))]$. It is however sound when every ciphertext constructed by honest parties uses a fresh key, as in our case study; covering the full theory is a long-standing, still open, problem for protocol verification • the same shared key could be used for the exchange of several witnesses within the timeframe chosen for the time commitment; • contrary to $\mathcal{L}_{\text{hash}}$, the zero-knowledge proof cannot be discarded by \mathcal{B} after verification, since it is necessary for extracting the witness; • on $\mathcal{L}_{\text{hash}}$, \mathcal{B} sets the ledger timeout and \mathcal{S} can accept to proceed; on $\mathcal{L}_{\text{base}}$ it is the other way around with respect to crypto timeout.

6 Homomorphism and abelian group reduction

We take a class of intruder theories that covers the one of Fig. 4; \mathcal{F} contains a set of homomorphic functions \mathcal{F}_{hom} . We reduce any \mathcal{I} from this class to \mathcal{I}_Δ such that: \mathcal{I}_Δ is simpler than \mathcal{I} ; \mathcal{I}_Δ is sound wrt \mathcal{I} . First, given any trace τ wrt \mathcal{I} , we show that there is \mathcal{I}_Δ generating τ and where: (i) the homomorphic properties are restricted by arguments from honest parties in τ ; (ii) the abelian group is degenerated, allowing to obtain any factors from products. Second, we augment any set of rules \mathcal{S} to \mathcal{S}_Δ , which records as facts the homomorphic arguments of \mathcal{S} , and \mathcal{I}_Δ is generalized to cover any trace of \mathcal{S}_Δ .

Definition 7. A base for \mathcal{F} is a function Δ with $\text{dom}(\Delta) = \mathcal{F}_{\text{hom}}$ and $\forall f \in \mathcal{F}_{\text{hom}}^{(n)}$. $\Delta(f) \subseteq \mathcal{T}^n$. We assume that Δ is closed modulo AC, i.e. $\Delta_f(u * v, \bar{w}) \Rightarrow \Delta_f(v * u, \bar{w})$ and similarly for associativity, and closed by: $\Delta_f(u * v, \bar{w}) \Rightarrow \Delta_f(u, \bar{w})$.

We extend intruder deduction to rules of the form $[\Delta_f(\bar{x}), M] \Rightarrow [N]$, which have the same semantics as $[M] \Rightarrow [N]$ with the additional constraint that $\bar{x}\theta \in \Delta(f)$ holds for the substitution θ that instantiates the rule.

Definition 8. We consider the class of intruder theories as defined below (left):

Initial theory \mathcal{I} (with Hom for all $f \in \mathcal{F}_{\text{hom}}$)	Reduced theory \mathcal{I}_Δ for base Δ
$\text{Hom} : [K(f(x, \bar{z})), K(y)] \Rightarrow [K(f(x * y, \bar{z}))]$ $\text{AG} : x * i(x) = 1, x * 1 = x$ $x * y = y * x, (x * y) * z = x * (y * z)$ $\mathcal{R}_0 : \{l_1 \rightarrow r_1, \dots, l_k \rightarrow r_k\}$	$\text{Hom}_\Delta : [\Delta_f(x, \bar{z}), K(y)] \Rightarrow [K(f(x * y, \bar{z}))]$ $\text{AP} : [K(x * y)] \Rightarrow [K(x)]$ $x * y = y * x, (x * y) * z = x * (y * z)$ $\mathcal{R}_0 : \{l_1 \rightarrow r_1, \dots, l_k \rightarrow r_k\}$

We assume that every $l \rightarrow r \in \mathcal{R}_0$ satisfies

H1: $\text{top}(l), \text{top}(r) \notin \mathcal{F}_{\text{hom}} \cup \{*, i\}$ **H2:** $\forall t \in \text{st}(r) \setminus \text{st}(l). \text{top}(t) \cap (\mathcal{F}_{\text{hom}} \cup \{*, i\}) = \emptyset$
Given such \mathcal{I} and a base Δ , we define the reduced theory \mathcal{I}_Δ as above (right). $\mathcal{I}, \mathcal{I}_\Delta$ also contain the deduction rules $\forall f. [K(x_1), \dots, K(x_k)] \Rightarrow [K(f(x_1, \dots, x_k))]$.

H1 and H2 help in proofs [40]; \mathcal{R}_0 from Figure 4 respects them. Intuitively, we split the homomorphic argument of f in two parts, e.g. $f(u * v, \bar{w})$, where the factors of v are known by the adversary, while the factors of u are provided by honest parties (in \mathcal{S}). When the adversary applies Hom to such a term, to get e.g. $f(u * v * t, \bar{w})$, there is a *smaller* term $f(u, \bar{w})$ that can be used to obtain the same result, since the adversary knows $v * t$. The term u will be added by \mathcal{S}_Δ to $\Delta(f)$ so Hom_Δ can be applied on it.

Proposition 4. For any \mathcal{S}, M_0, M_1 s.t. M_1 can be derived from M_0 using rules in $\mathcal{S} \cup \mathcal{I}$, there is Δ s.t. M_1 can be derived from M_0 using $\mathcal{S} \cup \mathcal{I}_\Delta$; Δ can be iteratively constructed by a set \mathcal{S}_Δ - augmenting each rule in \mathcal{S} with a constant number of facts.

Corollary 1. For any \mathcal{S} and formulas Ψ, Φ , we have $\mathcal{S}_\Delta, \mathcal{I}_\Delta; \Psi \vdash \Phi \Rightarrow \mathcal{S}, \mathcal{I}; \Psi \vdash \Phi$

Scope. The reduction is sound for any set of protocol rules. However, since \mathcal{I}_Δ allows to freely decompose products, it gives too much power to the adversary (leading to false attacks) for certain classes of protocols, e.g. when a nonce r protects a secret s in $s * r$. The reduction is useful for proofs only when secret data is protected by (homomorphic) cryptographic constructions, e.g. exponentiation, encryption, etc.

7 Related and future work

Several works extend the scope of Tamarin to new cryptographic primitives [41–43] or infrastructure features [44, 45]. Our models contribute to both of these directions. On the crypto side, an open question is to cover deductions like $\text{enc}(u, k), \text{enc}(v, k) \Rightarrow \text{enc}(u * v, k)$, which would allow to model e.g. homomorphic tallying for voting [46]. Protocol verification modulo this theory is studied in [47], where abstractions different from ours are used for reducing the theory, but the case studies are limited to unification problems and relatively simple protocols.

Works complementary to ours aim to provide formal guarantees for code executed on the blockchain [48–50]. Our ledger models are, on one hand, grounded on such guarantees and, on the other hand, they allow to reason about the properties of higher-level

protocols and applications. In future work, we can extend our models to cover more general smart contracts, hybrid ledgers and applications [18,32,51]. Current ZKCP protocols don't allow seller/buyer unlinkability, while the security properties leave scope for it. An open problem is ZKCP on ledgers with more privacy [52–54] and appropriate unlinkability notions.

Acknowledgment

The research leading to these results has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreements No 645865-SPOOC).

References

1. Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. 2008. Available at <https://bitcoin.org/bitcoin.pdf>.
2. Gavin Wood. Ethereum: A secure decentralised generalised transaction ledger, 2014. Available at <https://gavwood.com/paper.pdf>.
3. L. M. Goodman. Tezos - a self-amending crypto-ledger, 2014. Available at https://tezos.com/static/white_paper-2dc8c02267a8fb86bd67a108199441bf.pdf.
4. Timo Hanke, Mahnush Movahedi, and Dominic Williams. DFINITY technology overview series, consensus system. *CoRR*, abs/1805.04548, 2018.
5. Juan A. Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol: Analysis and applications. In *Advances in Cryptology - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT'15)*, volume 9057 of *Lecture Notes in Computer Science*, pages 281–310. Springer, 2015.
6. Rafael Pass, Lior Seeman, and Abhi Shelat. Analysis of the blockchain protocol in asynchronous networks. In *Advances in Cryptology - 36th Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT'17)*, volume 10211 of *Lecture Notes in Computer Science*, pages 643–673, 2017.
7. Mohammad Torabi Dashti and Sjouke Mauw. Fair exchange. In Burton Rosenberg, editor, *Handbook of Financial Cryptography and Security*, pages 109–132. Chapman and Hall/CRC, 2010.
8. N. Asokan, Victor Shoup, and Michael Waidner. Optimistic fair exchange of digital signatures (extended abstract). In *Advances in Cryptology - International Conference on the Theory and Application of Cryptographic Techniques (EUROCRYPT '98)*, volume 1403 of *Lecture Notes in Computer Science*, pages 591–606. Springer, 1998.
9. Christian Cachin and Jan Camenisch. Optimistic fair secure computation. In *Advances in Cryptology - 20th Annual International Cryptology Conference (CRYPTO'00)*, volume 1880 of *Lecture Notes in Computer Science*, pages 93–111. Springer, 2000.
10. Silvio Micali. Simple and fast optimistic protocols for fair electronic exchange. In *22nd ACM Symposium on Principles of Distributed Computing (PODC'03)*, pages 12–19. ACM, 2003.
11. Yehuda Lindell. Legally-enforceable fairness in secure two-party computation. In *Topics in Cryptology—CT-RSA 2008*, pages 121–137. Springer, 2008.

12. Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and Lukasz Mazurek. Fair two-party computations via bitcoin deposits. In *Financial Cryptography and Data Security Workshops (BITCOIN and WAHC'14)*, volume 8438 of *Lecture Notes in Computer Science*, pages 105–121. Springer, 2014.
13. Iddo Bentov and Ranjit Kumaresan. How to use bitcoin to design fair protocols. In *Advances in Cryptology - 34th Annual Cryptology Conference (CRYPTO'14)*, volume 8617 of *Lecture Notes in Computer Science*, pages 421–439. Springer, 2014.
14. Bitcoin wiki: Zero Knowledge Contingent Payment. https://en.bitcoin.it/wiki/Zero_Knowledge_Contingent_Payment.
15. Wacław Banasik, Stefan Dziembowski, and Daniel Malinowski. Efficient zero-knowledge contingent payments in cryptocurrencies without scripts. In *21st European Symposium on Research in Computer Security, Part II (ESORICS'16)*, volume 9879 of *Lecture Notes in Computer Science*, pages 261–280. Springer, 2016.
16. Matteo Campanelli, Rosario Gennaro, Steven Goldfeder, and Luca Nizzardo. Zero-knowledge contingent payments revisited: Attacks and payments for services. In *ACM SIGSAC Conference on Computer and Communications Security (CCS'17)*, pages 229–243. ACM, 2017.
17. Steven Goldfeder, Joseph Bonneau, Rosario Gennaro, and Arvind Narayanan. Escrow protocols for cryptocurrencies: How to buy physical goods using bitcoin. In *21st International Conference on Financial Cryptography and Data Security (FC'17)*, volume 10322 of *Lecture Notes in Computer Science*, pages 321–339. Springer, 2017.
18. Stefan Dziembowski, Lisa Ekey, and Sebastian Faust. Fairswap: How to fairly exchange digital goods. In *ACM SIGSAC Conference on Computer and Communications Security (CCS'18)*, pages 967–984. ACM, 2018.
19. Katriel Cohn-Gordon, Cas J. F. Cremers, and Luke Garratt. On post-compromise security. In *IEEE 29th Computer Security Foundations Symposium (CSF'16)*, pages 164–178. IEEE Computer Society, 2016.
20. Katriel Cohn-Gordon, Cas J. F. Cremers, Benjamin Dowling, Luke Garratt, and Douglas Stebila. A formal security analysis of the signal messaging protocol. In *IEEE European Symposium on Security and Privacy (EuroS&P'17)*, pages 451–466. IEEE Computer Society, 2017.
21. Karthikeyan Bhargavan, Bruno Blanchet, and Nadim Kobeissi. Verified models and reference implementations for the TLS 1.3 standard candidate. In *IEEE Symposium on Security and Privacy (SP'17)*, pages 483–502. IEEE Computer Society, 2017.
22. Cas Cremers, Marko Horvat, Jonathan Hoyland, Sam Scott, and Thyla van der Merwe. A comprehensive symbolic analysis of TLS 1.3. In *ACM SIGSAC Conference on Computer and Communications Security (CCS'17)*, pages 1773–1788. ACM, 2017.
23. Charlie Jacomme and Steve Kremer. An extensive formal analysis of multi-factor authentication protocols. In *31st IEEE Computer Security Foundations Symposium (CSF'18)*, pages 1–15. IEEE Computer Society, 2018.
24. Simon Meier, Benedikt Schmidt, Cas Cremers, and David A. Basin. The TAMARIN prover for the symbolic analysis of security protocols. In *25th International Conference on Computer Aided Verification (CAV'13)*, volume 8044 of *Lecture Notes in Computer Science*, pages 696–701. Springer, 2013.
25. Iliano Cervesato, Nancy A. Durgin, John C. Mitchell, Patrick Lincoln, and Andre Scedrov. Relating strands and multiset rewriting for security protocol analysis. In *13th IEEE Computer Security Foundations Workshop, CSFW '00, Cambridge, England, UK, July 3-5, 2000*, pages 35–51. IEEE Computer Society, 2000.
26. Benedikt Schmidt, Simon Meier, Cas J. F. Cremers, and David A. Basin. Automated analysis of diffie-hellman protocols and advanced security properties. In *25th IEEE Computer Security Foundations Symposium, (CSF'12)*, pages 78–94. IEEE Computer Society, 2012.

27. Nachum Dershowitz and Jean-Pierre Jouannaud. Rewrite systems. In *Handbook of Theoretical Computer Science, Volume B: Formal Models and Semantics (B)*, pages 243–320. MIT Press, 1990.
28. Serge Vaudenay. The security of DSA and ECDSA. In *6th International Workshop on Theory and Practice in Public Key Cryptography (PKC'03)*, volume 2567 of *Lecture Notes in Computer Science*, pages 309–323. Springer, 2003.
29. Bitcoin wiki: Hashed Timelock Contracts. https://en.bitcoin.it/wiki/Hashed_Timelock_Contracts.
30. Bitcoin wiki: Payment channels. https://en.bitcoin.it/wiki/Payment_channels.
31. Bitcoin wiki: Lightning Network. https://en.bitcoin.it/wiki/Lightning_Network.
32. Mike Hearn. Corda: A distributed ledger.
33. Ron L. Rivest, Adi Shamir, and David A. Wagner. Time-lock puzzles and timed-release crypto. Technical report, MIT, Cambridge, MA, USA, 1996.
34. Dan Boneh and Moni Naor. Timed commitments. In *Advances in Cryptology - 20th Annual International Cryptology Conference (CRYPTO'00)*, volume 1880 of *Lecture Notes in Computer Science*, pages 236–254. Springer, 2000.
35. Tamarin code for ZKCP protocol verification. <https://www.dropbox.com/sh/ahzbbobjm5z0e6a9/AAB6-Pz-RK3xwVznlaqaitfca?dl=0>.
36. Pascal Paillier. Public-key cryptosystems based on composite degree residuosity classes. In *Advances in Cryptology - International Conference on the Theory and Application of Cryptographic Techniques (EUROCRYPT'99)*, volume 1592 of *Lecture Notes in Computer Science*, pages 223–238. Springer, 1999.
37. Yehuda Lindell. Fast secure two-party ECDSA signing. In *Advances in Cryptology - 37th Annual International Cryptology Conference (CRYPTO'17)*, volume 10402 of *Lecture Notes in Computer Science*, pages 613–644. Springer, 2017.
38. Yehuda Lindell and Ariel Nof. Fast secure multiparty ECDSA with practical distributed key generation and applications to cryptocurrency custody. In *ACM SIGSAC Conference on Computer and Communications Security (CCS'18)*, pages 1837–1854. ACM, 2018.
39. Rosario Gennaro and Steven Goldfeder. Fast multiparty threshold ECDSA with fast trustless setup. In *ACM SIGSAC Conference on Computer and Communications Security (CCS'18)*, pages 1179–1194. ACM, 2018.
40. Additional material: Tamarin code and long paper version. <https://www.dropbox.com/sh/t74k3q4gxrm0pw/AADvx0e8WDaZgyf0OQF1El1Ca?dl=0>.
41. Benedikt Schmidt, Ralf Sasse, Cas Cremers, and David A. Basin. Automated verification of group key agreement protocols. In *IEEE Symposium on Security and Privacy (SP'14)*, pages 179–194, 2014.
42. Jannik Dreier, Lucca Hirschi, Sasa Radomirovic, and Ralf Sasse. Automated unbounded verification of stateful cryptographic protocols with exclusive OR. In *31st IEEE Computer Security Foundations Symposium, CSF'18*, pages 359–373. IEEE Computer Society, 2018.
43. Jannik Dreier, Charles Duménil, Steve Kremer, and Ralf Sasse. Beyond subterm-convergent equational theories in automated verification of stateful protocols. In *6th International Conference on Principles of Security and Trust (POST'17)*, volume 10204 of *Lecture Notes in Computer Science*, pages 117–140. Springer, 2017.
44. Steve Kremer and Robert Künnemann. Automated analysis of security protocols with global state. *Journal of Computer Security*, 24(5):583–616, 2016.
45. Michael Backes, Jannik Dreier, Steve Kremer, and Robert Künnemann. A novel approach for reasoning about liveness in cryptographic protocols and its application to fair exchange. In *IEEE European Symposium on Security and Privacy (EuroS&P'17)*, pages 76–91. IEEE Computer Society, 2017.

46. Olivier Baudron, Pierre-Alain Fouque, David Pointcheval, Jacques Stern, and Guillaume Poupard. Practical multi-candidate election system. In *20th annual (ACM) symposium on Principles of Distributed Computing (PODC'01)*, pages 274–283. ACM, 2001.
47. Fan Yang, Santiago Escobar, Catherine A. Meadows, José Meseguer, and Paliath Narendran. Theories of homomorphic encryption, unification, and the finite variant property. In *Proceedings of the 16th International Symposium on Principles and Practice of Declarative Programming, Kent, Canterbury, United Kingdom, September 8-10, 2014*, pages 123–133, 2014.
48. Massimo Bartoletti and Roberto Zunino. Bitml: A calculus for bitcoin smart contracts. In *ACM SIGSAC Conference on Computer and Communications Security (CCS'18)*, pages 83–100, 2018.
49. Karthikeyan Bhargavan, Antoine Delignat-Lavaud, Cédric Fournet, Anitha Gollamudi, Georges Gonthier, Nadim Kobeissi, Natalia Kulatova, Aseem Rastogi, Thomas Sibut-Pinote, Nikhil Swamy, and Santiago Zanella Béguelin. Formal verification of smart contracts. In *ACM Workshop on Programming Languages and Analysis for Security (PLAS@CCS'16)*, pages 91–96. ACM, 2016.
50. Everett Hildenbrandt, Manasvi Saxena, Nishant Rodrigues, Xiaoran Zhu, Philip Daian, Dwight Guth, Brandon M. Moore, Daejun Park, Yi Zhang, Andrei Stefanescu, and Grigore Rosu. KEVM: A complete formal semantics of the ethereum virtual machine. In *31st IEEE Computer Security Foundations Symposium (CSF'18)*, pages 204–217. IEEE Computer Society, 2018.
51. Stefan Dziembowski, Sebastian Faust, and Kristina Hostáková. General state channel networks. In *ACM SIGSAC Conference on Computer and Communications Security (CCS'18)*, pages 949–966. ACM, 2018.
52. Eli Ben-Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. Zerocash: Decentralized anonymous payments from bitcoin. In *IEEE Symposium on Security and Privacy, SP'14*, pages 459–474. IEEE Computer Society, 2014.
53. Guy Zyskind, Oz Nathan, and Alex Pentland. Enigma: Decentralized computation platform with guaranteed privacy. *CoRR*, abs/1506.03471, 2015.
54. Giulio Malavolta, Pedro Moreno-Sanchez, Aniket Kate, Matteo Maffei, and Srivatsan Ravi. Concurrency and privacy with payment-channel networks. In *ACM SIGSAC Conference on Computer and Communications Security (CCS'17)*, pages 455–471. ACM, 2017.