

Hawk: The Blockchain Model of Cryptography and Privacy-Preserving Smart Contracts

Ahmed Kosba

Andrew Miller

Elaine Shi

Zikai Wen

Charalampos Papamanthou

University of Maryland and Cornell University

ABSTRACT

Emerging smart contract systems over decentralized cryptocurrencies allow mutually distrustful parties to transact safely with each other without trusting a third-party intermediary. In the event of contractual breaches or aborts, the decentralized blockchain ensures that other honest parties obtain commensurate remuneration. Existing systems, however, lack transactional privacy. All transactions, including flow of money between pseudonyms and amount transacted, are exposed in the clear on the blockchain.

We present *Hawk*, a decentralized smart contract system that does not store financial transactions in the clear on the blockchain, thus retaining transactional privacy from the public’s view. A *Hawk* programmer can write a private smart contract in an intuitive manner without having to implement cryptography, and our compiler automatically generates an efficient cryptographic protocol where contractual parties interact with the blockchain, using cryptographic primitives such as succinct zero-knowledge proofs.

To formally define and reason about the security of our protocols, we are the first to formalize the blockchain model of secure computation. The formal modeling is of independent interest. We advocate the community to adopt such a formal model when designing interesting applications atop decentralized blockchains.

1. INTRODUCTION

Decentralized cryptocurrencies such as Bitcoin [45] and altcoins [18] have rapidly gained popularity, and are often quoted as a glimpse into our future [5]. These emerging cryptocurrency systems build atop a novel *blockchain* technology where *miners* run distributed consensus whose security is ensured if no adversary wields a large fraction of the computational (or other forms of) resource. The terms “blockchain” and “miners” are therefore often used interchangeably.

Blockchains like Bitcoin reach consensus not only about a stream of *data* (e.g., payment transactions), but also about *computations* involving this data (e.g., applying transactional semantics and updating the ledger). When we generalize the blockchain’s computation to arbitrary Turing-complete logic, we obtain an expressive *smart contract* system such as the soon-to-be-launched Ethereum [52]. Smart contracts are programs executed by all miners. Assuming that the decentralized consensus protocol is secure, smart contracts can be thought of as a conceptual party (in reality decentralized) that can be *trusted for correctness but not for privacy*. Specifically, a blockchain’s data and computation are publicly visible. Further, since the blockchain advances

in well-defined block intervals defined by the “mining” process, a discrete notion of *time* exists.

These features combined make the blockchain a new computation model empowered to enforce notions of *financial fairness* even in the presence of aborts. As is well-known in the cryptography literature, fairness against aborts is in general impossible in standard models of interactive protocol execution (e.g., secure multi-party computation), when the majority of parties can be corrupted [8, 16, 23]. In the blockchain model, protocol aborts can be made evident by *timeouts*. Therefore, the blockchain can enforce financial remuneration to the honest counterparties in the presence of protocol breach or aborts. In summary, decentralized smart contracts allow parties mutually unbeknownst to transact securely with each other, without trusting any central intermediary or incurring high legal and transactional cost.

Despite the expressiveness and power of blockchain and smart contracts, the present form of these technologies *lacks transactional privacy*. The entire sequence of actions taken in a smart contract are propagated across the network and/or recorded on the blockchain, and therefore are publicly visible. Even though parties can create new pseudonymous public keys to increase their anonymity, the values of all transactions and balances for each (pseudonymous) public key are publicly visible. Further, recent works have also demonstrated deanonymization attacks by analyzing the transactional graph structures of cryptocurrencies [40, 48]. The privacy requirements of many financial transactions will likely preclude the use of existing smart contract systems. Although there has been progress in designing privacy-preserving cryptocurrencies such as Zerocash [10] and several others [24, 41, 50], these systems forgo programmability, and it is unclear *a priori* how to enable programmability without exposing transactions and data in cleartext to miners.

1.1 Hawk Overview

We propose *Hawk*, a framework for building privacy-preserving smart contracts. With *Hawk*, a *non-specialist* programmer can easily write a *Hawk* program without having to implement any cryptography. Our *Hawk* compiler automatically compiles the program to a cryptographic protocol between the blockchain and the users. As shown in Figure 1, a *Hawk* program contains two parts:

- A *private contract* program denoted ϕ_{priv} which takes in parties’ input data (e.g., choice in a “rock, paper, scissors” game) as well as currency units (e.g., bids in an auction). The program ϕ_{priv} performs computation to determine the payout distribution amongst the parties. For exam-




Figure 1: Hawk overview.

ple, in an auction, winner’s bid goes to the seller, and others’ bids are refunded. The private contract ϕ_{priv} is meant to protect the participants’ data and the exchange of money.

- A *public contract* program denoted ϕ_{pub} that does not touch private data or money.

Our Hawk compiler will compile the public contract to execute directly on the public blockchain, and compile the private contract ϕ_{priv} into a cryptographic protocol involving the following pieces: *i*) protocol logic to be executed by the blockchain; and *ii*) protocol logic to be executed by contractual parties.

Security guarantees. Hawk’s security guarantees encompass two aspects:

- *On-chain privacy.* On-chain privacy stipulates that transactional privacy be provided against the public (i.e., against any party *not* involved in the contract) – unless the contractual parties themselves voluntarily disclose information. Although in Hawk protocols, users exchange data with the blockchain, and rely on it to ensure fairness against aborts, the flow of money and amount transacted in the private contract ϕ_{priv} is cryptographically hidden from the public’s view.
- *Contractual security.* While on-chain privacy protects contractual parties’ privacy against the public (i.e., parties not involved in the contract), contractual security protects parties in the same contractual agreement from *each other*. Hawk assumes that contractual parties act *selfishly* to maximize their own financial interest. In particular, they can *arbitrarily* deviate from the prescribed protocol or even *abort* prematurely. Therefore, contractual security is a multi-faceted notion that encompasses not only cryptographic notions of confidentiality and authenticity, but also financial fairness in the presence of cheating and aborting behavior. The best way to understand contractual security is through a concrete example, and we refer the reader to Section 1.2 for a more detailed explanation.

Minimally trusted manager. Hawk-generated protocols assume a special contractual party called a manager (e.g., an auction manager) besides the normal users. The manager aids the efficient execution of our cryptographic protocols while being minimally trusted. A skeptical reader may worry that our use of such a party trivializes our security guarantees; we dispel this notion by directly explaining what a corrupt manager can and cannot do: Although a manager can see the transactions that take place in a private contract, *it cannot affect the outcome of the contract, even when it colludes with other users*. In the event that a manager aborts the protocol, it can be financially penalized, and users obtain remuneration accordingly. The manager

```

1 HawkDeclareParties(Seller, /* N parties */);
2 HawkDeclareTimeouts(/* hardcoded timeouts */);

3 // Private contract  $\phi_{\text{priv}}$ 
4 private contract auction(Inp &in, Outp &out) {
5   int winner = -1;
6   int bestprice = -1;
7   int secondprice = -1;

8   for (int i = 0; i < N; i++) {
9     if (in.party[i].$val > bestprice) {
10      secondprice = bestprice;
11      bestprice = in.party[i].$val;
12      winner = i;
13     } else if (in.party[i].$val > secondprice) {
14       secondprice = in.party[i].$val;
15     }
16   }

17   // Winner pays secondprice to seller
18   // Everyone else is refunded
19   out.Seller.$val = secondprice;
20   out.party[winner].$val = bestprice - secondprice;
21   out.winner = winner;
22   for (int i = 0; i < N; i++) {
23     if (i != winner)
24       out.party[i].$val = in.party[i].$val;
25   }
26 }

27 // Public contract  $\phi_{\text{pub}}$ 
28 public contract deposit {
29   // Manager deposits $N
30   def check():
31     send $N to Manager
32   def managerTimeOut():
33     for (i in range($N)):
34       send $1 to party[i]
35 }

```

Figure 2: Hawk contract for a second-price sealed auction. Code described in this paper is an approximation of our real implementation. In the public contract, the syntax “send $\$N$ to P ” corresponds to the following semantics in our cryptographic formalism: $\text{ledger}[P] := \text{ledger}[P] + \N – see Section 2.

also need not be trusted to maintain the security or privacy of the underlying currency (e.g., it cannot double-spend, inflate the currency, or deanonymize users). Furthermore, if multiple contract instances run concurrently, each contract may specify a different manager and the effects of a corrupt manager are confined to that instance. Finally, the manager role may be instantiated with trusted computing hardware like Intel SGX, or replaced with a multiparty computation among the users themselves, as we describe in Section 4.3.

1.2 Example: Sealed Auction

Example program. Figure 2 shows a Hawk program for implementing a sealed, second-price auction where the highest bidder wins, but pays the second highest price. Second-price auctions are known to incentivize truthful bidding under certain assumptions, [51] and it is important that bidders submit bids without knowing the bid of the other people. Our example auction program contains a private contract ϕ_{priv} that determines the winning bidder and the price to be paid; and a public contract ϕ_{pub} that relies on public deposits to protect bidders from an aborting manager.

Contractual security requirements. Hawk will compile this auction program to a cryptographic protocol. As mentioned earlier, as long as the bidders and the manager do not voluntarily disclose information, transaction privacy is maintained against the public. Hawk also guarantees the following contractual security requirements for parties in the contract:

- *Input independent privacy.* Each user does not see others' bids before committing to their own. This way, users bids are independent of others' bids. Hawk guarantees input independent privacy even against a malicious manager.
- *Posterior privacy.* As long as the manager does not disclose information, users' bids are kept private from each other (and from the public) even after the auction.
- *Financial fairness.* If a party aborts or if the auction manager aborts, the aborting party should be financially penalized while the remaining parties receive compensation. Such fairness guarantees are not attainable in general by off-chain only protocols such as secure multi-party computation [7, 16]. As explained later, Hawk offers built-in mechanisms for enforcing refunds of private bids after certain timeouts. Hawk also allows the programmer to define additional rules, as part of the Hawk contract, that govern financial fairness.
- *Security against a dishonest manager.* We ensure *authenticity* against a dishonest manager: besides aborting, a dishonest manager cannot affect the outcome of the auction and the redistribution of money, even when it colludes with a subset of the users. We stress that to ensure the above, input independent privacy against a faulty manager is a prerequisite. Moreover, if the manager aborts, it can be financially penalized, and the participants obtain corresponding remuneration.

An auction with the above security and privacy requirements cannot be trivially implemented atop existing cryptocurrency systems such as Ethereum [52] or Zerocash [10]. The former allows for programmability but does not guarantee transactional privacy, while the latter guarantees transactional privacy but at the price of even reduced programmability than Bitcoin.

Aborting and timeouts. Aborting are dealt with using timeouts. A Hawk program such as Figure 2 declares timeout parameters using the `HawkDeclareTimeouts` special syntax. Three timeouts are declared where $T_1 < T_2 < T_3$:

T_1 : The Hawk contract stops collecting bids after T_1 .

T_2 : All users should have opened their bids to the manager within T_2 ; if a user submitted a bid but fails to open by T_2 , its input bid is treated as 0 (and any other potential input data treated as \perp), such that the manager can continue.

T_3 : If the manager aborts, users can reclaim their private bids after time T_3 .

The public Hawk contract ϕ_{pub} can additionally implement incentive structures. Our sealed auction program redistributes the manager's public deposit if it aborts. More specifically, in our sealed auction program, ϕ_{pub} defines two functions, namely `check` and `managerTimeOut`. The `check` function will be invoked when the Hawk contract completes execution within T_3 , i.e., manager did not abort. Otherwise,

if the Hawk contract does not complete execution within T_3 , the `managerTimeOut` function will be invoked. We remark that although not explicitly written in the code, all Hawk contracts have an implicit default entry point for accepting parties' deposits – these deposits are withheld by the contract till they are redistributed by the contract. Bidders should check that the manager has made a public deposit before submitting their bids.

Additional applications. Besides the sealed auction example, Hawk supports various other applications. We give more sample programs in Section 5.2.

1.3 Contributions

To the best of our knowledge, Hawk is the *first* to simultaneously offer transactional privacy and programmability in a decentralized cryptocurrency system.

Formal models for decentralized smart contracts. We are among the *first* ones to initiate a formal, academic treatment of the blockchain model of cryptography. We present a formal, UC-compatible model for the blockchain model of cryptography – this formal model is of independent interest, and can be useful in general for defining and modeling the security of protocols in the blockchain model. Our formal model has also been adopted by the Gyges work [35] in designing criminal smart contracts.

In defining for formal blockchain model, we rely on a notion called *wrappers* to modularize our protocol design and to simplify presentation. Wrappers handle a set of common details such as *timers*, *pseudonyms*, *global ledgers* in a centralized place such that they need not be repeated in every protocol.

New cryptography suite. We implement a new cryptography suite that binds private transactions with programmable logic. Our protocol suite contains three essential primitives `freeze`, `compute`, and `finalize`. The `freeze` primitive allows parties to commit to not only normal data, but also coins. Committed coins are frozen in the contract, and the payout distribution will later be determined by the program ϕ_{priv} . During `compute`, parties open their committed data and currency to the manager, such that the manager can compute the function ϕ_{priv} . Based on the outcome of ϕ_{priv} , the manager now constructs new private coins to be paid to each recipient. When the manager submits both the new private coins as well as zero-knowledge proofs of their well-formedness, the previously frozen coins are now redistributed among the users. Our protocol suite strictly generalizes Zerocash since Zerocash implements only private money transfers between users without programmability.

We define the security of our primitives using ideal functionalities, and formally prove security of our constructions under a simulation-based paradigm.

Implementation and evaluation. We built a Hawk prototype and evaluated its performance by implementing several example applications, including a *sealed-bid auction*, a *“rock, paper, scissors”* game, a *crowdfunding* application, and a *swap financial instrument*. We demonstrate interesting protocol optimizations that gained us a factor of **18**× in performance relative to a straightforward implementation. We show that for at about 100 parties (e.g., auction and crowdfunding), the manager's cryptographic computation (the most expensive part of the protocol) is under **2.8min**

using 4 cores, translating to under \$0.13 of EC2 time. Further, all on-chain computation (performed by all miners) is very cheap, and under 17ms for all cases. We will open source our Hawk framework in the near future.

1.4 Background and Related Work

1.4.1 Background

The original Bitcoin offers limited programmability through a scripting language that is neither Turing-complete nor user friendly. Numerous previous endeavors at creating smart contract-like applications atop Bitcoin (e.g., lottery [7, 16], micropayments [4], verifiable computation [38]) have demonstrated the difficulty of retrofitting Bitcoin’s scripting language – this serves well to motivate a Turing-complete, user-friendly smart contract language.

Ethereum is the first Turing-complete decentralized smart contract system. With Ethereum’s imminent launch, companies and hobbyists are already building numerous smart contract applications either atop Ethereum or by forking off Ethereum, such as prediction markets [3], supply chain provenance [6], crowd-based fundraising [1], and security and derivatives trading [27].

Security of the blockchain. Like earlier works that design smart contract applications for cryptocurrencies, we rely on the underlying decentralized blockchain to be secure. Therefore, we assume the blockchain’s consensus protocol attains security when an adversary does not wield a large fraction of the computational power. Existing cryptocurrencies are designed with heuristic security. On one hand, researchers have identified attacks on various aspects of the system [28, 33]; on the other, efforts to formally understand the security of blockchain consensus have begun [32, 43].

Minimizing on-chain costs. Since every miner will execute the smart contract programs while verifying each transaction, cryptocurrencies including Bitcoin and Ethereum collect transaction fees that roughly correlate with the cost of execution. While we do not explicitly model such fees, we design our protocols to minimize on-chain costs by performing most of the heavy-weight computation off-chain.

1.4.2 Additional Related Works

Leveraging blockchain for financial fairness. A few prior works have explored how to leverage the blockchain technology to achieve fairness in protocol design. For example, Bentov et al. [16], Andrychowicz et al. [7], Kumaresan et al. [38], Kiayias et al. [36], as well as Zyskind et al. [54], show how Bitcoin can be used to ensure fairness in secure multi-party computation protocols. These protocols also perform off-chain secure computation of various types, but do not guarantee transactional privacy (i.e., hiding the currency flows and amounts transacted). For example, it is not clear how to implement our sealed auction example using these earlier techniques. Second, these earlier works either do not offer system implementations or provide implementations only for specific applications (e.g., lottery). In comparison, Hawk provides a generic platform such that non-specialist programmers can easily develop privacy-preserving smart contracts.

Smart contracts. The conceptual idea of programmable electronic “smart contracts” dates back nearly twenty years [49].

Besides recent decentralized cryptocurrencies, which guarantee authenticity but not privacy, other smart contract implementations rely on trusted servers for security [44]. Our work therefore comes closest to realizing the original vision of parties interacting with a trustworthy “virtual computer” that executes programs involving money and data.

Programming frameworks for cryptography. Several works have developed programming frameworks that take in high-level program as specifications and generate cryptographic implementations. For example, previous works have developed compilers for secure multi-party computation [17, 37, 39, 47], authenticated data structures [42], and (zero-knowledge) proofs [12, 30, 31, 46]. Zheng et al. show how to generate secure distributed protocols such as sealed auctions, battleship games, and banking applications [53]. These works support various notions of security, but none of them interact directly with money or leverage public blockchains for ensuring financial fairness. Thus our work is among the first to combine the “correct-by-construction” cryptography approach with smart contracts.

Concurrent work. In concurrent and independent work, Kiayias et al. propose a simplified blockchain model in the Universal Composability framework [36]. Many aspects of their model, such as a shared global ledger and clock, are equivalent to ours. However, our model also captures many more important details that Kiayias et al. overlook, which are necessary to express the full gamut of protocols in the blockchain model. For example, their model fails to capture the pseudonymity and anonymity aspects of the blockchain, both of which are crucial for applications such as Hawk (this paper) and Gyges [35] (which also adopts the formal models proposed here). Their model also fails to capture the fact that the ordering of transactions in the same round can be subject to adversarial manipulation. Therefore, our work is the first that aims to provide a full-fledged formal model for decentralized blockchains as embodied by Bitcoin, Ethereum, and many other popular decentralized cryptocurrencies.

2. THE BLOCKCHAIN MODEL OF CRYPTOGRAPHY

We describe a formal model for smart contract execution under which we give formal proofs of security for our protocols. Our model conforms to the Universal Composability (UC) framework [19].

Protocols in the blockchain model. Our model makes use of a special party (in reality decentralized) called the *contract*¹. that is **entrusted to enforce correctness but not privacy**. All messages sent to the contract and all of its internal states are publicly visible. During the protocol, users interact with the contract by exchanging messages (also referred to as transactions). Money can be expressed as special bits stored on the blockchain interpreted as a ledger. Our contracts can access and update the ledger to implement money transfers between users (as represented by their pseudonymous public keys).

First, our model allow us to easily capture the time and

¹Disambiguation: the contract C in our cryptography formalism always refers to the native contract executed on the blockchain, not the user-facing Hawk program.

pseudonym features of cryptocurrencies. In cryptocurrencies such as Bitcoin and Ethereum, time progresses in block intervals, and a smart contract program is allowed to query the current time, and make decisions accordingly, e.g., make a refund operation after a timeout. Second, our model captures the role of a smart contract as a party trusted for correctness but not for privacy. Third, our formalism modularizes our notations by factoring out common specifics related to the smart contract execution model, and implementing these in central wrappers.

In a real-life cryptocurrency system such as Bitcoin or Ethereum, users can make up any number of identities by generating new public keys. In our formal model, for simplicity, we assume that there can be any number of identities in the system, and that they are fixed a priori. It is easy to extend our model to capture registration of new identities dynamically. As mentioned later, we allow each identity to generate an arbitrary (polynomial) number of pseudonyms.

2.1 Programs, Functionalities, and Wrappers

To make our notations simple for writing ideal functionalities and smart contracts, we make a conscious notational choice of introducing *wrapper* functionalities. Wrapper functionalities implement in a central place a set of common features (e.g., timer, ledger, pseudonyms) that are applicable to all ideal functionalities and contracts in our smart contract model of execution. In this way, we can modularize our notational system such that these common and tedious details need not be repeated in writing ideal functionalities and contract programs.

In particular, we will have two types of wrapper functionalities:

Contract functionality wrapper \mathcal{G} : A contract functionality wrapper $\mathcal{G}(C)$ takes in a *contract program* denoted C , and produces a contract functionality. Our real world protocols will be defined in the $\mathcal{G}(C)$ -hybrid world. Our contract functionality wrapper is formally presented in Figure 4.

Like the ideal functionality wrapper $\mathcal{F}(\cdot)$, the contract wrapper also implements standard features such as a timer, a global ledger, and money transfers between parties. In addition, the contract wrapper also models the specifics of the smart contract execution model. We point out the following important facts about the $\mathcal{G}(\cdot)$ wrapper:

- *Trusted for correctness but not privacy.* The contract functionality wrapper $\mathcal{G}(\cdot)$ stipulates that a smart contract is trusted for correctness but not for privacy. In particular, the contract wrapper exposes the contract’s internal state to any party that makes a query.
- *Time and batched processing of messages.* In popular decentralized cryptocurrencies such as Bitcoin and Ethereum, time progresses in block intervals marked by the creation of each new block. Intuitively, our $\mathcal{G}(\cdot)$ wrapper captures the following fact. In each round (i.e., block interval), the smart contract program may receive multiple messages (also referred to as transactions in the cryptocurrency literature). The order of processing these transactions is determined by the miner who mines the next block. In our model, we allow the adversary to specify an ordering of the messages collected in a round, and our contract program will then process the messages in this adversary-specified ordering.

- *Rushing adversary.* The contract wrapper $\mathcal{G}(\cdot)$ naturally captures a rushing adversary. Specifically, the adversary can first see all messages sent to the contract by honest parties, and then decide its own messages for this round, as well as an ordering in which the contract should process the messages in the next round. Modeling a rushing adversary is important, since it captures a class of well-known front-running attacks, e.g., those that exploit transaction malleability [10, 25]. For example, in a “rock, paper, scissors” game, if inputs are sent in the clear, an adversary can decide its input based on the other party’s input. An adversary can also try to maul transactions submitted by honest parties to potentially redirect payments to itself. Since our model captures a rushing adversary, we can write ideal functionalities that preclude such front-running attacks.

Ideal functionality wrapper \mathcal{F} : An ideal functionality $\mathcal{F}(\text{idealP})$ takes in an *ideal program* denoted idealP . Specifically, the wrapper $\mathcal{F}(\cdot)$ part defines standard features such as time, pseudonyms, a public ledger, and money transfers between parties. Our ideal functionality wrapper is formally presented in Figure 3.

Protocol wrapper Π : Our protocol wrapper allows us to modularize the presentation of user protocols. Our protocol wrapper is formally presented in Figure 5.

Terminology. For disambiguation, we always refer to the user-defined portions as *programs*. Programs alone do not have complete formal meanings. However, when programs are wrapped with functionality wrappers (including $\mathcal{F}(\cdot)$ and $\mathcal{G}(\cdot)$), we obtain functionalities with well-defined formal meanings. Programs can also be wrapped by a protocol wrapper Π to obtain a full protocol with formal meanings.

2.2 Modeling Time

At a high level, we express time in a way that conforms to the Universal Composability framework [19]. In the ideal world execution, time is explicitly encoded by a variable T in an ideal functionality $\mathcal{F}(\text{idealP})$. In the real world execution, time is explicitly encoded by a variable T in our contract functionality $\mathcal{G}(C)$. Time progresses in rounds. The environment \mathcal{E} has the choice of when to advance the timer.

We assume the following convention: to advance the timer, the environment \mathcal{E} sends a “tick” message to all honest parties. Honest parties’ protocols would then forward this message to $\mathcal{F}(\text{idealP})$ in the ideal-world execution, or to the $\mathcal{G}(C)$ functionality in the real-world execution. On collecting “tick” messages from all honest parties, the $\mathcal{F}(\text{idealP})$ or $\mathcal{G}(C)$ functionality would then advance the time $T := T + 1$. The functionality also allows parties to query the current time T .

As mentioned earlier, when multiple messages arrive at the blockchain in a time interval, we allow the adversary to choose a permutation to specify the order in which the blockchain will process the messages. This captures potential network attacks such as delaying message propagation, and front-running attacks (a.k.a. rushing attacks) where an adversary determines its own message after seeing what other parties send in a round.

2.3 Modeling Pseudonyms

We model a notion of “pseudonymity” that provides a form

$\mathcal{F}(\text{idealP})$ functionality

Given an ideal program denoted idealP , the $\mathcal{F}(\text{idealP})$ functionality is defined as below:

Init: Upon initialization, perform the following:

Time. Set current time $T := 0$. Set the receive queue $\text{rqueue} := \emptyset$.

Pseudonyms. Set $\text{nyms} := \{(P_1, P_1), \dots, (P_N, P_N)\}$, i.e., initially every party's true identity is recorded as a default pseudonym for the party.

Ledger. A ledger dictionary structure $\text{ledger}[P]$ stores the endowed account balance for each identity $P \in \{P_1, \dots, P_N\}$. Before any new pseudonyms are generated, only true identities have endowed account balances. Send the array $\text{ledger}[]$ to the ideal adversary \mathcal{S} .

idealP.Init. Run the **Init** procedure of the idealP program.

Tick: Upon receiving tick from an honest party P : notify \mathcal{S} of (tick, P) . If the functionality has collected tick confirmations from all honest parties since the last clock tick, then

Call the **Timer** procedure of the idealP program.

Apply the adversarial permutation perm to the rqueue to reorder the messages received in the previous round.

For each $(m, \bar{P}) \in \text{rqueue}$ in the permuted order, invoke the **delayed actions (in gray background)** defined by ideal program idealP at the activation point named “*Upon receiving message m from pseudonym \bar{P}* ”. Notice that the program idealP speaks of pseudonyms instead of party identifiers. Set $\text{rqueue} := \emptyset$.

Set $T := T + 1$

Other activations: Upon receiving a message of the form (m, \bar{P}) from a party P :

Assert that $(\bar{P}, P) \in \text{nyms}$.

Invoke the immediate actions defined by ideal program idealP at the activation point named “*Upon receiving message m from pseudonym \bar{P}* ”.

Queue the message by calling $\text{rqueue.add}(m, \bar{P})$.

Permute: Upon receiving $(\text{permute}, \text{perm})$ from the adversary \mathcal{S} , record perm .

GetTime: On receiving gettime from a party P , notify the adversary \mathcal{S} of $(\text{gettime}, P)$, and return the current time T to party P .

GenNym: Upon receiving gennym from an honest party P : Notify the adversary \mathcal{S} of gennym . Wait for \mathcal{S} to respond with a new nym \bar{P} such that $\bar{P} \notin \text{nyms}$. Now, let $\text{nyms} := \text{nyms} \cup \{(P, \bar{P})\}$, and send \bar{P} to P . Upon receiving (gennym, \bar{P}) from a corrupted party P : if $\bar{P} \notin \text{nyms}$, let $\bar{P} := \text{nyms} \cup \{(P, \bar{P})\}$.

Ledger operations: // inner activation

Transfer: Upon receiving $(\text{transfer}, \text{amount}, \bar{P}_r)$ from some pseudonym \bar{P}_s :

Notify $(\text{transfer}, \text{amount}, \bar{P}_r, \bar{P}_s)$ to the ideal adversary \mathcal{S} .

Assert that $\text{ledger}[\bar{P}_s] \geq \text{amount}$.

$\text{ledger}[\bar{P}_s] := \text{ledger}[\bar{P}_s] - \text{amount}$

$\text{ledger}[\bar{P}_r] := \text{ledger}[\bar{P}_r] + \text{amount}$

/ \bar{P}_s, \bar{P}_r can be pseudonyms or true identities. Note that each party's identity is a default pseudonym for the party. */*

Expose: On receiving exposeledger from a party P , return ledger to the party P .

Figure 3: The $\mathcal{F}(\text{idealP})$ functionality is parameterized by an ideal program denoted idealP . An ideal program idealP can specify two types of activation points, *immediate activations* and *delayed activations*. Activation points are invoked upon recipient of messages. Immediate activations are processed immediately, whereas delayed activations are collected and batch processed in the next round. The $\mathcal{F}(\cdot)$ wrapper allows the ideal adversary \mathcal{S} to specify an order perm in which the messages should be processed in the next round. For each delayed activation, we use the **leak** notation in an ideal program idealP to define the leakage which is *immediately* exposed to the ideal adversary \mathcal{S} upon recipient of the message.

$\mathcal{G}(\mathbf{C})$ functionality

Given a contract program denoted \mathbf{C} , the $\mathcal{G}(\mathbf{C})$ functionality is defined as below:

Init: Upon initialization, perform the following:

- A ledger data structure $\text{ledger}[\bar{P}]$ stores the account balance of party \bar{P} . Send the entire balance ledger to \mathcal{A} .
- Set current time $T := 0$. Set the receive queue $\text{rqueue} := \emptyset$.
- Run the **Init** procedure of the \mathbf{C} program.
- Send the \mathbf{C} program's internal state to the adversary \mathcal{A} .

Tick: Upon receiving **tick** from an honest party, if the functionality has collected **tick** confirmations from all honest parties since the last clock tick, then

- Apply the adversarial permutation **perm** to the **rqueue** to reorder the messages received in the previous round.
- Call the **Timer** procedure of the \mathbf{C} program.
- Pass the reordered messages to the \mathbf{C} program to be processed. Set $\text{rqueue} := \emptyset$.
- Set $T := T + 1$

Other activations:

- *Authenticated receive:* Upon receiving a message (**authenticated**, m) from party P :
 - Send (m, P) to the adversary \mathcal{A}
 - Queue the message by calling $\text{rqueue.add}(m, P)$.
- *Pseudonymous receive:* Upon receiving a message of the form (**pseudonymous**, m, \bar{P}, σ) from any party:
 - Send (m, \bar{P}, σ) to the adversary \mathcal{A}
 - Parse $\sigma := (\text{nonce}, \sigma')$, and assert $\text{Verify}(\bar{P}.\text{spk}, (\text{nonce}, T, \bar{P}.\text{epk}, m), \sigma') = 1$
 - If message (**pseudonymous**, m, \bar{P}, σ) has not been received earlier in this round, queue the message by calling $\text{rqueue.add}(m, \bar{P})$.
- *Anonymous receive:* Upon receiving a message (**anonymous**, m) from party P :
 - Send m to the adversary \mathcal{A}
 - If m has not been seen before in this round, queue the message by calling $\text{rqueue.add}(m)$.

Permute: Upon receiving (**permute**, **perm**) from the adversary \mathcal{A} , record **perm**.

Expose: On receiving **exposestate** from a party P , return the functionality's internal state to the party P . Note that this also implies that a party can query the functionality for the current time T .

Ledger operations: // inner activation

Transfer: Upon recipient of (**transfer**, **amount**, \bar{P}_r) from some pseudonym \bar{P}_s :

- Assert $\text{ledger}[\bar{P}_s] \geq \text{amount}$
- $\text{ledger}[\bar{P}_s] := \text{ledger}[\bar{P}_s] - \text{amount}$
- $\text{ledger}[\bar{P}_r] := \text{ledger}[\bar{P}_r] + \text{amount}$

Figure 4: The $\mathcal{G}(\mathbf{C})$ functionality is parameterized by a contract program denoted \mathbf{C} . The $\mathcal{G}(\cdot)$ wrapper mainly performs the following: *i*) exposes all of its internal states and messages received to the adversary; *ii*) makes the functionality time-aware: messages received in one round and queued and processed in the next round. The $\mathcal{G}(\cdot)$ wrapper allows the adversary to specify an ordering to the messages received by the contract in one round.

$\Pi(\text{prot})$ protocol wrapper in the $\mathcal{G}(\mathcal{C})$ -hybrid world

Given a party's local program denoted prot , the $\Pi(\text{prot})$ functionality is defined as below:

Pseudonym related:

GenNym: Upon receiving input gennym from the environment \mathcal{E} , generate $(\text{epk}, \text{esk}) \leftarrow \text{Keygen}_{\text{enc}}(1^\lambda)$, and $(\text{spk}, \text{ssk}) \leftarrow \text{Keygen}_{\text{sign}}(1^\lambda)$. Call $\text{payload} := \text{prot.GenNym}(1^\lambda, (\text{epk}, \text{spk}))$. Store $\text{nym} := \text{nym} \cup \{(\text{epk}, \text{spk}, \text{payload})\}$, and output $(\text{epk}, \text{spk}, \text{payload})$ as a new pseudonym.

Send: Upon receiving internal call $(\text{send}, m, \bar{P})$:

If $\bar{P} == P$: send $(\text{authenticated}, m)$ to $\mathcal{G}(\mathcal{C})$. // *this is an authenticated send*

Else, // *this is a pseudonymous send*

Assert that pseudonym \bar{P} has been recorded in nym ;

Query current time T from $\mathcal{G}(\mathcal{C})$. Compute $\sigma' := \text{Sign}(\text{ssk}, (\text{nonce}, T, \text{epk}, m))$ where ssk is the recorded secret signing key corresponding to \bar{P} , nonce is a freshly generated random string, and epk is the recorded public encryption key corresponding to \bar{P} . Let $\sigma := (\text{nonce}, \sigma')$.

Send $(\text{pseudonymous}, m, \bar{P}, \sigma)$ to $\mathcal{G}(\mathcal{C})$.

AnonSend: Upon receiving internal call $(\text{anonsend}, m, \bar{P})$: send $(\text{anonymous}, m)$ to $\mathcal{G}(\mathcal{C})$.

Timer and ledger transfers:

Transfer: Upon receiving input $(\text{transfer}, \$\text{amount}, \bar{P}_r, \bar{P})$ from the environment \mathcal{E} :

Assert that \bar{P} is a previously generated pseudonym.

Send $(\text{transfer}, \$\text{amount}, \bar{P}_r)$ to $\mathcal{G}(\mathcal{C})$ as pseudonym \bar{P} .

Tick: Upon receiving tick from the environment \mathcal{E} , forward the message to $\mathcal{G}(\mathcal{C})$.

Other activations:

Act as pseudonym: Upon receiving any input of the form (m, \bar{P}) from the environment \mathcal{E} :

Assert that \bar{P} was a previously generated pseudonym.

Pass (m, \bar{P}) the party's local program to process.

Others: Upon receiving any other input from the environment \mathcal{E} , or any other message from a party: Pass the input/message to the party's local program to process.

Figure 5: Protocol wrapper.

of privacy, similar to that provided by typical cryptocurrencies such as Bitcoin. Any user can generate an arbitrary (polynomially-bounded) number of pseudonyms, and each pseudonym is “owned” by the party who generated it. The correspondence of pseudonyms to real identities is hidden from the adversary.

Effectively, a pseudonym is a public key for a digital signature scheme, the corresponding private key to which is known by the party who “owns” the pseudonym. The public contract functionality allows parties to publish authenticated messages that are bound to a pseudonym of their choice. Thus each interaction with the public contract is, in general, associated with a pseudonym but not to a user’s real identity.

We abstract away the details of pseudonym management by implementing them in our wrappers. This allows user-defined applications to be written very simply, as though using ordinary identities, while enjoying the privacy benefits of pseudonymity.

Our wrapper provides a user-defined hook, “gennym”, that is invoked each time a party creates a pseudonym. This allows the application to define an additional per-pseudonym payload, such as application-specific public keys. From the point-of-view of the application, this is simply an initialization subroutine invoked once for each participant.

Our wrapper provides several means for users to communicate with a contract. The most common way is for a user to publish an authenticated message associated with one of their pseudonyms, as described above. Additionally, “anon-send” allows a user to publish a message without reference to any pseudonym at all.

In spite of pseudonymity, it is sometimes desirable to assign a particular user to a specific role in a contract (e.g., “auction manager”). The alternative is to assign roles on a “first-come first-served” basis (e.g., as the bidders in an auction). To this end, we allow each party to define generate a single “default” pseudonym which is publicly-bound to their real identity. We allow applications to make use of this through a convenient abuse of notation, by simply using a party identifier as a parameter or hardcoded string. Strictly speaking, the pseudonym string is not determined until the “gennym” subroutine is executed; the formal interpretation is that whenever such an identity is used, the default pseudonym associated with the identity is fetched from the contract. (This approach is effectively the same as taken by Canetti [20], where a functionality \mathcal{F}_{CA} allows each party to bind their real identity to a single public key of their choice).

2.4 Modeling Money

We model money as a public ledger, which associates quantities of money to pseudonyms. Users can transfer funds to each other (or among their own pseudonyms) by sending “transfer” messages to the public contract (as with other messages, these are delayed until the next round and may be delivered in any order). The ledger state is public knowledge, and can be queried immediately using the “exposeledger” instruction.

There are many conceivable policies for introducing new currency into such a system: for example, Bitcoin “mints” new currency as a reward for each miner who solves a proof-of-work puzzles. We take a simple approach of defining an arbitrary, publicly visible (i.e., common knowledge) ini-

tial allocation that associates a quantity of money to each party’s real identity. Except for this initial allocation, no money is created or destroyed.

2.5 Conventions for Writing Programs

Thanks to our wrapper-based modularized notational system, The ideal program and the contract program are the main locations where user-supplied, custom program logic is defined. We use the following conventions for writing the ideal program and the contract program.

Timer activation points. Every time $\mathcal{F}(\text{idealP})$ or $\mathcal{G}(C)$ advances the timer, it will invoke a **Timer** interrupt call. Therefore, by convention, we allow the ideal program or the contract program can define a **Timer** activation point. Timeout operations (e.g., refunding money after a certain timeout) can be implemented under the **Timer** activation point.

Delayed processing in ideal programs. When writing the contract program, every message received by the contract program is already delayed by a round due to the $\mathcal{G}(\cdot)$ wrapper.

When writing the ideal program, we introduce a simple convention to denote delayed computation. Program instructions that are written in gray background denote computation that does not take place immediately, but is deferred to the beginning of the next timer click. This is a convenient shorthand because in our real-world protocol, effectively any computation done by a contract functionality will be delayed. For example, in our $\text{Ideal}_{\text{cash}}$ ideal program (see Figure 7), whenever the ideal functionality receives a `mint` or `pour` message, the ideal adversary \mathcal{S} is notified immediately; however, processing of the messages is deferred till the next timer click. Formally, delayed processing can be implemented simply by storing state and invoking the delayed program instructions on the next **Timer** click. To avoid ambiguity, we assume that by convention, the delayed instructions are invoked at the beginning of the **Timer** call. In other words, upon the next timer click, the delayed instructions are executed first.

Pseudonymity. All party identifiers that appear in ideal programs, contract programs, and user-side programs by default refer to *pseudonyms*. When we write “upon receiving message from *some P*”, this accepts a message from any pseudonym. Whenever we write “upon receiving message from *P*”, without the keyword *some*, this accepts a message from a fixed pseudonym *P*, and typically which pseudonym we refer to is clear from the context.

Whenever we write “send *m* to $\mathcal{G}(\text{Contract})$ as nym *P*” inside a user program, this sends an internal message (“send”, *m*, *P*) to the protocol wrapper Π . The protocol wrapper will then authenticate the message appropriately under pseudonym *P*. When the context is clear, we avoid writing “as nym *P*”, and simply write “send *m* to $\mathcal{G}(\text{Contract})$ ”. Our formal system also allows users to send messages anonymously to a contract – although this option will not be used in this paper.

Ledger and money transfers. A public dedger is denoted **ledger** in our ideal programs and contract programs. When a party sends `$amt` to an ideal program or a contract program, this represents an ordinary message transmission. Money transfers only take place when ideal programs or contract

programs update the public ledger `ledger`. In other words, the symbol $\$$ is only adopted for readability (to distinguish variables associated with money and other variables), and does not have special meaning or significance. One can simply think of this variable as having the money type.

2.6 Composability and Multiple Contracts

Extending to multiple contracts. So far, our formalism only models a single running instance of a user-specified contract $(\phi_{\text{priv}}, \phi_{\text{pub}})$. It will not be too hard to extend the wrappers to support multiple contracts sharing a global ledger, clock, pseudonyms, and `ContractCash` (i.e, private cash). While such an extension is straightforward, one obvious drawback is that it would result in a monolithic functionality consisting all contract instances. This means that the proof also has to be done in a monolithic manner simultaneously proving all active contracts in the system.

Future work. To further modularize our functionality and proof, new composition theorems will be needed that are not covered by the current UC [19] or extended models such as GUC [21] and GNUC [34]. We give a brief discussion of the issues below. Since our model is expressed in the Universal Composability framework, we could apply to our functionalities and protocols standard composition operators, such as the multi-session extension [22]. However, a direct application of this operator to the wrapped functionality $\mathcal{F}(\text{Ideal}_{\text{hawk}})$ would give us multiple instances of *separate* timers and ledgers, one for each contract - which is not what we want! The Generalized UC (GUC) framework [21] is a better starting point; it provides a way to compose multiple instances of arbitrary functionalities along with a single instance of a shared functionality as a common resource. To apply this to our scenario, we would model the timer and ledger as a single shared functionality, composed with an arbitrary number of instances of Hawk contracts. However, even the GUC framework is inadequate for our needs since it does not allow interaction *between* the shared functionality and others, so this approach cannot be applied directly. In our ongoing work, we further generalize GUC and overcome these technical obstacles and more. As these details are intricate and unrelated to our contributions here, we defer further discussion to a future manuscript.

3. CRYPTOGRAPHY ABSTRACTIONS

We now describe our cryptography abstraction in the form of ideal programs. Ideal programs define the correctness and security requirements we wish to attain by writing a specification assuming the existence of a fully trusted party. We will later prove that our real-world protocols (based on smart contracts) securely emulate the ideal programs. As mentioned earlier, an ideal program must be combined with a wrapper \mathcal{F} to be endowed with exact execution semantics.

Overview. Hawk realizes the following specifications:

- *Private ledger and currency transfer.* Hawk relies on the existence of a private ledger that supports private currency transfers. We therefore first define an ideal functionality called $\text{Ideal}_{\text{cash}}$ that describes the requirements of a private ledger (see Figure 6). Informally speaking, earlier works such as Zerocash [10] are meant to realize (approximations of) this ideal functionality – although technically this ought to be interpreted with the caveat

$\text{Ideal}_{\text{cash}}$	
Init:	Coins: a multiset of coins, each of the form $(\mathcal{P}, \$\text{val})$
Mint:	Upon receiving $(\text{mint}, \$\text{val})$ from some \mathcal{P} : send $(\text{mint}, \mathcal{P}, \$\text{val})$ to \mathcal{A} <div style="background-color: #f0f0f0; padding: 2px;"> assert $\text{ledger}[\mathcal{P}] \geq \\val $\text{ledger}[\mathcal{P}] := \text{ledger}[\mathcal{P}] - \\val append $(\mathcal{P}, \\$\text{val})$ to Coins </div>
Pour:	On $(\text{pour}, \$\text{val}_1, \$\text{val}_2, \mathcal{P}_1, \mathcal{P}_2, \$\text{val}'_1, \$\text{val}'_2)$ from \mathcal{P} : assert $\$ \text{val}_1 + \$ \text{val}_2 = \$ \text{val}'_1 + \$ \text{val}'_2$ if \mathcal{P} is honest, assert $(\mathcal{P}, \$\text{val}_i) \in \text{Coins}$ for $i \in \{1, 2\}$ assert $\mathcal{P}_i \neq \perp$ for $i \in \{1, 2\}$ remove one $(\mathcal{P}, \$\text{val}_i)$ from Coins for $i \in \{1, 2\}$ for $i \in \{1, 2\}$, if \mathcal{P}_i is corrupted, send $(\text{pour}, i, \mathcal{P}_i, \$\text{val}'_i)$ to \mathcal{A} ; else send $(\text{pour}, i, \mathcal{P}_i)$ to \mathcal{A} <div style="background-color: #f0f0f0; padding: 2px;"> if \mathcal{P} is corrupted: assert $(\mathcal{P}, \\$\text{val}_i) \in \text{Coins}$ for $i \in \{1, 2\}$ remove one $(\mathcal{P}, \\$\text{val}_i)$ from Coins for $i \in \{1, 2\}$ for $i \in \{1, 2\}$: add $(\mathcal{P}_i, \\$\text{val}'_i)$ to Coins for $i \in \{1, 2\}$: if $\mathcal{P}_i \neq \perp$, send $(\text{pour}, \\$\text{val}'_i)$ to \mathcal{P}_i </div>

Figure 6: Definition of $\text{Ideal}_{\text{cash}}$. Notation: `ledger` denotes the public ledger, and `Coins` denotes the private pool of coins. As mentioned in Section 2.5, gray background denotes batched and delayed activation. All party names correspond to pseudonyms due to notations and conventions defined in Section 2.

that these earlier works prove indistinguishability or game-based security instead UC-based simulation security.

- *Hawk-specific primitives.* With a private ledger specified, we then define Hawk-specific primitives including *freeze*, *compute*, and *finalize* that are essential for enabling transactional privacy and programmability simultaneously.

3.1 Private Cash Specification $\text{Ideal}_{\text{cash}}$

At a high-level, the $\text{Ideal}_{\text{cash}}$ specifies the requirements of a private ledger and currency transfer. We adopt the same “mint” and “pour” terminology from Zerocash [10].

Mint. The `mint` operation allows a user \mathcal{P} to transfer money from the public ledger denoted `ledger` to the private pool denoted `Coins`[\mathcal{P}]. With each transfer, a private coin for user \mathcal{P} is created, and associated with a value `val`.

For correctness, the ideal program $\text{Ideal}_{\text{cash}}$ checks that the user \mathcal{P} has sufficient funds in its public ledger `ledger`[\mathcal{P}] before creating the private coin.

Pour. The `pour` operation allows a user \mathcal{P} to spend money in its private bank privately. For simplicity, we define the simple case with two input coins and two output coins. This is sufficient for users to transfer any amount of money by “making change,” although it would be straightforward to support more efficient batch operations as well.

For correctness, the ideal program $\text{Ideal}_{\text{cash}}$ checks the following: 1) for the two input coins, party \mathcal{P} indeed possesses private coins of the declared values; and 2) the two input

coins sum up to equal value as the two output coins, i.e., coins neither get created or vanish.

Privacy. When an honest party \mathcal{P} mints, the ideal-world adversary \mathcal{A} learns the pair $(\mathcal{P}, \text{val})$ – since minting is raising coins from the public pool to the private pool. Operations on the public pool are observable by \mathcal{A} .

When an honest party \mathcal{P} pours, however, the adversary \mathcal{A} learns only the output pseudonyms \mathcal{P}_1 and \mathcal{P}_2 . It does not learn which coin in the private pool Coins is being spent nor the name of the spender. Therefore, the spent coins are anonymous with respect to the private pool Coins . To get strong anonymity, new pseudonyms \mathcal{P}_1 and \mathcal{P}_2 can be generated on the fly to receive each pour. We stress that as long as **pour** hides the sender, this “breaks” the transaction graph, thus preventing linking analysis.

If a corrupted party is the recipient of a pour, the adversary additionally learns the value of the coin it receives.

Additional subtleties. Later in our protocol, honest parties keep track of a wallet of coins. Whenever an honest party pours, it first checks if an appropriate coin exists in its local wallet – and if so it immediately removes the coin from the wallet (i.e., without delay). In this way, if an honest party makes multiple pour transactions in one round, it will always choose distinct coins for each pour transaction. Therefore, in our $\text{Ideal}_{\text{cash}}$ functionality, honest pourers’ coins are immediately removed from Coins . Further, an honest party is not able to spend a coin paid to itself until the next round. By contrast, corrupted parties are allowed to spend coins paid to them in the same round – this is due to the fact that any message is routed immediately to the adversary, and the adversary can also choose a permutation for all messages received by the contract in the same round (see Section 2).

Another subtlety in the $\text{Ideal}_{\text{cash}}$ functionality is that honest parties will always pour to existing pseudonyms. However, the functionality allows the adversary to pour to non-existing pseudonyms denoted \perp – in this case, effectively the private coin goes into a blackhole and cannot be retrieved. This enables a performance optimization in our ProtCash and ContractCash protocol later – where we avoid including the ct_i ’s in the NIZK of $\mathcal{L}_{\text{POUR}}$ (see Section 4.1). If a malicious pourer chooses to compute the wrong ct_i , it is as if the recipient \mathcal{P}_i did not receive the pour, i.e., the pour is made to \perp .

3.2 Hawk Specification $\text{Ideal}_{\text{hawk}}$

To enable transactional privacy and programmability simultaneously, we now describe the specifications of new **Hawk** primitives, including *freeze*, *compute*, and *finalize*. The formal specification of the ideal program $\text{Ideal}_{\text{hawk}}$ is provided in Figure 7. Below, we provide some explanations. We also refer the reader to Section 1.3 for higher-level explanations.

Freeze. In **freeze**, a party tells $\text{Ideal}_{\text{hawk}}$ to remove one coin from the private coins pool Coins , and freeze it in the contract by adding it to ContractCoins . The party’s private input denoted in is also recorded in ContractCoins . $\text{Ideal}_{\text{hawk}}$ checks that \mathcal{P} has not called **freeze** earlier, and that a coin $(\mathcal{P}, \text{val})$ exists in Coins before proceeding with the freeze.

Compute. When a party \mathcal{P} calls **compute**, its private input in and the value of its frozen coin val are disclosed to the manager $\mathcal{P}_{\mathcal{M}}$.

$\text{Ideal}_{\text{hawk}}(\mathcal{P}_{\mathcal{M}}, \{\mathcal{P}_i\}_{i \in [N]}, T_1, T_2, \phi_{\text{priv}}, \phi_{\text{pub}})$

Init: Call $\text{Ideal}_{\text{cash}}.\text{Init}$. Additionally:

ContractCoins: a set of coins and private inputs received by this contract, each of the form $(\mathcal{P}, \text{in}, \text{\$val})$. Initialize $\text{ContractCoins} := \emptyset$.

Freeze: Upon receiving $(\text{freeze}, \text{\$val}_i, \text{in}_i)$ from \mathcal{P}_i for some $i \in [N]$:

assert current time $T < T_1$

assert \mathcal{P}_i has not called **freeze** earlier.

assert at least one copy of $(\mathcal{P}_i, \text{\$val}_i) \in \text{Coins}$

send $(\text{freeze}, \mathcal{P}_i)$ to \mathcal{A}

add $(\mathcal{P}_i, \text{\$val}_i, \text{in}_i)$ to ContractCoins

remove one $(\mathcal{P}_i, \text{\$val}_i)$ from Coins

Compute: Upon receiving **compute** from \mathcal{P}_i for some $i \in [N]$:

assert current time $T_1 \leq T < T_2$

if $\mathcal{P}_{\mathcal{M}}$ is corrupted, send $(\text{compute}, \mathcal{P}_i, \text{\$val}_i, \text{in}_i)$ to \mathcal{A}

else send $(\text{compute}, \mathcal{P}_i)$ to \mathcal{A}

let $(\mathcal{P}_i, \text{\$val}_i, \text{in}_i)$ be the item in ContractCoins corresponding to \mathcal{P}_i

send $(\text{compute}, \mathcal{P}_i, \text{\$val}_i, \text{in}_i)$ to $\mathcal{P}_{\mathcal{M}}$

Finalize: Upon receiving $(\text{finalize}, \text{in}_{\mathcal{M}}, \text{out})$ from $\mathcal{P}_{\mathcal{M}}$:

assert current time $T \geq T_2$

assert $\mathcal{P}_{\mathcal{M}}$ has not called **finalize** earlier

for $i \in [N]$:

let $(\text{\$val}_i, \text{in}_i) := (0, \perp)$ if \mathcal{P}_i has not called **compute**

$(\{\text{\$val}'_i\}, \text{out}^\dagger) := \phi_{\text{priv}}(\{\text{\$val}_i, \text{in}_i\}, \text{in}_{\mathcal{M}})$

assert $\text{out}^\dagger = \text{out}$

assert $\sum_{i \in [N]} \text{\$val}_i = \sum_{i \in [N]} \text{\$val}'_i$

send $(\text{finalize}, \text{in}_{\mathcal{M}}, \text{out})$ to \mathcal{A}

for each corrupted \mathcal{P}_i that called **compute**: send $(\mathcal{P}_i, \text{\$val}'_i)$ to \mathcal{A}

call $\phi_{\text{pub}}.\text{check}(\text{in}_{\mathcal{M}}, \text{out})$

for $i \in [N]$ such that \mathcal{P}_i called **compute**:

add $(\mathcal{P}_i, \text{\$val}'_i)$ to Coins

send $(\text{finalize}, \text{\$val}'_i)$ to \mathcal{P}_i

ϕ_{pub} : Run a local instance of public contract ϕ_{pub} . Messages between the adversary to ϕ_{pub} , and from ϕ_{pub} to parties are forwarded directly.

Upon receiving message (pub, m) from party \mathcal{P} :

notify \mathcal{A} of (pub, m)

send m to ϕ_{pub} on behalf of \mathcal{P}

$\text{Ideal}_{\text{cash}}$: include $\text{Ideal}_{\text{cash}}$ (Figure 6).

Figure 7: Definition of $\text{Ideal}_{\text{hawk}}$. Notations: ContractCoins denotes frozen coins owned by the contract; Coins denotes the global private coin pool defined by $\text{Ideal}_{\text{cash}}$; and $(\text{in}_i, \text{val}_i)$ denotes the input data and frozen coin value of party \mathcal{P}_i .

Finalize. In `finalize`, the manager $\mathcal{P}_{\mathcal{M}}$ submits a public input `inM` to `Idealhawk`. `Idealhawk` now computes the outcome of ϕ_{priv} on all parties’ inputs and frozen coin values, and redistribute the `ContractCoins` based on the outcome of ϕ_{priv} . To ensure money conservation, the ideal program `Idealhawk` checks that the sum of frozen coins is equal to the sum of output coins.

Interaction with public contract. The `Idealhawk` functionality is parameterized by an ordinary public contract ϕ_{pub} , which is included in `Idealhawk` as a sub-module. During a `finalize`, `Idealhawk` calls $\phi_{\text{pub}}.\text{check}$. The public contract ϕ_{pub} typically serves the following purposes:

- *Check the well-formedness of the manager’s input in_M.* For example, in our financial derivatives application (Section 5.2), the public contract ϕ_{pub} asserts that the input corresponds to the price of a stock as reported by the stock exchange’s authentic data feed.
- *Redistribute public deposits.* If parties or the manager have aborted, or if a party has provided invalid input (e.g., less than a minimum bet) the public contract ϕ_{pub} can now redistribute the parties’ public deposits to ensure financial fairness. For example, in our “Rock, Paper, Scissors” example (see Section 5.2), the private contract ϕ_{priv} checks if each party has frozen the minimal bet. If not, ϕ_{priv} includes that information in `out` so that ϕ_{pub} pays that party’s public deposit to others.

Security and privacy requirements. The `Idealhawk` specifies the following privacy guarantees. When an honest party \mathcal{P} freezes money (e.g., a bid), the adversary should not observe the amount frozen. However, the adversary can observe the party’s pseudonym \mathcal{P} . We note that leaking the pseudonym \mathcal{P} does not hurt privacy, since a party can simply create a new pseudonym \mathcal{P} and pour to this new pseudonym immediately before the freeze.

When an honest party calls `compute`, the manager $\mathcal{P}_{\mathcal{M}}$ gets to observe its input and frozen coin’s value. However, the public and other contractual parties do not observe anything (unless the manager voluntarily discloses information).

Finally, during a `finalize` operation, the output `out` is declassified to the public – note that `out` can be empty if we do not wish to declassify any information to the public.

It is not hard to see that our ideal program `Idealhawk` satisfies *input independent privacy* and *authenticity* against a dishonest manager. Further, it satisfies *posterior privacy* as long as the manager does not voluntarily disclose information. Intuitive explanations of these security/privacy properties were provided in Section 1.2.

Timing and aborts. Our ideal program `Idealhawk` requires that `freeze` operations be done by time T_1 , and that `compute` operations be done by time T_2 . If a user froze coins but did not open by time T_2 , our ideal program `Idealhawk` treats $(\text{in}_i, \text{val}_i) := (0, \perp)$, and the user \mathcal{P}_i essentially forfeits its frozen coins. Managerial aborts is not handled inside `Idealhawk`, but by the public portion of the contract.

Simplifying assumptions. For clarity, our basic version of `Idealhawk` is a stripped down version of our implementation. Specifically, our basic `Idealhawk` and protocols do not realize refunds of frozen coins upon managerial abort. As mentioned in Section 4.3, it is not hard to extend our protocols to support such refunds.

Other simplifying assumptions we made include the following. Our basic `Idealhawk` assumes that the set of pseudonyms participating in the contract as well as timeouts T_1 and T_2 are hard-coded in the program. This can also be easily relaxed as mentioned in Section 4.3.

4. CRYPTOGRAPHIC PROTOCOLS

4.1 Protocol Description

Our protocols are broken down into two parts: 1) the private cash part that implements direct money transfers between users; and 2) the `Hawk`-specific part that binds transactional privacy with programmable logic. The formal protocol descriptions are given in Figures 8 and 9. Below we explain the high-level intuition.

4.1.1 Warmup: Private Cash and Money Transfers

Our construction adopts a Zerocash-like protocol for implementing private cash and private currency transfers. For completeness, we give a brief explanation below, and we mainly focus on the `pour` operation which is technically more interesting. The contract `ContractCash` maintains a set `Coins` of private coins. Each private coin is stored in the format

$$(\mathcal{P}, \text{coin} := \text{Comm}_s(\$val))$$

where \mathcal{P} denotes a party’s pseudonym, and `coin` commits to the coin’s value `$val` under randomness s .

During a `pour` operation, the spender \mathcal{P} chooses two coins in `Coins` to spend, denoted $(\mathcal{P}, \text{coin}_1)$ and $(\mathcal{P}, \text{coin}_2)$ where $\text{coin}_i := \text{Comm}_{s_i}(\$val_i)$ for $i \in \{1, 2\}$. The `pour` operation pays `val’1` and `val’2` amount to two output pseudonyms denoted \mathcal{P}_1 and \mathcal{P}_2 respectively, such that `val1` + `val2` = `val’1` + `val’2`. The spender chooses new randomness s'_i for $i \in \{1, 2\}$, and computes the output coins as

$$(\mathcal{P}_i, \text{coin}_i := \text{Comm}_{s'_i}(\$val'_i))$$

The spender gives the values s'_i and `val’i` to the recipient \mathcal{P}_i for \mathcal{P}_i to be able to spend the coins later.

Now, the spender computes a zero-knowledge proof to show that the output coins are constructed appropriately, where correctness compasses the following aspects:

- *Existence of coins being spent.* The coins being spent $(\mathcal{P}, \text{coin}_1)$ and $(\mathcal{P}, \text{coin}_2)$ are indeed part of the private pool `Coins`. We remark that here the zero-knowledge property allows the spender to hide which coins it is spending – this is the key idea behind transactional privacy. To prove this efficiently, `ContractCash` maintains a Merkle tree `MT` over the private pool `Coins`. Membership in the set can be demonstrated by a Merkle branch consistent with the root hash, and this is done in zero-knowledge.
- *No double spending.* Each coin $(\mathcal{P}, \text{coin})$ has a cryptographically unique serial number `sn` that can be computed as a pseudorandom function of \mathcal{P} ’s secret key and `coin`. To pour a coin, its serial number `sn` must be disclosed, and a zero-knowledge proof given to show the correctness of `sn`. `ContractCash` checks that no `sn` is used twice.
- *Money conservation.* The zero-knowledge proof also attests to the fact that the input coins and the output coins have equal total value.

We make some remarks about the security of the scheme. Intuitively, when an honest party pours to an honest party,

ContractCash	
Init:	<p>crs: a reference string for the underlying NIZK system Coins: a set of coin commitments, initially \emptyset SpentCoins: set of spent serial numbers, initially \emptyset</p>
Mint:	<p>Upon receiving (mint, \$val, s) from some party \mathcal{P}, coin := Comm_s(\$val) assert ($\mathcal{P}$, coin) \notin Coins assert ledger[\mathcal{P}] \geq \$val ledger[$\mathcal{P}$] := ledger[$\mathcal{P}$] - \$val add (\mathcal{P}, coin) to Coins</p>
Pour:	<p>Anonymous receive (pour, π, {sn_i, \mathcal{P}_i, coin_i, ct_i}_{i∈{1,2}}}) let MT be a merkle tree built over Coins statement := (MT.root, {sn_i, \mathcal{P}_i, coin_i}_{i∈{1,2}}) assert NIZK.Verify($\mathcal{L}_{\text{POUR}}$, π, statement) for $i \in \{1, 2\}$, assert sn_i \notin SpentCoins assert (\mathcal{P}_i, coin_i) \notin Coins add sn_i to SpentCoins add (\mathcal{P}, coin_i) to Coins send (pour, coin_i, ct_i) to \mathcal{P}_i,</p>
<p>Relation (statement, witness) $\in \mathcal{L}_{\text{POUR}}$ is defined as: parse statement as (MT.root, {sn_i, \mathcal{P}_i, coin'_i}_{i∈{1,2}}) parse witness as (\mathcal{P}, sk_{prf}, {branch_i, s_i, \$val_i, s'_i, r_i, \$val'_i}_i) assert $\mathcal{P}.\text{pk}_{\text{prf}} = \text{PRF}_{\text{sk}_{\text{prf}}}(0)$ assert \$val₁ + \$val₂ = \$val'₁ + \$val'₂ for $i \in \{1, 2\}$, coin_i := Comm_{s_i}(\$val_i) assert MerkleBranch(MT.root, branch_i, (\mathcal{P} coin_i)) assert sn_i = PRF_{sk_{prf}}(\mathcal{P} coin_i) assert coin'_i = Comm_{s'_i}(\$val'_i)</p>	

Protocol ProtCash	
Init:	<p>Wallet: stores \mathcal{P}'s spendable coins, initially \emptyset</p>
GenNym:	<p>sample a random seed sk_{prf} pk_{prf} := PRF_{sk_{prf}}(0) return pk_{prf}</p>
Mint:	<p>On input (mint, \$val), sample a commitment randomness s coin := Comm_s(\$val) store (s, \$val, coin) in Wallet send (mint, \$val, s) to $\mathcal{G}(\text{ContractCash})$</p>
Pour (as sender):	<p>On input (pour, \$val₁, \$val₂, \mathcal{P}_1, \mathcal{P}_2, \$val'₁, \$val'₂), assert \$val₁ + \$val₂ = \$val'₁ + \$val'₂ for $i \in \{1, 2\}$, assert (s_i, \$val_i, coin_i) \in Wallet for some (s_i, coin_i) let MT be a merkle tree over ContractCash.Coins for $i \in \{1, 2\}$: remove one (s_i, \$val_i, coin_i) from Wallet sn_i := PRF_{sk_{prf}}(\mathcal{P} coin_i) let branch_i be the branch of (\mathcal{P}, coin_i) in MT sample randomness s'_i, r_i coin'_i := Comm_{s'_i}(\$val'_i) ct_i := ENC($\mathcal{P}_i$.epk, r_i, \$val'_i s'_i) statement := (MT.root, {sn_i, \mathcal{P}_i, coin'_i}_{i∈{1,2}}) witness := (\mathcal{P}, sk_{prf}, {branch_i, s_i, \$val_i, s'_i, r_i, \$val'_i}_i) π := NIZK.Prove($\mathcal{L}_{\text{POUR}}$, statement, witness) AnonSend(pour, π, {sn_i, \mathcal{P}_i, coin'_i, ct_i}_{i∈{1,2}}) to $\mathcal{G}(\text{ContractCash})$</p>
Pour (as recipient):	<p>On receive (pour, coin, ct) from $\mathcal{G}(\text{ContractCash})$: let (\$val s) := DEC(sk_{enc}, ct) assert Comm_s(\$val) = coin store (s, \$val, coin) in Wallet output (pour, \$val)</p>

Figure 8: ProtCash construction. A trusted setup phase generates the NIZK's common reference string crs. For notational convenience, we omit writing the crs explicitly in the construction. The Merkle tree MT is stored in the contract and not computed on the fly – we omit stating this in the protocol for notational simplicity. The protocol wrapper $\mathcal{H}(\cdot)$ invokes GenNym whenever a party creates a new pseudonym.

<p>ContractHawk($\mathcal{P}_M, \{\mathcal{P}_i\}_{i \in [N]}, T_1, T_2, \phi_{\text{priv}}, \phi_{\text{pub}}$)</p> <p>Init: See $\text{Ideal}_{\text{hawk}}$ for description of parameters Call ContractCash.Init.</p> <p>Freeze: Upon receiving (freeze, $\pi, \text{sn}_i, \text{cm}_i$) from \mathcal{P}_i: assert current time $T \leq T_1$ assert this is the first freeze from \mathcal{P}_i let MT be a merkle tree built over Coins assert $\text{sn}_i \notin \text{SpentCoins}$ statement := (MT.root, sn_i, cm_i) assert $\text{NIZK.Verify}(\mathcal{L}_{\text{FREEZE}}, \pi, \text{statement})$ add sn_i to SpentCoins and store cm_i for later</p> <p>Compute: Upon receiving (compute, π, ct) from \mathcal{P}_i: assert $T_1 \leq T < T_2$ for current time T assert $\text{NIZK.Verify}(\mathcal{L}_{\text{COMPUTE}}, \pi, (\mathcal{P}_M, \text{cm}_i, \text{ct}))$ send (compute, \mathcal{P}_i, ct) to \mathcal{P}_M</p> <p>Finalize: On receiving (finalize, $\pi, \text{in}_M, \text{out}, \{\text{coin}'_i, \text{ct}_i\}_{i \in [N]}$) from \mathcal{P}_M: assert current time $T \geq T_2$ for every \mathcal{P}_i that has not called compute, set $\text{cm}_i := \perp$ statement := ($\text{in}_M, \text{out}, \{\text{cm}_i, \text{coin}'_i, \text{ct}_i\}_{i \in [N]}$) assert $\text{NIZK.Verify}(\mathcal{L}_{\text{FINALIZE}}, \pi, \text{statement})$ for $i \in [N]$: assert $\text{coin}'_i \notin \text{Coins}$ add coin'_i to Coins send (finalize, $\text{coin}'_i, \text{ct}_i$) to \mathcal{P}_i Call $\phi_{\text{pub}}.\text{check}(\text{in}_M, \text{out})$</p> <hr/> <p>ContractCash: include ContractCash ϕ_{pub} : include user-defined public contract ϕ_{pub}</p> <hr/> <p>Relation (statement, witness) $\in \mathcal{L}_{\text{FREEZE}}$ is defined as: parse statement as (MT.root, sn, cm) parse witness as (\mathcal{P}, coin, sk_{prf}, branch, $s, \\$\text{val}, \text{in}, k, s'$) coin := $\text{Comm}_s(\\$ \text{val})$ assert $\text{MerkleBranch}(\text{MT.root}, \text{branch}, (\mathcal{P} \parallel \text{coin}))$ assert $\mathcal{P}.\text{pk}_{\text{prf}} = \text{sk}_{\text{prf}}(0)$ assert $\text{sn} = \text{PRF}_{\text{sk}_{\text{prf}}}(\mathcal{P} \parallel \text{coin})$ assert $\text{cm} = \text{Comm}_{s'}(\\$ \text{val} \parallel \text{in} \parallel k)$</p> <hr/> <p>Relation (statement, witness) $\in \mathcal{L}_{\text{COMPUTE}}$ is defined as: parse statement as ($\mathcal{P}_M, \text{cm}, \text{ct}$) parse witness as ($\\$ \text{val}, \text{in}, k, s', r$) assert $\text{cm} = \text{Comm}_{s'}(\\$ \text{val} \parallel \text{in} \parallel k)$ assert $\text{ct} = \text{ENC}(\mathcal{P}_M.\text{epk}, r, (\\$ \text{val} \parallel \text{in} \parallel k \parallel s'))$</p> <hr/> <p>Relation (statement, witness) $\in \mathcal{L}_{\text{FINALIZE}}$ is defined as: parse statement as ($\text{in}_M, \text{out}, \{\text{cm}_i, \text{coin}'_i, \text{ct}_i\}_{i \in [N]}$) parse witness as $\{s_i, \\$ \text{val}_i, \text{in}_i, s'_i, k_i\}_{i \in [N]}$ $(\{\\$ \text{val}'_i\}_{i \in [N]}, \text{out}) := \phi_{\text{priv}}(\{\\$ \text{val}_i, \text{in}_i\}_{i \in [N]}, \text{in}_M)$ assert $\sum_{i \in [N]} \\$ \text{val}_i = \sum_{i \in [N]} \\$ \text{val}'_i$ for $i \in [N]$: assert $\text{cm}_i = \text{Comm}_{s_i}(\\$ \text{val}_i \parallel \text{in}_i \parallel k_i)$ $\forall (\\$ \text{val}_i, \text{in}_i, k_i, s_i, \text{cm}_i) = (0, \perp, \perp, \perp, \perp)$ assert $\text{ct}_i = \text{SENC}_{k_i}(s'_i \parallel \\$ \text{val}'_i)$ assert $\text{coin}'_i = \text{Comm}_{s'_i}(\\$ \text{val}'_i)$</p>	<p>Protocol ProtHawk($\mathcal{P}_M, \{\mathcal{P}_i\}_{i \in [N]}, T_1, T_2, \phi_{\text{priv}}, \phi_{\text{pub}}$)</p> <p>Init: Call ProtCash.Init.</p> <hr/> <p>Protocol for a party $\mathcal{P} \in \{\mathcal{P}_i\}_{i \in [N]}$:</p> <p>Freeze: On input (freeze, $\\$ \text{val}, \text{in}$) as party \mathcal{P}: assert current time $T < T_1$ assert this is the first freeze input let MT be a merkle tree over Contract.Coins assert that some entry $(s, \\$ \text{val}, \text{coin}) \in \text{Wallet}$ for some (s, coin) remove one $(s, \\$ \text{val}, \text{coin})$ from Wallet $\text{sn} := \text{PRF}_{\text{sk}_{\text{prf}}}(\mathcal{P} \parallel \text{coin})$ let branch be the branch of $(\mathcal{P}, \text{coin})$ in MT sample a symmetric encryption key k sample a commitment randomness s' cm := $\text{Comm}_{s'}(\\$ \text{val} \parallel \text{in} \parallel k)$ statement := (MT.root, sn, cm) witness := (\mathcal{P}, coin, sk_{prf}, branch, $s, \\$ \text{val}, \text{in}, k, s'$) $\pi := \text{NIZK.Prove}(\mathcal{L}_{\text{FREEZE}}, \text{statement}, \text{witness})$ send (freeze, $\pi, \text{sn}, \text{cm}$) to $\mathcal{G}(\text{ContractHawk})$ store in, cm, $\\$ \text{val}, s'$, and k to use later (in compute)</p> <p>Compute: On input (compute) as party \mathcal{P}: assert current time $T_1 \leq T < T_2$ sample encryption randomness r ct := $\text{ENC}(\mathcal{P}_M.\text{epk}, r, (\\$ \text{val} \parallel \text{in} \parallel k \parallel s'))$ $\pi := \text{NIZK.Prove}((\mathcal{P}_M, \text{cm}, \text{ct}), (\\$ \text{val}, \text{in}, k, s', r))$ send (compute, π, ct) to $\mathcal{G}(\text{ContractHawk})$</p> <p>Finalize: Receive (finalize, coin, ct) from $\mathcal{G}(\text{ContractHawk})$: decrypt $(s \parallel \\$ \text{val}) := \text{SDEC}_k(\text{ct})$ store $(s, \\$ \text{val}, \text{coin})$ in Wallet output (finalize, $\\$ \text{val}$)</p> <hr/> <p>Protocol for manager \mathcal{P}_M:</p> <p>Compute: On receive (compute, \mathcal{P}_i, ct) from $\mathcal{G}(\text{ContractHawk})$: decrypt and store $(\\$ \text{val}_i \parallel \text{in}_i \parallel k_i \parallel s_i) := \text{DEC}(\text{epk}, \text{ct})$ store $\text{cm}_i := \text{Comm}_{s_i}(\\$ \text{val}_i \parallel \text{in}_i \parallel k_i)$ output $(\mathcal{P}_i, \\$ \text{val}_i, \text{in}_i)$ If this is the last compute received: for $i \in [N]$ such that \mathcal{P}_i has not called compute, $(\\$ \text{val}_i, \text{in}_i, k_i, s_i, \text{cm}_i) := (0, \perp, \perp, \perp, \perp)$ $(\{\\$ \text{val}'_i\}_{i \in [N]}, \text{out}) := \phi_{\text{priv}}(\{\\$ \text{val}_i, \text{in}_i\}_{i \in [N]}, \text{in}_M)$ store and output $(\{\\$ \text{val}'_i\}_{i \in [N]}, \text{out})$</p> <p>Finalize: On input (finalize, in_M, out): assert current time $T \geq T_2$ for $i \in [N]$: sample a commitment randomness s'_i $\text{coin}'_i := \text{Comm}_{s'_i}(\\$ \text{val}'_i)$ $\text{ct}_i := \text{SENC}_{k_i}(s'_i \parallel \\$ \text{val}'_i)$ statement := ($\text{in}_M, \text{out}, \{\text{cm}_i, \text{coin}'_i, \text{ct}_i\}_{i \in [N]}$) witness := $\{s_i, \\$ \text{val}_i, \text{in}_i, s'_i, k_i\}_{i \in [N]}$ $\pi := \text{NIZK.Prove}(\text{statement}, \text{witness})$ send (finalize, $\pi, \text{in}_M, \text{out}, \{\text{coin}'_i, \text{ct}_i\}_{i \in [N]}$) to $\mathcal{G}(\text{ContractHawk})$</p> <hr/> <p>ProtCash: include ProtCash.</p>
--	--

Figure 9: ProtHawk construction.

the adversary \mathcal{A} does not learn the values of the output coins assuming that the commitment scheme Comm is hiding, and the NIZK scheme we employ is computational zero-knowledge. The adversary \mathcal{A} can observe the nyms that receive the two output coins. However, as we remarked earlier, since these nyms can be one-time, leaking them to the adversary would be okay. Essentially we only need to break linkability at spend time to ensure transactional privacy.

When a corrupted party \mathcal{P}^* pours to an honest party \mathcal{P} , even though the adversary knows the opening of the coin, it cannot spend the coin $(\mathcal{P}, \text{coin})$ once the transaction takes effect by the `ContractCash`, since \mathcal{P}^* cannot demonstrate knowledge of \mathcal{P} 's secret key. We stress that since the contract binds the owner's nym \mathcal{P} to the coin, only the owner can spend it even when the opening of coin is disclosed.

Technical subtleties. Our `ContractCash` uses a modified version of Zerocash to achieve stronger security in the simulation paradigm. In comparison, Zerocash adopts a strictly weaker, indistinguishability-based privacy notion called ledger indistinguishability. In multi-party protocols, indistinguishability-based security notions are strictly weaker than simulation security. Not only so, the particular ledger indistinguishability notion adopted by Zerocash [10] appears subtly questionable upon a careful examination, which we elaborate on in the Appendix. This does not imply that the Zerocash construction is necessarily insecure – however, there is no obvious path to proving their scheme secure under a simulation based paradigm.

4.1.2 Binding Privacy and Programmable Logic

So far, `ContractCash`, similar to Zerocash [10], only supports *direct* money transfers between users. We allow transactional privacy and programmable logic simultaneously.

Freeze. We support a new operation called `freeze`, that does not spend directly to a user, but commits the money as well as an accompanying private input to a smart contract. This is done using a `pour`-like protocol:

- The user \mathcal{P} chooses a private coin $(\mathcal{P}, \text{coin}) \in \text{Coins}$, where $\text{coin} := \text{Comm}_s(\$val)$. Using its secret key, \mathcal{P} computes the serial number `sn` for coin – to be disclosed with the `freeze` operation to prevent double-spending.
- The user \mathcal{P} computes a commitment $(\text{val}||\text{in}||k)$ to the contract where `in` denotes its input, and k is a symmetric encryption key that is introduced due to a practical optimization explained later in Section 4.2.
- The user \mathcal{P} now makes a zero-knowledge proof attesting to similar statements as in a `pour` operation, i.e., that the spent coin exists in the pool `Coins`, the `sn` is correctly constructed, and that the `val` committed to the contract equals the value of the coin being spent. See $\mathcal{L}_{\text{FREEZE}}$ in Figure 9 for details of the NP statement being proven.

Compute. Next, computation takes place off-chain to compute the payout distribution $\{\text{val}'_i\}_{i \in [n]}$ and a proof of correctness. In `Hawk`, we rely on a minimally trusted manager $\mathcal{P}_{\mathcal{M}}$ to perform computation. All parties would open their inputs to the manager $\mathcal{P}_{\mathcal{M}}$, and this is done by encrypting the opening to the manager's public key:

$$\text{ct} := \text{ENC}(\mathcal{P}_{\mathcal{M}}.\text{epk}, r, (\$val||\text{in}||k||s'))$$

The ciphertext `ct` is submitted to the smart contract along

with appropriate zero-knowledge proofs of correctness. While the user can also directly send the opening to the manager off-chain, passing the ciphertext `ct` through the smart contract would make any aborts evident such that the contract can financially punish an aborting user.

After obtaining the openings, the manager now computes the payout distribution $\{\text{val}'_i\}_{i \in [n]}$ and public output `out` by applying the private contract ϕ_{priv} . The manager also constructs a zero-knowledge proof attesting to the outcomes.

Finalize. When the manager submits the outcome of ϕ_{priv} and a zero-knowledge proof of correctness to `ContractHawk`, `ContractHawk` verifies the proof and redistributes the frozen money accordingly. Here `ContractHawk` also passes the manager's public input `inM` and public output `out` to a public contract denoted `C`. The public contract `C` can be invoked to check the validity of the manager's input, as well as redistribute public collateral deposit.

THEOREM 1. *Assuming that the hash function in the Merkle tree is collision resistant, the commitment scheme Comm is perfectly binding and computationally hiding, the NIZK scheme is computationally zero-knowledge and simulation sound extractable, the encryption schemes ENC and SENC are perfectly correct and semantically secure, the PRF scheme PRF is secure, then, our protocols in Figures 8 and 9 securely emulate the ideal functionality $\mathcal{F}(\text{Ideal}_{\text{hawk}})$.*

PROOF. Deferred to the Appendix. \square

4.2 Practical Considerations

Our scheme's main performance bottleneck is computing NIZK proofs. Specifically, NIZK proofs must be computed whenever *i*) a user invokes a `pour`, `freeze`, or a `compute` operation; and *ii*) the manager calls `finalize`. In our implementation, we use state-of-the-art SNARK constructions to efficiently instantiate the NIZK proofs.

Efficient SNARK circuits. A SNARK prover's performance is mainly determined by the number of multiplication gates in the algebraic circuit to be proven [12,46]. To achieve efficiency, we designed optimized circuits through two ways: 1) using cryptographic primitives that are SNARK-friendly, i.e. efficiently realizable as arithmetic circuits under a specific SNARK parametrization. For example, we use a SNARK-friendly collision-resistant hash function [14] to realize the Merkle tree circuit. 2) Building customized circuit generators instead of relying on compilers to translate higher level implementation. For example, our circuits rely also on standard SHA-256, RSA-OAEP encryption, and RSA signature verification (with 2048-bit keys for both), so hand-optimized these components to reduce circuit size.

Subtleties related to the NIZK proofs. Some subtle technicalities arise when we use SNARKs to instantiate NIZK proofs. As mentioned in Theorem 1, we assume that our NIZK scheme satisfies simulation sound extractability. Unfortunately, ordinary SNARKs do not offer simulation sound extractability – and our simulator cannot simply use the SNARK's extractor since the SNARK extractor is non-blackbox, and using the SNARK extractor in a UC-style simulation proof would require running the environment in the extractor!

We therefore rely a generic transformation (see Appendix B.2) to build a simulation sound extractable NIZK from an ordinary SNARK scheme.

Optimizations. We describe an optimization that greatly enhances our performance. We therefore focus on optimizing the $O(N)$ -sized `finalize` circuit since this is our main performance bottleneck. All other SNARK proofs in our scheme are for $O(1)$ -sized circuits. Two key observations allow us to greatly improve the performance of the proof generation during `finalize`.

Optimization 1: Minimize SSE-secure NIZKs. First, we observe that in our proof, the simulator need not extract any new witnesses when a corrupted manager submits proofs during a `finalize` operation. All witnesses necessary will have been learned or extracted by the simulator at this point. Therefore, we can employ an ordinary SNARK instead of a stronger simulation sound extractable NIZK during `finalize`. For `freeze` and `compute`, we still use the stronger NIZK. This optimization reduces our SNARK circuit sizes by $1.6\times$ as shown in Table 2 of Section 5.

Optimization 2: Minimize public-key encryption in SNARKs. Second, during `finalize`, the manager encrypts each party \mathcal{P}_i 's output coins to \mathcal{P}_i 's key, resulting in a ciphertext ct_i . The ciphertexts $\{ct_i\}_{i \in [N]}$ would then be submitted to the contract along with appropriate SNARK proofs of correctness. Here, if a public-key encryption is employed to generate the ct_i 's, it would result in relatively large SNARK circuit size. Instead, we rely on a symmetric-key encryption scheme denoted SENC in Figure 9. This requires that the manager and each \mathcal{P}_i perform a key exchange to establish a symmetric key k_i . During an `compute`, the user encrypts this k_i to the manager's public key $\mathcal{P}_M.epk$, and prove that the k encrypted is consistent with the k committed to earlier in cm_i . The SNARK proof during `finalize` now only needs to include commitments and symmetric encryptions instead of public key encryptions in the circuit – the latter much more expensive. This second optimization additionally gains us a factor of $11\times$ as shown in Table 2 of Section 5.

4.3 Extensions and Discussions

Open enrollment of pseudonyms. In our current formalism, parties' pseudonyms are hardcoded and known a priori. We can easily relax this to allow open enrollment of any pseudonym that joins the contract (e.g., in an auction). Our implementation supports open enrollment.

Refunding frozen coins to users. In our implementation, we extend our basic scheme to allow the users to reclaim their frozen money after a timeout $T_3 > T_2$. To achieve this, user \mathcal{P} simply sends the contract a newly constructed coin ($\mathcal{P}, coin := Comm_s(\$val)$) and proves in zero-knowledge that its value $\$val$ is equal to that of the frozen coin. In this case, the user can identify the previously frozen coin in the clear, i.e., there is no need to compute a zero-knowledge proof of membership within the frozen pool as is needed in a `pour` transaction.

Instantiating the manager with trusted hardware. In some applications, it may be a good idea to instantiate the manager using trusted hardware such as the emerging Intel SGX. In this case, the off-chain computation can take place in a secret SGX enclave that is not visible to any untrusted software or users. Alternatively, in principle, the manager role can also be split into two or more parties that jointly run a secure computation protocol – although this approach is likely to incur higher overhead.

We stress that our model is fundamentally different from placing full trust in any centralized node. *Trusted hardware cannot serve as a replacement of the blockchain.* Any off-chain only protocol that does not interact with the blockchain cannot offer financial fairness in the presence of aborts – even when trusted hardware is employed.

Furthermore, even the use of SGX does not obviate the need for our cryptographic protocol. If the SGX is trusted only by a subset of parties (e.g., just the parties to a particular private contract), rather than globally, then those users can benefit from the efficiency of an SGX-managed private contract, while still utilizing the more widely trusted underlying currency.

Pouring anonymously to long-lived pseudonyms. In our basic formalism of $\text{Ideal}_{\text{cash}}$, the `pour` operation discloses the recipient's pseudonyms to the adversary. This means that $\text{Ideal}_{\text{cash}}$ only retains full privacy if the recipient generates a fresh, new pseudonym every time. In comparison, Zerocash [10] provides an option of anonymously spending to a long-lived pseudonym (in other words, having $\text{Ideal}_{\text{cash}}$ not reveal recipients' pseudonyms to the adversary).

It is not too hard to extend our protocols to allow pouring anonymously to long-lived pseudonyms. However, in most applications, since the transfer is typically subsequent to some interaction between the sender and recipient anyway, we opted for the approach of revealing recipients' pseudonyms. This not only keeps our core formalism simple, but also achieves a constant factor speedup for our protocols.

Other modes of off-chain computation. In our current Hawk compiler, off-chain computation of the private contract ϕ_{priv} (e.g., determining the winner of an auction) is performed by a manager. It is also not hard to extend our formalism to support other modes of off-chain computation, i.e., through an interactive protocol among participants such as secure multiparty computation. We defer the implementation of this to future work.

Remarks about the common reference string. SNARK schemes require the generation of a common reference string (CRS) during a pre-processing step. This common reference string consists of an evaluation key for the prover, and a verification key for the verifier. Unless we employ recursively composed SNARKs whose costs are significantly higher, the evaluation key is circuit-dependent, and its size is proportional to the circuit's size. In comparison, the verification key is $O(|in| + |out|)$ in size, i.e., depends on the total length of inputs and outputs, but independent of the circuit size. We stress that *only the verification key portion of the CRS needs to be included in the public contract that lives on the blockchain.*

We also remark that the CRS for protocol `ProtCash` is shared globally, and can be generated in a one-time setup. In comparison, the CRS for each Hawk contract would depend on the Hawk contract, and therefore exists per instance of Hawk contract. To minimize the trust necessary in the CRS generation, one can employ either trusted hardware or use secure multi-party computation techniques as described by Ben-Sasson et al. [13].

5. IMPLEMENTATION

To demonstrate feasibility and estimate the practical costs, we developed a prototype of Hawk.




Figure 10: Compiler overview. Circuit augmentation for finalize.

5.1 Compiler Implementation

Our compiler consists of several steps, which we illustrate in Figure 10 and describe below:

Preprocessing: First, the input Hawk program is split into its *public contract* and *private contract* components. The public contract is Serpent code, and can be executed directly atop an ordinary cryptocurrency platform such as Ethereum. The private contract is written in a subset of the C language, and is passed as input to the Pinocchio arithmetic circuit compiler [46]. Keywords such as `HawkDeclareParties` are implemented as C preprocessors macros, and serve to define the input (`Inp`) and output (`Outp`) datatypes. Currently, our private contract inherits the limitations of the Pinocchio compiler, e.g., cannot support dynamic-length loops. In the future, we can relax these limitations by employing recursively composition of SNARKs.

Circuit Augmentation: After compiling the preprocessed private contract code with Pinocchio, we have an arithmetic circuit representing the input/output relation ϕ_{priv} . This becomes a subcomponent of a larger arithmetic circuit, which we assemble using a customized circuit assembly tool. This tool is parameterized by the number of parties and the input/output datatypes, and attaches cryptographic constraints, such as computing commitments and encryptions over each party’s output value, and asserting that the input and output values satisfy the balance property.

Cryptographic Protocol: Finally, the augmented arithmetic circuit is used as input to a state-of-the-art zkSNARK library, Libsnark [15]. To avoid implementing SNARK verification in Ethereum’s Serpent language, we must add a SNARK verification opcode to Ethereum’s stack machine. We finally compile an executable program for the parties to compute the Libsnark proofs according to our protocol.

5.2 Additional Examples

Besides our running example of a sealed-bid auction (Figure 2), we implemented several other examples in Hawk, demonstrating various capabilities:

Crowdfunding: A Kickstarter-style crowdfunding campaign, (also known as an assurance contract in economics literature [9]) overcomes the “free-rider problem,” allowing a large number of parties to contribute funds towards some social good. If the minimum donation target is reached before the deadline, then the donations are transferred to a designated party (the entrepreneur); otherwise, the donations are refunded. Hawk preserves privacy in the following sense:

	pour freeze compute			finalize					
	swap	rps	auction	crowdfund					
#Parties	-	-	-	2	2	10	100	10	100
KeyGen(s)	MUL 97.0	85.1	116.3	8.7	8.1	32.3	298.6	34.7	298.5
	ONE 323.6	289.3	407.7	27.6	24.7	122.8	973.8	122.4	970.0
Prove(s)	MUL 40.1	30.6	49.9	3.6	2.8	13.3	166.6	13.9	159.3
	ONE 100.6	85.1	132.8	7.5	7.3	39.5	369.2	39.4	375.9
Verify(ms)	9.7	9.6	9.8	8.1	8.0	9.7	16.6	9.3	13.8
EvalKey(GB)	0.60	0.54	0.76	0.04	0.04	0.21	1.92	0.21	1.91
VerKey(KB)	13.0	11.3	14.4	3.3	2.9	12.9	113.8	12.9	113.8
Proof(B)	288	288	288	288	288	288	288	288	288

Table 1: Performance of the zk-SNARK circuits for pour, freeze and compute (same for all applications), and for finalize in four different applications. The units of quantity are stated in the leftmost column of the table. MUL denotes multiple (4) cores, and ONE denotes a single core. The mint operation does not involve any SNARKs, and can be computed within tens of microseconds.

	Mul gates	Ratio	Encryption	Signature
Naive	155M	18.2×	85%	9%
Opt1	96M	11.3×	91%	-
Opt1,2	8.5M	1×	-	-

Table 2: Gains attained by optimizations. Opt 1 and Opt 2 are two practical optimizations detailed in Section 4.2.

a) the donations pledged are kept private until the deadline; and b) if the contract fails, only the manager learns the amount by which the donations were insufficient. These privacy properties may conceivably have a positive effect on the willingness of entrepreneurs to launch a crowdfunding campaign and its likelihood of success.

Rock Paper Scissors: A two-player lottery game, and naturally generalized to an N -player version. Our Hawk implementation provides the same notion of financial fairness as in [7, 16] and provides stronger security/privacy guarantees. If any party (including the manager), cheats or aborts, the remaining honest parties receive the maximum amount they might have won otherwise. Furthermore, we go beyond prior works [7, 16] by concealing the players’ moves and the pseudonym of the winner to everyone except the manager.

“Swap” Financial Instrument: An individual with a risky investment portfolio (e.g, one who owns a large number of Bitcoins) may hedge his risks by purchasing insurance (e.g., by effectively betting against the price of Bitcoin with another individual). Our example implements a simple swap instrument where the price of a stock at some future date (as reported by a trusted authority specified in the public contract) determines which of two parties receives a payout. The private contract ensures the privacy of both the details of the agreement (i.e., the price threshold) and the outcome.

The full Hawk programs for these examples are provided in the Appendix.

5.3 Performance Evaluation

We evaluated the performance for various examples, using an Amazon EC2 `r3.8xlarge` virtual machine. Our benchmarks actually consume at most 23.7GB of memory and 4 cores in the most expensive case. Table 1 illustrates the results – we focus on evaluating the zk-SNARK performance

since all other computation time is negligible in comparison. We highlight some important observations:

- **On-chain computation** (dominated by zk-SNARK verification time) is very small in all cases, ranging from **8 to 16.6 milliseconds**, while the cryptographic proof is constant (288 bytes). This makes the running time of the verification algorithm just linearly dependent on the size of the public statement, which is far smaller than the size of the computation, resulting into small verification time.
- **On-chain public parameters:** As mentioned in Section 4.3, not the entire SNARK common reference string (CRS) need to be on the blockchain, but only the verification key part of the CRS needs to be on-chain. Our implementation suggests the following: the private cash protocol requires a verification key of 39KB to be stored on-chain – this verification key is globally shared and there is only a single instance. Besides the globally shared public parameters, each Hawk contract will additionally require **13-114 KB** of verification key to be stored on-chain, for 10 to 100 users. This per-contract verification key is circuit-dependent, i.e., depends on the contract program. We refer the readers to Section 4.3 for more discussions on techniques for performing trusted setup.
- **Manager computation:** Running private auction or crowdfunding protocols with 100 participants requires under 6.25min proof time for the manager on a single core, and under **2.8min** on 4 cores. This translates to under **\$0.13** of EC2 time [2].
- **User computation:** Users’ proof times for `pour`, `freeze` and `compute` are under one minute, and independent of the number of parties.

Savings from protocol optimizations. Table 2 evaluates the performance gains attained by our protocol optimizations described in Section 4.2. This table shows the sealed-bid auction example at 100 bidders. We show that our two optimizations combined significantly reduce the SNARK circuit sizes, and achieve a gain of **18×** relative to a straight-forward implementation.

6. FREQUENTLY ASKED QUESTIONS

In this section, we address some of the frequently asked questions. Some of this content repeats what is already stated earlier in the paper, but we hope that addressing these points again in a centralized section will help reiterate some important points that may be missed by a reader.

6.1 Motivational

“Does privacy come at the expense of transparency?”

Our work preserves the most important aspects of transparency; in particular it is *efficiently* and *publicly* verifiable that the balance of all coins is conserved. It is also possible to engineer the cryptography to make other properties publicly auditable too if needed. In general, it is well-known that cryptography allows one to achieve privacy and transparency simultaneously.

We stress that privacy is of *utmost importance* in a variety of financial transactions. For example, Section 5.2 and Appendix A, Figure 13 describe a financial swap/insurance contract. In real life, companies do not wish to reveal their financial dealings to the public, and the present blockchain

technology will preclude any such applications that require privacy. Privacy may also be important in our auction or crowdfunding examples. This is certainly not an exhaustive list, and there are a large variety of financial instruments and contracts (e.g., futures, options, multilateral credit settlement transactions, and double-auctions) that may be modeled in our framework.

“Why are the recipient pseudonyms \mathcal{P}_1 and \mathcal{P}_2 revealed to the adversary? And what about Zerocash’s persistent addresses feature?” For simplicity, we define our protocol (see Section 4.1.1) to reveal the recipients’ pseudonyms. This only provides full privacy when the recipient generates a fresh pseudonym for each transaction. Zerocash takes an alternative approach, where the recipient has a long-term master pseudonym and the sender derives from this an ephemeral public key. It would be straightforward to add this feature to our system as well (at the cost of a constant factor blowup in performance); however, we believe in most applications (e.g., a payment made after receiving an invoice), the transfer is subsequent to some interaction between the recipient and sender. We have added a paragraph to Section 4.3 to discuss this issue.

“How does Hawk’s programming model differ from Ethereum?” Our high-level approach may be superior than Ethereum: Ethereum’s language only defines the public on-chain contract. By contrast, our idea is to allow the programmer to write a single program in a centralized fashion, and we auto-generate not only the on-chain contract, but also the protocols for user agents. Having said this, we remark that presently, for the public portion of the Hawk contract, we support Ethereum’s Serpent language as is.

A gold standard for smart contract languages is non-existent, and it remains an open research question how to design safe and usable programming languages for smart contracts. As Delmolino et al. [26] and Juels et al. [35] demonstrate, smart contract programming is an error-prone process. In future work, we would like to give a more systematic and in-depth study on how to design programming languages for smart contracts.

“Should this be two papers?” We prefer a *single strong paper* with mutually-supporting theoretical and practical results, rather than two thinner papers. The blockchain UC formalization could be presented on its own, but we gain evidence of its usefulness by implementing it and applying it to interesting practical examples. Likewise our system implementation benefits from the formalism because we can use our framework to provide provable security.

6.2 Technical

“SNARKs do not offer simulation extractability required for UC.” Our UC proofs require our NIZKs to be simulation extractable. In general, SNARKs are not simulation extractable and cannot suffice for UC proofs. We stress that our simulator is *not* using the SNARK extractor to extract, since doing so would require running the environment’s algorithm inside the extractor! Section 4.2 discusses our generic transformation from SNARK to simulation extractable NIZK (see paragraph entitled “Subtleties related to the NIZK proofs”) and we formalize this in the Appendix (see Theorem 2). Our experimental results include this transformation, and suggest the practicality of

this approach. This issue is also related to the technical incorrectness of Zerocash (see paragraph entitled “Technical subtleties” in Section 4.1 and Appendix C).

Trust in the manager. “Can’t a dishonest manager forge proofs?” As we mention upfront in Sections 1.1 and 1.2, the manager need not be trusted for correctness and input independence. Further, each contract instance can choose its own manager, and the manager of one contract instance cannot affect the security of another contract instance. Similarly, the manager also need not be trusted to retain the security of the crypto-currency as a whole. Therefore, the only thing we trust the manager for is posterior privacy. To see that the manager need not be trusted for correctness, notice that in our cryptographic construction, the manager must give zero-knowledge proofs to attest to the correctness of its computation, as well as the conservation of money. *The unforgeability of the proofs is cryptographically enforced* by the soundness of the NIZK scheme, and a dishonest manager cannot forge proofs.

As mentioned in Section 4.3 we note that one can possibly rely on secure multi-party computation (MPC) to avoid having to trust the manager even for posterior privacy – however such a solution is unlikely to be practical in the near future, especially when a large number of parties are involved. The theoretical formulation of this full-generality MPC-based approach is detailed in Appendix G.1.

In our implementation, we made a *conscious* design choice and opted for the approach with a minimally trusted manager (rather than MPC), since we believe that this is a desirable *sweet-spot* that simultaneously attains practical efficiency and strong enough security for realistic applications. We stress that practical efficiency is an important goal of Hawk’s design.

In Section 4.3, we also discuss practical considerations for instantiating this manager. For the reader’s convenience, we iterate: we think that a particularly promising choice is to rely on trusted hardware such as Intel SGX to obtain higher assurance of posterior privacy. We stress again that even when we use the SGX to realize the manager, the SGX should not have to be trusted for retaining the global security of the cryptocurrency. In particular, it is a very strong assumption to require all participants to globally trust a single or a handful of SGX processor(s). With Hawk’s design, the SGX is only very minimally trusted, and is only trusted within the scope of the current contract instance.

“Is fair MPC impossible without the blockchain?” As is well known in cryptography, fairness against aborts is in general impossible in the plain model of secure multi-party computation (without a blockchain), when the majority of parties are corrupted [8, 16, 23]. There have been a plethora of results on this topic, beginning with the well-known result by Cleve [23], to more recent works that aim to establish a complete characterization on when fairness is possible and when not [8].

Acknowledgments

We gratefully acknowledge Jonathan Katz, Rafael Pass, and abhi shelat for helpful technical discussions about the zero-knowledge proof constructions. We also acknowledge Ari Juels and Dawn Song for general discussions about cryptocurrency smart contracts. This research is partially sup-

ported by NSF grants CNS-1314857, CNS-1445887, CNS-1518765, a Sloan Fellowship, and two Google Research Awards.

7. REFERENCES

- [1] <http://koinify.com>.
- [2] Amazon ec2 pricing. <http://aws.amazon.com/ec2/pricing/>.
- [3] Augur. <http://www.augur.net/>.
- [4] bitcoinj. <https://bitcoinj.github.io/>.
- [5] The rise and rise of bitcoin. Documentary.
- [6] Skuchain. <http://www.skuchain.com/>.
- [7] M. Andrychowicz, S. Dziembowski, D. Malinowski, and L. Mazurek. Secure Multiparty Computations on Bitcoin. In *S&P*, 2013.
- [8] G. Asharov, A. Beimel, N. Makriyannis, and E. Omri. Complete characterization of fairness in secure two-party computation of boolean functions. In *Theory of Cryptography Conference (TCC)*, pages 199–228, 2015.
- [9] M. Bagnoli and B. L. Lipman. Provision of public goods: Fully implementing the core through private contributions. *The Review of Economic Studies*, 1989.
- [10] E. Ben-Sasson, A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, and M. Virza. Zerocash: Decentralized anonymous payments from Bitcoin. In *S&P*, 2014.
- [11] E. Ben-Sasson, A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, and M. Virza. Zerocash: Decentralized anonymous payments from bitcoin. In *Security and Privacy (SP), 2014 IEEE Symposium on*. IEEE, 2014.
- [12] E. Ben-Sasson, A. Chiesa, D. Genkin, E. Tromer, and M. Virza. Snarks for C: verifying program executions succinctly and in zero knowledge. In *CRYPTO*, 2013.
- [13] E. Ben-Sasson, A. Chiesa, M. Green, E. Tromer, and M. Virza. Secure sampling of public parameters for succinct zero knowledge proofs. In *IEEE Symposium on Security and Privacy*, 2015.
- [14] E. Ben-Sasson, A. Chiesa, E. Tromer, and M. Virza. Scalable zero knowledge via cycles of elliptic curves. In *CRYPTO*, 2014.
- [15] E. Ben-Sasson, A. Chiesa, E. Tromer, and M. Virza. Succinct non-interactive zero knowledge for a von neumann architecture. In *USENIX Security*, 2014.
- [16] I. Bentov and R. Kumaresan. How to Use Bitcoin to Design Fair Protocols. In *CRYPTO*, 2014.
- [17] D. Bogdanov, S. Laur, and J. Willemson. Sharemind: A Framework for Fast Privacy-Preserving Computations. In *ESORICS*. 2008.
- [18] J. Bonneau, A. Miller, J. Clark, A. Narayanan, J. A. Kroll, and E. W. Felten. Research Perspectives and Challenges for Bitcoin and Cryptocurrencies. In *S&P*, 2015.
- [19] R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *FOCS*, 2001.
- [20] R. Canetti. Universally composable signature, certification, and authentication. In *Computer Security Foundations Workshop, 2004. Proceedings. 17th IEEE*, pages 219–233. IEEE, 2004.
- [21] R. Canetti, Y. Dodis, R. Pass, and S. Walfish. Universally composable security with global setup. In

- Theory of Cryptography*, pages 61–85. Springer, 2007.
- [22] R. Canetti and T. Rabin. Universal composition with joint state. In *Advances in Cryptology-Crypto 2003*, pages 265–281. Springer, 2003.
- [23] R. Cleve. Limits on the security of coin flips when half the processors are faulty. In *Proceedings of the Eighteenth Annual ACM Symposium on Theory of Computing*, pages 364–369, 1986.
- [24] G. Danezis, C. Fournet, M. Kohlweiss, and B. Parno. Pinocchio Coin: building Zerocoin from a succinct pairing-based proof system. In *PETShop*, 2013.
- [25] C. Decker and R. Wattenhofer. Bitcoin transaction malleability and mtgox. In *Computer Security-ESORICS 2014*, pages 313–326. Springer, 2014.
- [26] K. Delmolino, M. Arnett, A. Kosba, A. Miller, and E. Shi. Step by step towards creating a safe smart contract: Lessons and insights from a cryptocurrency lab. <https://eprint.iacr.org/2015/460>.
- [27] A. K. R. Dermody and O. Slama. Counterparty announcement. <https://bitcointalk.org/index.php?topic=395761.0>.
- [28] I. Eyal and E. G. Sirer. Majority is not enough: Bitcoin mining is vulnerable. In *FC*, 2014.
- [29] M. Fischlin, A. Lehmann, T. Ristenpart, T. Shrimpton, M. Stam, and S. Tessaro. Random oracles with (out) programmability. In *Advances in Cryptology-ASIACRYPT 2010*, pages 303–320. Springer, 2010.
- [30] C. Fournet, M. Kohlweiss, G. Danezis, and Z. Luo. Zql: A compiler for privacy-preserving data processing. In *USENIX Security*, 2013.
- [31] M. Fredrikson and B. Livshits. Z ϕ : An optimizing distributing zero-knowledge compiler. In *USENIX Security*, 2014.
- [32] J. A. Garay, A. Kiayias, and N. Leonardos. The bitcoin backbone protocol: Analysis and applications. In *Eurocrypt*, 2015.
- [33] E. Heilman, A. Kendler, A. Zohar, and S. Goldberg. Eclipse attacks on bitcoin’s peer-to-peer network. 2015.
- [34] D. Hofheinz and V. Shoup. GNUC: A new universal composability framework. *J. Cryptology*, 28(3):423–508, 2015.
- [35] A. Juels, A. Kosba, and E. Shi. The ring of gyges: Using smart contracts for crime. Manuscript, 2015.
- [36] A. Kiayias, H.-S. Zhou, and V. Zikas. Fair and robust multi-party computation using a global transaction ledger. <http://eprint.iacr.org/2015/574>, 2015.
- [37] B. Kreuter, B. Mood, A. Shelat, and K. Butler. PCF: A portable circuit format for scalable two-party secure computation. In *Usenix Security*, 2013.
- [38] R. Kumaresan and I. Bentov. How to Use Bitcoin to Incentivize Correct Computations. In *CCS*, 2014.
- [39] C. Liu, X. S. Wang, K. Nayak, Y. Huang, and E. Shi. OblivM: A programming framework for secure computation. In *S&P*, 2015.
- [40] S. Meiklejohn, M. Pomarole, G. Jordan, K. Levchenko, D. McCoy, G. M. Voelker, and S. Savage. A fistful of bitcoins: characterizing payments among men with no names. In *IMC*, 2013.
- [41] I. Miers, C. Garman, M. Green, and A. D. Rubin. Zerocoin: Anonymous Distributed E-Cash from Bitcoin. In *S&P*, 2013.
- [42] A. Miller, M. Hicks, J. Katz, and E. Shi. Authenticated data structures, generically. In *POPL*, 2014.
- [43] A. Miller and J. J. LaViola Jr. Anonymous Byzantine Consensus from Moderately-Hard Puzzles: A Model for Bitcoin, 2014.
- [44] M. S. Miller, C. Morningstar, and B. Frantz. Capability-based financial instruments. In *FC*, 2001.
- [45] S. Nakamoto. Bitcoin: A Peer-to-Peer Electronic Cash System. <http://bitcoin.org/bitcoin.pdf>, 2009.
- [46] B. Parno, C. Gentry, J. Howell, and M. Raykova. Pinocchio: Nearly practical verifiable computation. In *S&P*, 2013.
- [47] A. Rastogi, M. A. Hammer, and M. Hicks. Wysteria: A programming language for generic, mixed-mode multiparty computations. In *S&P*, 2014.
- [48] D. Ron and A. Shamir. Quantitative Analysis of the Full Bitcoin Transaction Graph. In *FC*, 2013.
- [49] N. Szabo. Formalizing and securing relationships on public networks. *First Monday*, 1997.
- [50] N. van Saberhagen. Cryptonote v 2.0. <https://cryptonote.org/whitepaper.pdf>, 2013.
- [51] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of finance*, 1961.
- [52] G. Wood. Ethereum: A secure decentralized transaction ledger. <http://gavwood.com/paper.pdf>.
- [53] L. Zheng, S. Chong, A. C. Myers, and S. Zdancewic. Using replication and partitioning to build secure distributed systems. In *S&P*, 2003.
- [54] G. Zyskind, O. Nathan, and A. Pentland. Enigma: Decentralized computation platform with guaranteed privacy.

APPENDIX

A. ADDITIONAL EXAMPLE PROGRAMS

Crowdfunding example. In the crowdfunding example in Figure 12, parties donate money for a kickstarter project. If the total raised funding exceeds a pre-set budget denoted BUDGET, then the campaign is successful and the kickstarter obtains the total donations. Otherwise, all donations are returned to the donors after a timeout. In this case, no public deposit is necessary to ensure the incentive compatibility of the contract. If a party does not open after freezing its money, the money is unrecoverable by anyone.

Swap instrument example. In this financial swap instrument, Alice is betting on the stock price exceeding a certain threshold at a future point of time, while Bob is betting on the reverse. If the stock price is below the threshold, Alice obtains \$20; else Bob obtains \$20. As mentioned earlier in Section 5.2, such a financial swap can be used as a means of insurance to hedge investment risks. This swap contract makes use of public deposits to provide financial fairness when either Alice or Bob cheats.

This swap assumes that the manager is a well-known public entity such as a stock exchange. Therefore, the contract does not protect against the manager aborting. In the event that the manager aborts, the aborting event can be ob-

```

1  typedef enum {ROCK, PAPER, SCISSORS} Move;
2  typedef enum {DRAW, WIN, LOSE} Outcome;
3  typedef enum {OK, A_CHEAT, B_CHEAT} Output;

4  // Private Contract parameters
5  HawkDeclareParties(Alice, Bob);
6  HawkDeclareTimeouts(/* hardcoded timeouts */);
7  HawkDeclareInput(Move move);

8  Outcome outcome(Move a, Move b) {
9      return (a - b) % 3;
10 }

11 private contract game(Inp &in, Outp &out) {
12     if (in.Alice.$val != $1) out.out = A_CHEAT;
13     if (in.Bob.$val != $1) out.out = B_CHEAT;
14     Outcome o = outcome(in.Alice.move, in.Bob.move);
15     if (o == WIN) out.Alice.$val = $2;
16     else if (o == LOSE) out.Bob.$val = $2;
17     else out.Alice.$val = out.Bob.$val = $1;
18 }

19 public contract deposit() {
20     // Alice and Bob each deposit $2
21     // Manager deposits $4
22     def check(Output o):
23         send $4 to Manager
24         if (o == A_CHEAT): send $4 to Bob
25         if (o == B_CHEAT): send $4 to Alice
26         if (o == OK):
27             send $2 to Alice
28             send $2 to Bob
29     def managerTimedOut():
30         send $4 to Bob
31         send $4 to Alice
32 }

```

Figure 11: Hawk program for a rock-paper-scissors game. This program defines both a private contract and a public contract. The private contract guarantees that only Alice, Bob, and the Manager learn the outcome of the game. Public collateral deposits are used to guarantee financial fairness such that if any of the parties cheat, the remaining honest parties receive monetary compensation.

served in public, and therefore external mechanisms (e.g., legal enforcement or reputation) can be leveraged to punish the manager.

We provide the Hawk programs for the applications used in our evaluation in Section 5. For the sealed auction contract, please refer to Section 1.2.

Rock-Paper-Scissors example. In this lottery game in Figure 11, each party deposits \$3 in total. In the case that all parties are honest, then each party has a 50% chance of leaving with \$4 (i.e., winning \$1) and a 50% chance of leaving with \$2 (i.e., losing \$2).

The lottery game is fair in the following sense: if any party cheats, then the remaining honest parties are guaranteed a payout distribution that *stochastically dominates* the payout distribution they would expect if every party was honest.

This is achieved using standard “collateral deposit” techniques [7, 16]. For example, if Alice aborts, then her deposit is used to compensate Bob by the maximum amount \$4. If the Manager aborts, then both Alice and Bob receive \$8.

Unlike the lottery games found in Bitcoin and Ethereum [7, 16, 26], our contract also provides privacy. If the Manager and both parties do not voluntarily disclose information, then no one else in the system learns which of Alice or Bob won. Even when the Manager, Alice, and Bob are all

```

1  // Raise $10,000 from up to N donors
2  #define BUDGET $10000

3  HawkDeclareParties(Entrepreneur, /* N Parties */);
4  HawkDeclareTimeouts(/* hardcoded timeouts */);

5  private contract crowdfund(Inp &in, Outp &out) {
6      int sum = 0;
7      for (int i = 0; i < N; i++) {
8          sum += in.p[i].$val;
9      }
10     if (sum >= BUDGET) {
11         // Campaign successful
12         out.Entrepreneur.$val = sum;
13     } else {
14         // Campaign unsuccessful
15         for (int i = 0; i < N; i++) {
16             out.p[i].$val = in.p[i].$val; // refund
17         }
18     }
19 }

```

Figure 12: Hawk contract for a kickstarter-style crowdfunding contract. No public portion is required. An attacker who freezes but does not open would not be able to recover his money.

corrupted, the underlying ecash cash system still provides privacy for other contracts and guarantees that the total amount of money is conserved.

B. PRELIMINARIES

Notation. In the remainder of the paper, $f(\lambda) \approx g(\lambda)$ means that there exists a negligible function $\nu(\lambda)$ such that $|f(\lambda) - g(\lambda)| < \nu(\lambda)$.

B.1 Non-Interactive Zero-Knowledge Proofs

A non-interactive zero-knowledge proof system (NIZK) for an NP language \mathcal{L} consists of the following algorithms:

- $\text{crs} \leftarrow \mathcal{K}(1^\lambda, \mathcal{L})$, also written as $\text{crs} \leftarrow \text{KeyGen}_{\text{nizk}}(1^\lambda, \mathcal{L})$: Takes in a security parameter λ , a description of the language \mathcal{L} , and generates a common reference string crs .
- $\pi \leftarrow \mathcal{P}(\text{crs}, \text{stmt}, w)$: Takes in crs , a statement stmt , a witness w such that $(\text{stmt}, w) \in \mathcal{L}$, and produces a proof π .
- $b \leftarrow \mathcal{V}(\text{crs}, \text{stmt}, \pi)$: Takes in a crs , a statement stmt , and a proof π , and outputs 0 or 1, denoting accept or reject.

Perfect completeness. A NIZK system is said to be perfectly complete, if an honest prover with a valid witness can always convince an honest verifier. More formally, for any $(\text{stmt}, w) \in R$, we have that

$$\Pr \left[\text{crs} \leftarrow \mathcal{K}(1^\lambda, \mathcal{L}), \pi \leftarrow \mathcal{P}(\text{crs}, \text{stmt}, w) : \mathcal{V}(\text{crs}, \text{stmt}, \pi) = 1 \right] = 1$$

Computational zero-knowledge. Informally, a NIZK system is computationally zero-knowledge, if the proof does not reveal any information about the witness to any polynomial-time adversary. More formally, a NIZK system is said to be computationally zero-knowledge, if there exists a polynomial-time simulator $S = (\hat{\mathcal{K}}, \hat{\mathcal{P}})$, such that for all non-uniform

Table 3: Notations.

ϕ_{priv}	user-defined private Hawk contract. Specifically, $(\{\$val'_i\}_{i \in [N]}, \text{out}) := \phi(\{\$val_i, \text{in}_i\}_{i \in [N]}, \text{in}_{\mathcal{M}})$, i.e., ϕ takes in the parties' private inputs $\{\text{in}_i\}_{i \in [N]}$, private coin values $\{\text{val}_i\}_{i \in [N]}$, the manager's public input $\mathcal{P}_{\mathcal{M}}$, and outputs the payout of each party $\{\$val'_i\}_{i \in [N]}$, and a public output out .
ϕ_{pub}	user-defined public Hawk contract.
IdealP	ideal program
simP	simulator program
C, ContractP	contract program
prot, ProtP	user-side program
$\mathcal{F}(\cdot)$	ideal functionality wrapper, $\mathcal{F}(\text{IdealP})$ denotes an ideal functionality
$\mathcal{G}(\cdot)$	contract functionality wrapper, $\mathcal{G}(\text{C})$ denotes a contract functionality
$\Pi(\cdot)$	protocol wrapper, $\Pi(\text{prot})$ denotes user-side protocol
\mathcal{P}	party or its pseudonym
$\mathcal{P}_{\mathcal{M}}$	minimally trusted manager (or its pseudonym)
\mathcal{A}	adversary
\mathcal{E}	environment
T	current time
ledger	global public ledger
Coins (in ideal programs)	private ledger, maintained by the ideal functionality
Coins (in contract programs)	a set of cryptographic coins stored by a contract. Private spending (including pours and freezes) must demonstrate a zero-knowledge proof of the spent coin's membership in Coins . Further, private spending must demonstrate a cryptographic serial number sn that prevents double spending.

polynomial-time adversary \mathcal{A} , we have that

$$\begin{aligned} & \Pr \left[\text{crs} \leftarrow \mathcal{K}(1^\lambda, \mathcal{L}) : \mathcal{A}^{\mathcal{P}(\text{crs}, \cdot, \cdot)} = 1 \right] \\ & \approx \Pr \left[(\widehat{\text{crs}}, \tau) \leftarrow \widehat{\mathcal{K}}(1^\lambda, \mathcal{L}) : \mathcal{A}^{\widehat{\mathcal{P}}_1(\widehat{\text{crs}}, \tau, \cdot, \cdot)} = 1 \right] \end{aligned}$$

In the above, $\widehat{\mathcal{P}}_1(\widehat{\text{crs}}, \tau, \text{stmt}, w)$ verifies that $(\text{stmt}, w) \in \mathcal{L}$, and if so, outputs $\widehat{\mathcal{P}}(\widehat{\text{crs}}, \tau, \text{stmt})$ which simulates a proof without knowing a witness.

Computational soundness. A NIZK scheme for the language \mathcal{L} is said to be computationally sound, if for all polynomial-time adversaries \mathcal{A} ,

$$\Pr \left[\begin{array}{l} \text{crs} \leftarrow \mathcal{K}(1^\lambda, \mathcal{L}), (\text{stmt}, \pi) \leftarrow \mathcal{A}(\text{crs}) : \\ \mathcal{V}(\text{crs}, \text{stmt}, \pi) = 1 \wedge (\text{stmt} \notin \mathcal{L}) \end{array} \right] \approx 0$$

Simulation sound extractability. Simulation sound extractability says that even after seeing many simulated proofs, whenever the adversary makes a new proof, a simulator is able to extract a witness. Simulation extractability implies simulation soundness and non-malleability, since if the simulator can extract a valid witness from an adversary's proof, the statement must belong to the language. More formally, a NIZK system is said to be simulation sound extractable, if there exist polynomial-time algorithms $(\widehat{\mathcal{K}}, \widehat{\mathcal{P}}, \mathcal{E})$, such that for any polynomial-time adversary \mathcal{A} , it holds that

$$\Pr \left[\begin{array}{l} (\widehat{\text{crs}}, \tau, \text{ek}) \leftarrow \widehat{\mathcal{K}}(1^\lambda, \mathcal{L}); \\ (\text{stmt}, \pi) \leftarrow \mathcal{A}^{\widehat{\mathcal{P}}(\widehat{\text{crs}}, \tau, \cdot)}(\widehat{\text{crs}}, \text{ek}); \\ w \leftarrow \mathcal{E}(\widehat{\text{crs}}, \text{ek}, \text{stmt}, \pi) : \text{stmt} \notin Q \text{ and} \\ (\text{stmt}, w) \notin \mathcal{L} \text{ and } \mathcal{V}(\widehat{\text{crs}}, \text{stmt}, \pi) = 1 \end{array} \right] \approx 0$$

In the above, Q is the list of simulation queries. Here the $\widehat{\mathcal{K}}$ is identical to the zero-knowledge simulation setup algorithm when restricted to the first two terms.

Note that in the above definition, the adversary may be able to fake a (different) proof for a statement that has been queried, however, it is not able to forge a proof for any other invalid statement. There is a natural strengthening of the above notion where the adversary cannot even fake a different proof for a statement queried. In our paper, however, the weaker notion defined above would suffice.

B.2 Simulation Sound Extractable Proofs from Ordinary NIZKs

We give a generic transformation that turns an ordinary NIZK into one that satisfies simulation sound extractability. In our implementation, we use Zero-Knowledge Succinct Non-interactive ARGuments of Knowledge (SNARKs) as the underlying NIZK, and apply our transformation here to make it a NIZK with simulation sound extractability. For this reason, in the transformation described below, we simply use **snark** to denote the underlying NIZK.

- $\mathcal{K}(1^\lambda, \mathcal{L})$: Run $(\text{pk}, \text{sk}) \leftarrow \Sigma.\text{Gen}(1^\lambda)$. Run $(\text{pk}_e, \text{sk}_e) \leftarrow \text{KeyGen}_{\text{enc}}(1^\lambda)$.

Let \mathcal{L}' be the following language: $((\text{stmt}, c), (r, w, \sigma)) \in \mathcal{L}'$ iff

$$\begin{aligned} & (c = \text{Enc}(\text{pk}_e, (w, \sigma), r)) \wedge \\ & ((\text{stmt}, w) \in \mathcal{L} \vee (\Sigma.\mathcal{V}(\text{pk}, \text{stmt}, \sigma) = 1)) \end{aligned}$$

Run $\text{snark.crs} \leftarrow \text{snark}.\mathcal{K}(1^\lambda, \mathcal{L}')$.

Publish $\text{crs} := (\text{snark.crs}, \text{pk}, \text{pk}_e)$ as the common reference string.

- $\mathcal{P}(\text{crs}, \text{stmt}, w)$: Parse $\text{crs} := (\text{snark.crs}, \text{pk})$. Choose random r , and compute $c := \text{Enc}(\text{pk}_e, (w, \sigma), r)$. Call $\pi := \text{snark}.\mathcal{P}(\text{snark.crs}, (\text{stmt}, c), (r, w, \perp))$, and output $\pi' := (c, \pi)$.
- $\mathcal{V}(\text{crs}, \text{stmt}, \pi')$: Parse $\pi' := (c, \pi)$, and output $\text{snark}.\mathcal{V}(\text{snark.crs}, (\text{stmt}, c), \pi)$.

```

1  typedef enum {OK, A_CHEAT, B_CHEAT} Output
2  HawkDeclareParties(Alice, Bob);
3  HawkDeclareTimeouts(/* hardcoded timeouts */);
4  HawkDeclarePublicInput(int stockprice,
                        int threshold[5]);
5  HawkDeclareOutput(Output o);

6  int threshold_comm[5] = {/* hardcoded */};

6  private contract swap(Inp &in, Outp &out) {
7    if (sha1(in.Alice.threshold) != threshold_comm)
        out.o = A_CHEAT;
7    if (in.Alice.$val != $10) out.o = A_CHEAT;
8    if (in.Bob.$val != $10) out.o = B_CHEAT;
8
9    if (in.stockprice < in.Alice.threshold[0])
        out.Alice.$val = $20;
10   else out.Bob.$val = $20;
11 }

12 public contract deposit {
13   def receiveStockPrice(stockprice):
14     // Alice and Bob each deposits $10
15     // Assume the stock price authority is trusted
16     // to send this contract the price
17     assert msg.sender == StockPriceAuthority
18     self.stockprice = stockprice
19   def check(int stockprice, Output o):
20     assert stockprice == self.stockprice
21     if (o == A_CHEAT): send $20 to Bob
22     if (o == B_CHEAT): send $20 to Alice
23     if (o == OK):
24       send $10 to Alice
25       send $10 to Bob
26 }

```

Figure 13: Hawk program for a risk-swap financial instrument. In this case, we assume that the manager is a well-known entity such as a stock exchange, and therefore the contract does not protect against the manager defaulting. An aborting manager (e.g., a stock exchange) can be held accountable through external means such as legal enforcement or reputation, since aborting is observable by the public.

THEOREM 2. *Assume that the SNARK scheme satisfies perfect completeness, computational soundness, and computational zero-knowledge, and that the encryption scheme is perfectly correct, then the above construction is a zero-knowledge proof system satisfying perfect completeness, computational zero-knowledge, and simulation sound extractability.*

PROOF. The proofs of perfect completeness and computational zero-knowledge are obvious. We now show that this transformation gives a simulation sound extractable NIZK. We construct the following simulation and extractor:

- $\widehat{\mathcal{K}}(1^\lambda, \mathcal{L})$: Run the honest \mathcal{K} algorithm, but retain the signing key sk as the simulation trapdoor $\tau := \text{sk}$. The extraction key $\text{ek} := \text{sk}_e$, the simulated $\widehat{\text{crs}} := \text{crs} = (\text{snark.crs}, \text{pk}, \text{pk}_e)$.

- $\widehat{\mathcal{P}}(\widehat{\text{crs}}, \tau, \text{stmt})$: the simulator calls

$$\pi := \text{snark.P}(\text{snark.crs}, (\text{stmt}, c), (\perp, \perp, \sigma))$$

where c is an encryption of 0, and

$$\sigma := \Sigma.\text{Sign}(\text{sk}, \text{stmt})$$

Output (c, π) .

- $\mathcal{E}(\widehat{\text{crs}}, \text{ek}, \text{stmt}, \pi')$: parse $\pi' := (c, \pi)$, and let $(w, \sigma) := \text{Dec}(\text{sk}_e, c)$. Output w .

We now show that no polynomial-time adversary \mathcal{A} can win the simulation sound extractable game except with negligible probability. Given that the encryption scheme is perfectly correct, and that snark is computationally sound, the witness (w, σ) decrypted by \mathcal{E} must satisfy one of the following two cases except with negligible probability: 1) w is a valid witness for stmt under language \mathcal{L} ; or 2) σ is a valid signature for stmt . If stmt has not been queried by the adversary \mathcal{A} , then it must be that w is a valid witness for stmt , since otherwise, the simulator can easily leverage the adversary to break the security of the signature scheme. \square

C. TECHNICAL SUBTLETIES IN ZEROCASH

In general, a simulation-based security definition is more straightforward to write and understand than ad-hoc indistinguishability games – although it is often more difficult to prove or require a protocol with more overhead. Below we highlight a subtle weakness with Zerocash’s security definition [10], which motivates our stronger definition.

Ledger Indistinguishability leaks unintended information. The privacy guarantees of Zerocash [10] are defined by a “Ledger Indistinguishability” game (in [10], Appendix C.1). In this game, the attacker (adaptively) generates two sequences of queries, Q_{left} and Q_{right} . Each query can either be a raw “insert” transaction (which corresponds in our model to a transaction submitted by a corrupted party) or else a “mint” or “pour” query (which corresponds in our model to an instruction from the environment to an honest party). The attacker receives (incrementally) a pair of views of protocol executions, V_{left} and V_{right} , according to one of the following two cases, and tries to discern which case occurred: either V_{right} is generated by applying all the queries in Q_{right} and respectively for V_{right} ; or else V_{left} is generated by interleaving the “insert” queries of Q_{left} with the “mint” and “pour” queries of Q_{right} , and V_{right} is generated by interleaving the “insert” queries of Q_{right} with the “mint” and “pour” queries of Q_{left} . The two sequences of queries are constrained to be “publicly consistent”, which effectively defines the information leaked to the adversary. For example, the i^{th} queries in both sequences must be of the same type (either “mint”, “pour”, or “insert”), and if a “pour” query includes an output to a corrupted recipient, then the output value must be the same in both queries.

However, the definition of “public consistency” is subtly overconstraining: it requires that if the i^{th} query in one sequence is an (honest) “pour” query that spends a coin previously created by a (corrupt) “insert” query, then the i^{th} queries in both sequences must spend coins of equal value created by prior “insert” queries. Effectively, this means that if a corrupted party sends a coin to an honest party, then the adversary may be alerted when the honest party spends it.

We stress that this does not imply any flaw with the Zerocash construction itself — however, there is no obvious path to proving their scheme secure under a simulation based paradigm.

D. SIMULATOR WRAPPER

Before we describe our proofs, we first define a simulator wrapper that will help modularize our proofs.

Simulator wrapper \mathcal{S} : The ideal adversary \mathcal{S} can typically be obtained by applying the simulator wrapper $\mathcal{S}(\cdot)$ to the user-defined portion of the simulator simP . The simulator wrapper modularizes the simulator construction by factoring out the common part of the simulation pertaining to all protocols in this model of execution.

The simulator is defined formally in Figure 14.

E. FORMAL PROOF FOR PRIVATE CASH

We now prove that the protocol in Figure 8 is a secure and correct implementation of $\mathcal{F}(\text{Ideal}_{\text{hawk}})$. For any real-world adversary \mathcal{A} , we construct an ideal-world simulator \mathcal{S} , such that no polynomial-time environment \mathcal{E} can distinguish whether it is in the real or ideal world. We first describe the construction of the simulator \mathcal{S} and then argue the indistinguishability of the real and ideal worlds.

THEOREM 3. *Assuming that the hash function in the Merkle tree is collision resistant, the commitment scheme Comm is perfectly binding and computationally hiding, the NIZK scheme is computationally zero-knowledge and simulation sound extractable, the encryption schemes ENC and SENC are perfectly correct and semantically secure, the PRF scheme PRF is secure, then our protocol in Figure 8 securely emulates the ideal functionality $\mathcal{F}(\text{Ideal}_{\text{cash}})$.*

E.1 Ideal World Simulator

Due to Canetti [19], it suffices to construct a simulator \mathcal{S} for the dummy adversary that simply passes messages to and from the environment \mathcal{E} . The ideal-world simulator \mathcal{S} also interacts with the $\mathcal{F}(\text{Ideal}_{\text{cash}})$ ideal functionality. Below we construct the user-defined portion of our simulator simP . Our ideal adversary \mathcal{S} can be obtained by applying the simulator wrapper $\mathcal{S}(\text{simP})$. The simulator wrapper modularizes the simulator construction by factoring out the common part of the simulation pertaining to all protocols in this model of execution.

Recall that the simulator wrapper performs the ordinary setup procedure, but retains the “trapdoor” information used in creating the crs for the NIZK proof system, allowing it to forge proofs for false statement and to extract witnesses from valid proofs. Since the real world adversary would see the entire state of the contract, the simulator allows the environment to see the entire state of the local instance of the contract. The environment can also submit transactions directly to the contract on behalf of corrupt parties. Such a **pour** transaction contains a zero-knowledge proof involving the values of coins being spent or created; the simulator must rely on its ability to extract witnesses in order to learn these values and trigger $\mathcal{F}(\text{Ideal}_{\text{cash}})$ appropriately.

The environment may also send **mint** and **pour** instructions to honest parties that in the ideal world would be forwarded directly to $\mathcal{F}(\text{Ideal}_{\text{cash}})$. These activate the simulator, but only reveal partial information about the instruction – in particular, the simulator does not learn the values of the coins being spent. The simulator handles this by writing *bogus* (but plausible-seeming) information to the contract.

Thus the simulator must translate transactions submitted by corrupt parties to the contract into ideal world instructions, and must translate ideal world instructions into transactions published on the contract.

The simulator simP is defined in more detail below:

Init. The simulator simP runs $(\widehat{\text{crs}}, \tau, \text{ek}) \leftarrow \text{NIZK}.\widehat{\mathcal{K}}(1^\lambda)$, and gives $\widehat{\text{crs}}$ to the environment \mathcal{E} .

Simulating corrupted parties. The following messages are sent by the environment \mathcal{E} to the simulator $\mathcal{S}(\text{simP})$ which then forwards it on to both the internally simulated contract $\text{G}(\text{ContractCash})$ and the inner simulator simP .

- simP receives a pseudonymous mint message $(\text{mint}, \$\text{val}, r)$. No extra action is necessary.
- simP receives an anonymous pour message, $(\text{pour}, \{\text{sn}_i, \mathcal{P}_i, \text{coin}_i, \text{ct}_i\}_{i \in \{1,2\}}\})$. The simulator uses τ to extract the witness from π , which includes the sender \mathcal{P} and values $\$\text{val}_1, \$\text{val}_2, \$\text{val}'_1$ and $\$\text{val}'_2$. If \mathcal{P}_i is an uncorrupted party, then the simulator must check whether each encryption ct_i is performed correctly, since the NIZK proof does not guarantee that this is the case. The simulator performs a trial decryption using $\mathcal{P}_i.\text{esk}$; if the decryption is *not* a valid opening of coin_i , then the simulator must avoid causing \mathcal{P}_i in the ideal world to output anything (since \mathcal{P}_i in the real world would not output anything either). The simulator therefore substitutes some default value (e.g., the name of any corrupt party \mathcal{P}) for the recipient’s pseudonym. The simulator forwards $(\text{pour}, \$\text{val}_1, \$\text{val}_2, \mathcal{P}_1^\dagger, \mathcal{P}_2^\dagger, \$\text{val}'_1, \$\text{val}'_2)$ anonymously to $\mathcal{F}(\text{Ideal}_{\text{cash}})$, where $\mathcal{P}_i^\dagger = \mathcal{P}$ if \mathcal{P}_i is uncorrupted and decryption fails, and $\mathcal{P}_i^\dagger = \mathcal{P}_i$ otherwise.

Simulating honest parties. When the environment \mathcal{E} sends inputs to honest parties, the simulator \mathcal{S} needs to simulate messages that corrupted parties receive, from honest parties or from functionalities in the real world. The honest parties will be simulated as below:

- **GenNym**(epk, spk): The simulator simP generates and records the PRF keypair, $(\text{pk}_{\text{PRF}}, \text{sk}_{\text{PRF}})$ and returns $\text{payload} := \text{pk}_{\text{PRF}}$.
- Environment \mathcal{E} gives a **mint** instruction to party \mathcal{P} . The simulator simP receives $(\text{mint}, \mathcal{P}, \$\text{val}, r)$ from the ideal functionality $\mathcal{F}(\text{Ideal}_{\text{cash}})$. The simulator has enough information to run the honest protocol, and posts a valid **mint** transaction to the contract.
- Environment \mathcal{E} gives a **pour** instruction to party \mathcal{P} . The simulator simP receives $(\text{pour}, \mathcal{P}_1, \mathcal{P}_2)$ from $\mathcal{F}_{\text{CASH}}$. However, the simulator does not learn the name of the honest sender \mathcal{P} , or the correct values for each input coin val_i (for $i \in \{1, 2\}$). Instead, the simulator uses τ to create a false proof using arbitrary values for these values in the witness. To generate each serial number sn_i in the witness, the simulator chooses a random element from the codomain of PRF. For each recipient \mathcal{P}_i (for $i \in \{1, 2\}$), the simulator behaves differently depending on whether or not \mathcal{P}_i is corrupted:
 - Case 1: \mathcal{P}_i is *honest*. The simulator does not know the correct output value, so instead sets $\text{val}'_i := 0$, and computes coin'_i and ct_i as normal. The environment therefore sees a commitment and an encryption of 0, but without $\mathcal{P}_i.\text{esk}$ it cannot distinguish between an encryption of 0 or of the correct value.
 - Case 2: \mathcal{P}_i is *corrupted*. Since the ideal world recipient would receive $\$\text{val}'_i$ from $\mathcal{F}_{\text{CASH}}$, and since \mathcal{P}_i is corrupted, the simulator learns the correct value $\$\text{val}'_i$.

$\mathcal{S}(\text{simP})$

Init. The simulator \mathcal{S} simulates a $\mathcal{G}(\text{contract})$ instance internally. Here \mathcal{S} calls $\mathcal{G}(\text{contract}).\text{Init}$ to initialize the internal states of the contract functionality. \mathcal{S} also calls $\text{simP}.\text{Init}$.

Simulating honest parties.

- **Tick:** Environment \mathcal{E} sends input `tick` to an honest party P : simulator \mathcal{S} receives notification `(tick, P)` from the ideal functionality. Simulator forwards the tick message to the simulated $\mathcal{G}(\text{contract})$ functionality.
- **GenNym:** Environment \mathcal{E} sends input `gennym` to an honest party P : simulator \mathcal{S} receives notification `gennym` from the ideal functionality. Simulator \mathcal{S} honestly generates an encryption key and a signing key as defined in Figure 5, and remembers the corresponding secret keys. Simulator \mathcal{S} now calls $\text{simP}.\text{GenNym}(\text{epk}, \text{spk})$ and waits for the returned value `payload`. Finally, the simulator passes the nym $\bar{P} = (\text{epk}, \text{spk}, \text{payload})$ to the ideal functionality.
- **Other activations.** // *From the inner idealP*
If ideal functionality sends `(transfer, $amount, Pr, Ps)`, then update the ledger in the simulated $\mathcal{G}(\text{Contract})$ instance accordingly.
Else, forward the message to the inner `simP`.

Simulating corrupted parties.

- **Permute:** Upon receiving `(permute, perm)` from the environment \mathcal{E} , forward it to the internally simulated $\mathcal{G}(\text{contract})$ and the ideal functionality.
- **Expose.** Upon receiving `exposestate` from the environment \mathcal{E} , expose all states of the internally simulated $\mathcal{G}(\text{contract})$.
- **Other activations.**
 - Upon receiving `(authenticated, m)` from the environment \mathcal{E} on behalf of corrupted party P : Forward to internally simulated $\mathcal{G}(\text{contract})$. If the message is of the format `(transfer, $amount, Pr, Ps)`, then forward it to the ideal functionality. Otherwise, forward to `simP`.
 - Upon receiving `(pseudonymous, m, \bar{P} , σ)` from the environment \mathcal{E} on behalf of corrupted party P : Forward to internally simulated $\mathcal{G}(\text{contract})$. Now, assert that σ verifies just like in $\mathcal{G}(\text{contract})$. If the message is of the format `(transfer, $amount, Pr, Ps)`, then forward it to the ideal functionality. Else, forward to `simP`.
 - Upon receiving `(anonymous, m)` from the environment \mathcal{E} on behalf of corrupted party P : Forward to internally simulated $\mathcal{G}(\text{contract})$. If the message is of the format `(transfer, $amount, Pr, Ps)`, then forward it to the ideal functionality. Else, forward to `simP`.

Figure 14: Simulator wrapper.

directly. Hence `coini` is a correct encryption of `$vali` under \mathcal{P}_i 's registered encryption public key.

E.2 Indistinguishability of Real and Ideal Worlds

To prove indistinguishability of the real and ideal worlds from the perspective of the environment, we will go through a sequence of hybrid games.

Real world. We start with the real world with a dummy adversary that simply passes messages to and from the environment \mathcal{E} .

Hybrid 1. Hybrid 1 is the same as the real world, except that now the adversary (also referred to as the simulator) will call $(\widehat{\text{crs}}, \tau, \text{ek}) \leftarrow \text{NIZK}.\widehat{\mathcal{K}}(1^\lambda)$ to perform a simulated setup for the NIZK scheme. The simulator will pass the simulated `crs` to the environment \mathcal{E} . When an honest party \mathcal{P} publishes a NIZK proof, the simulator will replace the real proof with a simulated NIZK proof before passing it onto the environment \mathcal{E} . The simulated NIZK proof can be computed by calling the $\text{NIZK}.\widehat{\mathcal{P}}(\widehat{\text{crs}}, \tau, \cdot)$ algorithm which takes only the statement as input but does not require knowledge of a

witness.

FACT 1. It is immediately clear that if the NIZK scheme is computational zero-knowledge, then no polynomial-time environment \mathcal{E} can distinguish Hybrid 1 from the real world except with negligible probability.

Hybrid 2. The simulator simulates the $\mathcal{G}(\text{ContractCash})$ functionality. Since all messages to the $\mathcal{G}(\text{ContractCash})$ functionality are public, simulating the contract functionality is trivial. Therefore, Hybrid 2 is identically distributed as Hybrid 1 from the environment \mathcal{E} 's view.

Hybrid 3. Hybrid 3 is the same as Hybrid 2 except the following changes. When an honest party sends a message to the contract (now simulated by the simulator \mathcal{S}), it will sign the message with a signature verifiable under an honestly generated nym. In Hybrid 3, the simulator will replace all honest parties' nym and generate these nym itself. In this way, the simulator will simulate honest parties' signatures by signing them itself. Hybrid 3 is identically distributed as Hybrid 2 from the environment \mathcal{E} 's view.

Hybrid 4. Hybrid 4 is the same as Hybrid 3 except for the following changes:

- When an honest party \mathcal{P} produces a ciphertext ct_i for a recipient \mathcal{P}_i , and if the recipient is also uncorrupted, then the simulator will replace this ciphertext with an encryption of 0 before passing it onto the environment \mathcal{E} .
- When an honest party \mathcal{P} produces a commitment coin, then the simulator replaces this commitment with a commitment to 0.
- When an honest party \mathcal{P} computes a pseudorandom serial number sn , the simulator replaces this with a randomly chosen value from the codomain of PRF.

FACT 2. *It is immediately clear that if the encryption scheme is semantically secure, if PRF is a pseudorandom function, and if Comm is a perfectly hiding commitment scheme, then no polynomial-time environment \mathcal{E} can distinguish Hybrid 4 from Hybrid 3 except with negligible probability.*

Hybrid 5. Hybrid 5 is the same as Hybrid 4 except for the following changes. Whenever the environment \mathcal{E} passes to the simulator \mathcal{S} a message signed on behalf of an honest party's nym, if the message and signature pair was not among the ones previously passed to the environment \mathcal{E} , then the simulator \mathcal{S} aborts.

FACT 3. *Assume that the signature scheme employed is secure; then the probability of aborting in Hybrid 5 is negligible. Notice that from the environment \mathcal{E} 's view, Hybrid 5 would otherwise be identically distributed as Hybrid 4 modulo aborting.*

Hybrid 6. Hybrid 6 is the same as Hybrid 5 except for the following changes. Whenever the environment passes $(\text{pour}, \pi, \{\text{sn}_i, \mathcal{P}_i, \text{coin}_i, \text{ct}_i\})$ to the simulator (on behalf of corrupted party \mathcal{P}), if the proof π verifies under **statement**, then the simulator will call the NIZK's extractor algorithm \mathcal{E} to extract witness. If the NIZK π verifies but the extracted witness does not satisfy the relation $\mathcal{L}_{\text{POUR}}(\text{statement}, \text{witness})$, then abort the simulation.

FACT 4. *Assume that the NIZK is simulation sound extractable, then the probability of aborting in Hybrid 6 is negligible. Notice that from the environment \mathcal{E} 's view, Hybrid 6 would otherwise be identically distributed as Hybrid 5 modulo aborting.*

Finally, observe that Hybrid 6 is computationally indistinguishable from the ideal simulation \mathcal{S} unless one of the following bad events happens:

- A value val' decrypted by an honest recipient is different from that extracted by the simulator. However, given that the encryption scheme is perfectly correct, this cannot happen.
- A commitment coin is different than any stored in `ContractCash.coins`, yet it is valid according to the relation $\mathcal{L}_{\text{POUR}}$. Given that the merkle tree `MT` is computed using collision-resistant a hash function, this occurs with at most negligible probability.

- The honest public key generation algorithm results in key collisions. Obviously, this happens with negligible probability if the encryption and signature schemes are secure.

FACT 5. *Given that the encryption scheme is semantically secure and perfectly correct, and that the signature scheme is secure, then Hybrid 6 is computationally indistinguishable from the ideal simulation to any polynomial-time environment \mathcal{E} .*

F. FORMAL PROOF FOR HAWK

We now prove our main result, Theorem 1 (see Section 4.1.2). Just as we did for private cash in Theorem 3, we will construct an ideal-world simulator \mathcal{S} for every real-world adversary \mathcal{A} , such that no polynomial-time environment \mathcal{E} can distinguish whether it is in the real or ideal world.

F.1 Ideal World Simulator

Our ideal program (`Idealhawk`) and construction (`ContractHawk` and Π_{HAWK}) borrows from our private cash definition and construction in a non-blackbox way (i.e., by duplicating the relevant behaviors). As such, our simulator program `simP` also duplicates the behavior of the simulator from Appendix E.1 involving `mint` and `pour` interactions. Hence we will here explain the behavior involving the additional `freeze`, `compute`, and `finalize` interactions.

Init. Same as in Appendix E.

Simulating corrupted parties. The following messages are sent by the environment \mathcal{E} to the simulator $\mathcal{S}(\text{simP})$ which then forwards it on to both the internally simulated contract `G(ContractHawk)` and the inner simulator `simP`.

- Corrupt party \mathcal{P} submits a transaction $(\text{freeze}, \pi, \text{sn}, \text{cm})$ to the contract. The simulator forwards this transaction to the contract, but also uses the trapdoor τ to extract a witness from π , including $\$val$ and in . The simulator then sends $(\text{freeze}, \$val, \text{in})$ to $\mathcal{F}_{\text{HAWK}}$.
- Corrupt party \mathcal{P} submits a transaction $(\text{compute}, \pi, \text{ct})$ to the contract. The simulator forwards this to the contract and sends `compute` to $\mathcal{F}_{\text{HAWK}}$. The simulator also uses τ to extract a witness from π , including k_i , which is used later. These is stored as `CorruptOpeni := ki`.
- Corrupt party $\mathcal{P}_{\mathcal{M}}$ submits a transaction $(\text{finalize}, \pi, \text{in}_{\mathcal{M}}, \text{out}, \{\text{coin}'_i, \text{ct}_i\})$. The simulator forwards this to the contract, and simply sends $(\text{finalize}, \text{in}_{\mathcal{M}})$ to $\mathcal{F}_{\text{HAWK}}$.

Simulating honest parties. When the environment \mathcal{E} sends inputs to honest parties, the simulator \mathcal{S} needs to simulate messages that corrupted parties receive, from honest parties or from functionalities in the real world. The honest parties will be simulated as below:

- Environment \mathcal{E} gives a `freeze` instruction to party \mathcal{P} . The simulator `simP` receives $(\text{freeze}, \mathcal{P})$ from $\mathcal{F}(\text{Ideal}_{\text{hawk}})$. The simulator does not have any information about the actual committed values for $\$val$ or in . Instead, the simulator create a bogus commitment $\text{cm} := \text{Comm}_s(0 \parallel \perp \parallel \perp)$ that will later be opened (via a false proof) to an arbitrary value. To generate the

serial number sn , the simulator chooses a random element from the codomain of PRF. Finally, the simulator uses τ to generate a forged proof π and sends $(\text{freeze}, \pi, \text{sn}, \text{cm})$ to the contract.

- Environment \mathcal{E} gives a `compute` instruction to party \mathcal{P} . The simulator `simP` receives $(\text{compute}, \mathcal{P})$ from $\mathcal{F}(\text{Ideal}_{\text{hawk}})$. The simulator behaves differently depending on whether or not the manager $\mathcal{P}_{\mathcal{M}}$ is corrupted.

Case 1: $\mathcal{P}_{\mathcal{M}}$ is honest. The simulator does not know values $\$val$ or in . Instead, the simulator samples an encryption randomness r and generates an encryption of 0, $\text{ct} := \text{ENC}(\mathcal{P}_{\mathcal{M}}.\text{epk}, r, 0)$. Finally, the simulator uses the trapdoor τ to create a false proof π that the commitment cm and ciphertext ct are consistent. The simulator then passes $(\text{compute}, \pi, \text{ct})$ to the contract.

Case 2: $\mathcal{P}_{\mathcal{M}}$ is corrupted. Since the manager $\mathcal{P}_{\mathcal{M}}$ in the ideal world would learn $\$val$, in , and k at this point, the simulator learns these values instead. Hence it samples an encryption randomness r and computes a valid encryption $\text{ct} := \text{ENC}(\mathcal{P}_{\mathcal{M}}.\text{epk}, r, (\$val \parallel \text{in} \parallel k))$. The simulator next uses τ to create a proof π attesting that ct is consistent with cm . Finally, the simulator sends $(\text{compute}, \pi, \text{ct})$ to the contract.

- Environment \mathcal{E} gives a `finalize` instruction to party $\mathcal{P}_{\mathcal{M}}$. The simulator `simP` receives $(\text{finalize}, \text{in}_{\mathcal{M}}, \text{out})$ from $\mathcal{F}(\text{Ideal}_{\text{hawk}})$. The simulator generates the output coin'_i for each party \mathcal{P}_i depending on whether \mathcal{P}_i is corrupted or not:

- \mathcal{P}_i is honest: The simulator does not know the correct output value for \mathcal{P}_i , so instead creates a bogus commitment $\text{coin}'_i := \text{Comm}_{s'_i}(0)$ and a bogus ciphertext $\text{ct}'_i := \text{SENC}_{k_i}(s'_i \parallel 0)$ for sampled randomnesses k_i and s'_i .
- \mathcal{P}_i is corrupted: Since the ideal world recipient would receive $\$val'_i$ from $\mathcal{F}(\text{Ideal}_{\text{hawk}})$, the simulator learns the correct value $\$val'_i$ directly. Notice that since \mathcal{P}_i was corrupted, the simulator has access to $k_i := \text{CorruptOpen}_i$, which it extracted earlier. The simulator therefore draws a randomness s'_i , and computes $\text{coin}'_i := \text{Comm}_{s'_i}(\$val'_i)$ and $\text{ct}_i := \text{SENC}_{k_i}(s'_i \parallel \$val'_i)$.

The simulator finally constructs a forged proof π using the trapdoor τ , and then passes $(\text{finalize}, \pi, \text{in}_{\mathcal{M}}, \text{out}, \{\text{coin}'_i, \text{ct}_i\}_{i \in [N]})$ to the contract.

F.2 Indistinguishability of Real and Ideal Worlds

To prove indistinguishability of the real and ideal worlds from the perspective of the environment, we will go through a sequence of hybrid games.

Real world. We start with the real world with a dummy adversary that simply passes messages to and from the environment \mathcal{E} .

Hybrid 1. Hybrid 1 is the same as the real world, except that now the adversary (also referred to as the simulator) will call $(\widehat{\text{crs}}, \tau, \text{ek}) \leftarrow \text{NIZK}.\widehat{\mathcal{K}}(1^\lambda)$ to perform a simulated setup for the NIZK scheme. The simulator will pass the simulated $\widehat{\text{crs}}$ to the environment \mathcal{E} . When an honest party

\mathcal{P} publishes a NIZK proof, the simulator will replace the real proof with a simulated NIZK proof before passing it onto the environment \mathcal{E} . The simulated NIZK proof can be computed by calling the $\text{NIZK}.\widehat{\mathcal{P}}(\widehat{\text{crs}}, \tau, \cdot)$ algorithm which takes only the statement as input but does not require knowledge of a witness.

FACT 6. *It is immediately clear that if the NIZK scheme is computational zero-knowledge, then no polynomial-time environment \mathcal{E} can distinguish Hybrid 1 from the real world except with negligible probability.*

Hybrid 2. The simulator simulates the $\mathcal{G}(\text{ContractHawk})$ functionality. Since all messages to the $\mathcal{G}(\text{ContractHawk})$ functionality are public, simulating the contract functionality is trivial. Therefore, Hybrid 2 is identically distributed as Hybrid 1 from the environment \mathcal{E} 's view.

Hybrid 3. Hybrid 3 is the same as Hybrid 2 except the following changes. When an honest party sends a message to the contract (now simulated by the simulator \mathcal{S}), it will sign the message with a signature verifiable under an honestly generated nym. In Hybrid 3, the simulator will replace all honest parties' nym and generate these nym itself. In this way, the simulator will simulate honest parties' signatures by signing them itself. Hybrid 3 is identically distributed as Hybrid 2 from the environment \mathcal{E} 's view.

Hybrid 4. Hybrid 4 is the same as Hybrid 3 except for the following changes:

- When an honest party \mathcal{P} produces a ciphertext ct_i for a recipient \mathcal{P}_i , and if the recipient is also uncorrupted, then the simulator will replace this ciphertext with an encryption of 0 before passing it onto the environment \mathcal{E} .
- When an honest party \mathcal{P} produces a commitment coin or cm , then the simulator replaces this commitment with a commitment to 0.
- When an honest party \mathcal{P} computes a pseudorandom serial number sn , the simulator replaces this with a randomly chosen value from the codomain of PRF.

FACT 7. *It is immediately clear that if the encryption scheme is semantically secure, if PRF is a pseudorandom function, and if Comm is a perfectly hiding commitment scheme, then no polynomial-time environment \mathcal{E} can distinguish Hybrid 4 from Hybrid 3 except with negligible probability.*

Hybrid 5. Hybrid 5 is the same as Hybrid 4 except for the following changes. Whenever the environment \mathcal{E} passes to the simulator \mathcal{S} a message signed on behalf of an honest party's nym, if the message and signature pair was not among the ones previously passed to the environment \mathcal{E} , then the simulator \mathcal{S} aborts.

FACT 8. *Assume that the signature scheme employed is secure; then the probability of aborting in Hybrid 5 is negligible. Notice that from the environment \mathcal{E} 's view, Hybrid 5 would otherwise be identically distributed as Hybrid 4 modulo aborting.*

Hybrid 6. Hybrid 6 is the same as Hybrid 5 except for the following changes. Whenever the environment passes

(`pour`, π , $\{\text{sn}_i, \mathcal{P}_i, \text{coin}_i, \text{ct}_i\}$) (or (`freeze`, π , `sn`, `cm`)) to the simulator (on behalf of corrupted party \mathcal{P}), if the proof π verifies under `statement`, then the simulator will call the NIZK’s extractor algorithm \mathcal{E} to extract `witness`. If the NIZK π verifies but the extracted `witness` does not satisfy the relation $\mathcal{L}_{\text{POUR}}(\text{statement}, \text{witness})$ (or $\mathcal{L}_{\text{FREEZE}}(\text{statement}, \text{witness})$), then abort the simulation.

FACT 9. Assume that the NIZK is simulation sound extractable, then the probability of aborting in Hybrid 6 is negligible. Notice that from the environment \mathcal{E} ’s view, Hybrid 6 would otherwise be identically distributed as Hybrid 5 modulo aborting.

Finally, observe that Hybrid 6 is computationally indistinguishable from the ideal simulation \mathcal{S} unless one of the following bad events happens:

- A value `val'` decrypted by an honest recipient is different from that extracted by the simulator. However, given that the encryption scheme is perfectly correct, this cannot happen.
- A commitment `coin` is different than any stored in `ContractHawk.coins`, yet it is valid according to the relation $\mathcal{L}_{\text{POUR}}$. Given that the merkle tree `MT` is computed using collision-resistant a hash function, this occurs with at most negligible probability.
- The honest public key generation algorithm results in key collisions. Obviously, this happens with negligible probability if the encryption and signature schemes are secure.

FACT 10. Given that the encryption scheme is semantically secure and perfectly correct, and that the signature scheme is secure, then Hybrid 6 is computationally indistinguishable from the ideal simulation to any polynomial-time environment \mathcal{E} .

G. ADDITIONAL THEORETICAL RESULTS

In this section, we describe additional theoretical results for a more general model that “shares” the role of the (minimally trusted) manager among n designated parties. In contrast to our main construction, where posterior privacy relies on a specific party (the manager) following the protocol, in this section posterior privacy is guaranteed even if a majority of the designated parties follow the protocol. Just as in our main construction, even if *all* the manager parties are corrupted, the correctness of the outputs as well as the security and privacy of the underlying cryptocurrency remains in-tact.

G.1 Financially Fair MPC with Public Deposits

We describe a variant of the financially fair MPC result by Kumaresan et al. [38], reformulated under our formal model. We stress that while Bentov et al. [16], Kumaresan et al. [38], and Kiayias et al. [36] also introduce formal models for cryptocurrency-based secure computation, their models are somewhat restrictive and insufficient for reasoning about general protocols in the blockchain model of secure computation — especially protocols involving pseudonymity, anonymity, or financial privacy, including the protocols described in this paper, Zerocash-like protocols [10], and other protocols of interest [35].

$\text{Ideal}_{\text{sfe}}(\{\mathcal{P}_i\}_{i \in [n]}, \$\text{amt}, f, T_1)$

Deposit: Upon receiving (`deposit`, x_i) from \mathcal{P}_i :

send (`deposit`, \mathcal{P}_i) to the adversary \mathcal{A}
assert $T \leq T_1$ and $\text{ledger}[\mathcal{P}_i] \geq \amt
assert \mathcal{P}_i has not called `deposit` earlier

$\text{ledger}[\mathcal{P}_i] := \text{ledger}[\mathcal{P}_i] - \amt
record that \mathcal{P}_i has called `deposit`

Compute: Upon receiving (`compute`) from \mathcal{P}_i :

send (`compute`, \mathcal{P}_i) to the adversary \mathcal{A}
assert $T \leq T_1$
assert that all parties have called `deposit`
let $(y_1, \dots, y_n) := f(x_1, \dots, x_n)$.
if all honest parties have called `compute`, notify the adversary \mathcal{A} of $\{y_i\}_{i \in \mathcal{K}}$ where \mathcal{K} is the set of corrupt parties

record that \mathcal{P}_i has called `compute`
if all parties have called `compute`:

send each y_i to \mathcal{P}_i
for each party \mathcal{P}_i that deposited: let
 $\text{ledger}[\mathcal{P}_i] := \text{ledger}[\mathcal{P}_i] + \amt .

Timer: Assert $T > T_1$

If not all parties have deposited: for each \mathcal{P}_i that deposited: let $\text{ledger}[\mathcal{P}_i] := \text{ledger}[\mathcal{P}_i] + \amt .
Else, let $\$r := (k \cdot \$\text{amt}) / (n - k)$ where k is the number of parties who did not call `compute`. For each party \mathcal{P}_i that called `compute`: let $\text{ledger}[\mathcal{P}_i] := \text{ledger}[\mathcal{P}_i] + \$\text{amt} + \$r$.

Figure 15: Ideal program for fair secure function evaluation.

Therefore, to facilitate designing and reasoning about the security of general protocols in the blockchain model of secure computation, we propose a new and comprehensive model for blockchain-based secure computation in this paper.

G.1.1 Definitions

Our ideal program for fair secure function evaluation is given in Figure 15. We make the following remarks about this ideal program. First, in a `deposit` phase, parties are required to commit their inputs to the ideal functionality and make deposits of the amount `$amt`. Next, parties send a `compute` command to the ideal functionality. When all honest parties have issued a `compute` command, then the adversary learns the outputs of the corrupt parties. If all parties (including honest and corrupt) have issued a `compute` command, then all parties learn their respective outputs, and the deposits are returned. Finally, if a timeout happens defined by T_1 , the ideal functionality checks to see if all parties have deposited. If not, this means that the computation has not even started. Therefore, simply return the deposits to those who have deposited, and no one needs to be punished. However, if some corrupt parties called `deposit` but did not call `compute`, then these parties’ deposits are redistributed to honest parties.

G.1.2 Construction

ContractSFE($\{\mathcal{P}_i\}_{i \in [n]}, \amt)

Deposit: Upon receiving (**deposit**, $\{\text{com}_j\}_{j \in [n]}$) from \mathcal{P}_i :

assert $T \leq T_1$ and $\text{ledger}[\mathcal{P}_i] \geq \amt
 assert \mathcal{P}_i has not called **deposit** earlier
 $\text{ledger}[\mathcal{P}_i] := \text{ledger}[\mathcal{P}_i] - \amt
 record that \mathcal{P}_i has called **deposit**

Compute: Upon receiving (**compute**, s_i, r_i) from \mathcal{P}_i :

assert $T \leq T_1$
 assert that all \mathcal{P}_i s have deposited, and that they have all deposited the same set $\{\text{com}_j\}_{j \in [n]}$.
 assert that (s_i, r_i) is a valid opening of com_i
 record that \mathcal{P}_i has called **compute**
 if all parties have called **compute**:
 $\text{ledger}[\mathcal{P}_j] := \text{ledger}[\mathcal{P}_j] + \amt for each $j \in [n]$
 reconstruct ρ , send ρ_j to \mathcal{P}_j for each $j \in [n]$

Timer: Assert $T > T_1$

If not all parties have deposited or parties deposited different $\{\text{com}_j\}_{j \in [n]}$ sets:
 For each \mathcal{P}_i that deposited: let $\text{ledger}[\mathcal{P}_i] := \text{ledger}[\mathcal{P}_i] + \amt .
 Else, let $\$r := (k \cdot \$amt)/(n - k)$ where k is the number of parties whose did not send a valid opening. For each party \mathcal{P}_i that sent a valid opening: let $\text{ledger}[\mathcal{P}_i] := \text{ledger}[\mathcal{P}_i] + \$amt + \$r$.

Figure 16: Contract program for fair secure function evaluation.

We now describe how to construct a protocol that realizes the functionality $\mathcal{F}(\text{Ideal}_{\text{sfe}})$ in the most general case.

Our contract construction and user-side protocols are described in Figures 16 and 17 respectively. The protocol is a variant of Bentov et al. [16] and Kumaresan et al. [38], but reformulated under our formal framework. The intuition is that all parties first run an off-chain MPC protocol – at the end of this off-chain protocol, party \mathcal{P}_i obtains \hat{y}_i which is a secret share of its output y_i . The other share needed to recover output y_i is ρ_i , i.e., $y_i := \hat{y}_i \oplus \rho_i$. Denote $\rho := (\rho_1, \dots, \rho_n)$. All parties also obtain random shares of the vector ρ at the end of the off-chain MPC protocol. Then, in an on-chain fair exchange, all parties reconstruct ρ . Here, each party deposits some money, and can only redeem its deposit if it releases its share of ρ . If a party aborts without releasing its share of ρ , its deposit will be redistributed to other honest parties.

THEOREM 4. *Assume that the underlying MPC protocol $\Pi_{\hat{f}}$ is UC-secure against an arbitrary number of corruptions, that the secret sharing scheme is perfectly secret against any $n - 1$ collusions, and that the commitment scheme **commit** is perfectly binding, computationally hiding, and equivocal. Then, the protocols described in Figures 16 and 17 securely emulate $\mathcal{F}(\text{Ideal}_{\text{sfe}})$ in the presence of an arbitrary number of corruptions.*

PROOF. Suppose that $\Pi_{\hat{f}}$ securely emulates the ideal functionality $\mathcal{F}_{\text{SFE}}(\hat{f})$. For the proof, we replace the $\Pi_{\hat{f}}$ in Figure 17 with $\mathcal{F}_{\text{SFE}}(\hat{f})$, and prove the security of the protocol in the $(\mathcal{F}_{\text{SFE}}(\hat{f}), \mathcal{G}(\text{ContractSFE}))$ -hybrid world. We describe

ProtSFE($\{\mathcal{P}_i\}_{i \in [n]}, \amt, f)

Init: Let $\hat{f}(x_1, \dots, x_n)$ be the following function parameterized by f :

pick a random $\rho := (\rho_1, \dots, \rho_n) \in \{0, 1\}^{|y|}$, where each ρ_i is of bit length $|y_i|$
 additively secret share ρ into n shares s_1, \dots, s_n , where each share $s_i \in \{0, 1\}^{|y_i|}$
 for each $i \in [n]$, pick $r_i \in \{0, 1\}^\lambda$, and compute $\text{com}_i := \text{commit}(s_i, r_i)$
 the i -th party's output of \hat{f} is defined as:

$$\text{out}_i := \begin{pmatrix} \hat{y}_i := y_i \oplus \rho_i \\ \text{com}_1, \dots, \text{com}_n \\ s_i, r_i \end{pmatrix}$$

where y_i denotes the i -th coordinate of the output $f(x_1, \dots, x_n)$.

Let $\Pi_{\hat{f}}$ denote an MPC protocol for evaluating the function \hat{f} .

Deposit: Upon receiving the first input of the form (**deposit**, x_i),

assert $T \leq T_1$
 run the protocol $\Pi_{\hat{f}}$ off-chain with input x_i
 when receiving the output out_i from protocol $\Pi_{\hat{f}}$, send (**deposit**, $\{\text{com}_i\}_{i \in [n]}$) to $\mathcal{G}(\text{ContractSFE})$

Compute: Upon receiving the first (**compute**) input,

assert that all parties have deposited, and that they have deposited the same set $\{\text{com}_j\}_{j \in [n]}$ to $\mathcal{G}(\text{ContractSFE})$
 if $T \leq T_1$ and \mathcal{P}_i has not sent any **compute** instruction, then send (**compute**, s_i, r_i) to $\mathcal{G}(\text{ContractSFE})$.
 On receiving ρ_i from $\mathcal{G}(\text{ContractSFE})$, output $\hat{y}_i \oplus \rho_i$

Figure 17: User program for fair secure function evaluation.

the user-defined portion of the simulator program **simP**. The simulator wrapper was described earlier in Figure 14. During the simulation, **simP** will receive a **deposit** instruction from the environment on behalf of corrupt parties. The ideal functionality will also notify the simulator that an honest party has deposited (without disclosing honest parties' inputs). If the simulator has collected **deposit** instructions on behalf of all parties (from both the ideal functionality and environment), at this point the simulator

- Simulates $n - 1$ shares. Among these $|K|$ shares will be assigned to corrupt parties.
- Simulates all commitments $\{\text{com}_i\}_{i \in [n]}$. $n - 1$ of these commitments will be computed honestly from the simulated tokens. The last commitment will be simulated by committing to 0.

Now the simulator collects **compute** instructions from the ideal functionality on behalf of honest parties, and from the environment on behalf of corrupt parties. When the simulator receives a notification (**compute**, s_i, r_i) from the environment on behalf of a corrupt party \mathcal{P}_i , if s_i and r_i are not consistent with what was previously generated by the

simulator, ignore the message. Otherwise, send `compute` to the ideal functionality on behalf of corrupt party \mathcal{P}_i . When the simulator receives a notification $(\text{compute}, \mathcal{P}_i)$ from the ideal functionality for some honest \mathcal{P}_i , unless this is the last honest \mathcal{P}_i , the simulator returns one of the previously generated and unused (s_i, r_i) 's. If this is the last honest \mathcal{P}_i , then the simulator will also get the corrupt parties' outputs $\{y_i\}_{i \in \mathcal{K}}$ from the ideal functionality. At this point, the simulator simulates the last honest party's opening to be consistent with the corrupt parties' outputs – this can be done if the secret sharing scheme is perfectly simulatable (i.e., zero-knowledge) against $n - 1$ collusions and the commitment scheme is equivocable.

It is not hard to see that the environment cannot distinguish between the real world and the ideal world simulation. \square

Optimizations and on-chain costs. Since $\mathcal{F}(\text{Ideal}_{\text{sfe}})$ is simultaneously a generalization of Zerocash [11] and of earlier cryptocurrency-based MPC protocols [16, 36, 38], our construction satisfies the strongest definition so far. However, our construction above requires compiling a generic NIZK prover algorithm with a generic MPC compiler, it is likely slow. Our main construction, `ProtHawk` (see Section 4.1), can be seen as an optimization when $n = 1$ (i.e., the MPC is executed by only a single party). Similarly, the earlier off-chain MPC protocols [16, 36, 38] can be used in place of ours if the user-specified program does not involve any private money.

Even our general construction can be optimized in several ways. One obvious optimization is that not all parties need to send the commitment set $\{\text{com}_j\}_{j \in [n]}$ to the contract. After the first party sends the commitment set, all other parties can simply send a bit to indicate that they agree with the set.

If we adopt this optimization, the on-chain communication and computation cost would be $O(|y| + \lambda)$ per party. In the special case when all parties share the same output, i.e., $y_1 = y_2 = \dots = y_n$, it is not hard to see that the on-chain cost can be reduced to $O(|y_i| + \lambda)$.

If we were to rely on a (programmable) random oracle model, [29] we could further reduce the on-chain cost to $O(\lambda)$ per party (i.e., independent of the total output size). In a nutshell, we could modify the protocol to adopt a ρ of length λ . We then apply a random oracle to expand ρ to $|y|$ bits. Our simulation proof would still go through as long as the simulator can choose the outputs of the random oracle.

G.2 Fair MPC with Private Deposits

The construction above leaks nothing to the public except the *size of the public collateral deposit*. For some applications, even revealing this information may leak unintended details about the application. As an example, an appropriate deposit for a private auction might correspond to the seller's estimate of the item's value. Therefore, we now describe the same task as in Appendix G.1, but with private deposits instead.

7.2.1 Ideal Functionality

Figure 18 defines the ideal program for fair MPC with private deposits, $\text{Ideal}_{\text{sfe-priv}}$. Here, the deposit amount is known to all parties $\{\mathcal{P}_i\}_{i \in [n]}$ participating in the protocol, but it is not revealed to other users of the blockchain. In par-

$\text{Ideal}_{\text{sfe-priv}}(\{\mathcal{P}_i\}_{i \in [N]}, T_1, f)$

Init: Call $\text{Ideal}_{\text{cash}}$.**Init.** Additionally:

ContractCoins: a set of coins and private inputs received by this contract, each of the form $(\mathcal{P}, \text{in}, \$\text{val})$ Initialize $\text{ContractCoins} := \emptyset$

On receiving the first $\$amt$ from some \mathcal{P}_i , notify all parties of $\$amt$

Deposit: Upon receiving $(\text{deposit}, \$\text{val}_i, x_i)$ from \mathcal{P}_i for some $i \in [n]$:

assert $\$val_i \geq \amt and $T \leq T_1$

assert at least one copy of $(\mathcal{P}_i, \$\text{val}_i) \in \text{Coins}$

assert \mathcal{P}_i has not called `deposit` earlier

send $(\text{deposit}, \mathcal{P}_i)$ to \mathcal{A}

add $(\mathcal{P}_i, \$\text{val}_i, \text{in}_i)$ to ContractCoins

remove one $(\mathcal{P}_i, \$\text{val}_i)$ from Coins

record that \mathcal{P}_i has called `deposit`

Compute: Upon receiving `compute` from \mathcal{P}_i for some $i \in [N]$:

send $(\text{compute}, \mathcal{P}_i)$ to \mathcal{A}

assert current time $T \leq T_1$

assert that all parties called `deposit`

Let $(y_1, \dots, y_n) := f(x_1, \dots, x_n)$.

If all honest parties have called `compute`, notify the adversary \mathcal{A} of $\{y_i\}_{i \in \mathcal{K}}$ where \mathcal{K} is the set of corrupt parties.

record that \mathcal{P}_i has called `compute`

If all parties have called `compute`:

Send each y_i to \mathcal{P}_i .

For each party \mathcal{P}_i that deposited: add one $(\mathcal{P}_i, \$\text{val}_i)$ to Coins

Refund: Upon receiving (refund) from \mathcal{P}_i :

notify $(\text{refund}, \mathcal{P}_i)$ to \mathcal{A}

assert $T > T_1$

assert \mathcal{P}_i has not called `refund` earlier

assert \mathcal{P}_i has called `compute`

If not all parties have called `deposit`, add one $(\mathcal{P}_i, \$\text{val}_i)$ to Coins

Else $\$r := (k \cdot \$\text{val}) / (n - k)$ where k is the number of parties who did not call `compute`, and add one $(\mathcal{P}_i, \$\text{val}_i + \$r)$ to Coins

Ideal_{cash}: include $\text{Ideal}_{\text{cash}}$ (Figure 6).

Figure 18: Definition of $\text{Ideal}_{\text{sfe-priv}}$ with private deposit. Notations: ContractCoins denotes frozen coins owned by the contract; Coins denotes the global private coin pool defined by $\text{Ideal}_{\text{cash}}$.

ticular, if all parties behave honestly in the protocol, then the adversary will not learn the deposit amount. Therefore, in the **Init** part of this ideal functionality, some party \mathcal{P}_i sends the deposit amount $\$amt$ to the functionality, and the functionality notifies all parties of $\$amt$. Otherwise, the functionality in Figure 18 is very similar to Figure 15, except that when all of $\{\mathcal{P}_i\}_{i \in [n]}$ are honest, the adversary does not learn the deposit amount.

7.2.2 Protocol

Figures 19 and 20 depict the user-side program and the contract program for fair MPC with private deposits.

At the beginning of the protocol, all parties $\{\mathcal{P}_i\}_{i \in [n]}$ agree on a deposit amount $\$amt$, and cm_0 and publish a commitment to $\$amt$ on the blockchain. As in the case with public deposits, all parties first run an off-chain protocol after which each party \mathcal{P}_i obtains \hat{y}_i . \hat{y}_i is random by itself, and must be combined with another share ρ_i to recover y_i (i.e., the output is recovered as $y_i := \hat{y}_i \oplus \rho_i$). Denote $\rho := (\rho_1, \dots, \rho_n)$. All parties also obtain random shares of the vector ρ at the end of the off-chain MPC protocol. The vector ρ can be reconstructed when parties reveal their shares on the blockchain, such that each party \mathcal{P}_i can obtain its outcome y_i . To ensure fairness, parties make *private* deposits of $\$amt$ to the blockchain, and can only obtain their private deposit back if they reveal their share of ρ to the block chain. The private deposit and private refund protocols make use of commitment schemes and NIZKs in a similar fashion as Zerocash and Hawk.

THEOREM 5. *Assuming that the hash function in the Merkle tree is collision resistant, the commitment scheme Comm is perfectly binding and computationally hiding, the NIZK scheme is computationally zero-knowledge and simulation sound extractable, the encryption scheme ENC is perfectly correct and semantically secure, the PRF scheme PRF is secure, then, our protocols in Figures 19 and 20 securely emulate the ideal functionality $\mathcal{F}(\text{Ideal}_{\text{sfe-priv}})$ in Figure 18.*

PROOF. The proof can be done in a similar manner as that of Theorem 1 (see Appendix F). \square

ProtSFE-Priv($\{\mathcal{P}_i\}_{i \in [n]}, f$)

Init: Same as Figure 17. Additionally, let \mathcal{P} denote the present pseudonym, let crs denote an appropriate common reference string for the NIZK
 If current (pseudonymous) party is P_1 :
 send $(\$amt, r_0)$ to all $\{\mathcal{P}_i\}_{i \in [n]}$
 let $cm_0 := \text{Comm}_{r_0}(\$amt)$, and send (init, cm_0) to $\mathcal{G}(\text{ContractSFE-Priv})$
 Else, on receiving $(\$amt, r_0)$, store $(\$amt, r_0)$
 On receiving (init, cm_0) from $\mathcal{G}(\text{ContractSFE-Priv})$: verify that $cm_0 = \text{Comm}_{r_0}(\$amt)$

Deposit: Upon receiving the first input of the form $(\text{deposit}, \$val, x_i)$: Same as Figure 17. Additionally,
 assert initialization was successful
 assert current time $T < T_1$
 assert this is the first **deposit** input
 let MT be a merkle tree over Contract.Coins
 assert that some entry $(s, \$val, \text{coin}) \in \text{Wallet}$ where $\$val = \amt
 remove one such $(s, \$val, \text{coin})$ from Wallet
 $sn := \text{PRF}_{sk_{\text{prf}}}(\mathcal{P} \parallel \text{coin})$
 let branch be the branch of $(\mathcal{P}, \text{coin})$ in MT
 $\text{statement} := (MT.\text{root}, sn, cm_0)$
 $\text{witness} := (\mathcal{P}, \text{coin}, sk_{\text{prf}}, \text{branch}, s, \$val, r_0)$
 $\pi := \text{NIZK.Prove}(\mathcal{L}_{\text{DEPOSIT}}, \text{statement}, \text{witness})$
 send $(\text{deposit}, \pi, sn)$ to $\mathcal{G}(\text{ContractSFE-Priv})$

Compute: Same as Figure 17

Refund: On input (refund) from the environment,
 if not all parties called **deposit**, $k := 0$
 else $k := (\text{number of parties that aborted})$
 let $\$val' := \$amt + (k \cdot \$amt) / (n - k)$
 pick randomness s
 let $\text{coin} := \text{Comm}_s(\$val')$
 $\text{statement} := (\text{coin}, cm_0, k, n)$
 $\text{witness} := (s, r_0, \$val, \$val')$
 $\pi := \text{NIZK.Prove}(\mathcal{L}_{\text{REFUND}}, \text{statement}, \text{witness})$
 send $(\text{refund}, \pi, \text{coin})$ to $\mathcal{G}(\text{ContractSFE-Priv})$.

Figure 19: User program for fair SFE with private deposit.

ContractSFE-Priv($\{\mathcal{P}_i\}_{i \in [n]}$)

Init: Let crs denote an appropriate common reference string for the NIZK.
On first receiving (init, cm_0) from \mathcal{P}_i for some $i \in [n]$, send cm_0 to all $\{\mathcal{P}_i\}_{i \in [n]}$.

Deposit: On receive ($\text{deposit}, \{\text{com}_j\}_{j \in [n]}, \pi, \text{sn}$) from \mathcal{P}_i :
assert initialization was successful
assert $T \leq T_1$
assert $\text{sn} \notin \text{SpentCoins}$
statement := (MT.root, sn, cm₀)
assert NIZK.Verify($\mathcal{L}_{\text{DEPOSIT}}, \pi, \text{statement}$)
assert \mathcal{P}_i has not called **deposit** earlier
record that \mathcal{P}_i has called **deposit**

Compute: Upon receiving ($\text{compute}, s_i, r_i$) from \mathcal{P}_i :
assert $T \leq T_1$
assert that all \mathcal{P}_i s have deposited, and that they have all deposited the same set $\{\text{com}_j\}_{j \in [n]}$.
assert that (s_i, r_i) is a valid opening of com_i .
record that \mathcal{P}_i has called **compute**

Refund: Upon receiving ($\text{refund}, \pi, \text{coin}$) from \mathcal{P}_i :
assert $T > T_1$
assert \mathcal{P}_i did not call **refund** earlier
assert \mathcal{P}_i called **compute**
if not all parties have deposited or parties deposited different $\{\text{com}_j\}_{j \in [n]}$ sets, $k := 0$
else $k :=$ (number of aborting parties)
statement := (coin, cm₀, k, n)
assert NIZK.Verify($\mathcal{L}_{\text{REFUND}}, \pi, \text{statement}$)
add $(\mathcal{P}_i, \text{coin})$ to Coins

Relation $(\text{statement}, \text{witness}) \in \mathcal{L}_{\text{DEPOSIT}}$ is defined as:
parse **statement** := (MT.root, sn, cm₀)
parse **witness** := (\mathcal{P} , coin, sk_{prf}, branch, $s, \$\text{val}, r_0$)
coin := Comm _{s} (\$val)
cm₀ := Comm _{r_0} (\$val)
assert MerkleBranch(MT.root, branch, ($\mathcal{P}||\text{coin}$))
assert $\mathcal{P}.\text{pk}_{\text{prf}} = \text{sk}_{\text{prf}}(0)$
assert $\text{sn} = \text{PRF}_{\text{sk}_{\text{prf}}}(\mathcal{P}||\text{coin})$

Relation $(\text{statement}, \text{witness}) \in \mathcal{L}_{\text{REFUND}}$ is defined as:
parse **statement** := (coin, cm₀, k, n)
parse **witness** := ($s, r_0, \$\text{val}, \val')
assert **cm₀** := Comm _{r_0} (\$val)
assert $\text{\$val}' := \$\text{val} + (k \cdot \$\text{val}) / (n - k)$
assert **coin** := Comm _{s} (\$val')

Figure 20: Contract program for fair SFE with private deposit.