

Towards Fairness of Cryptocurrency Transactions

Jian Liu¹, Wenting Li¹, Ghassan O. Karame¹, and N. Asokan²

¹Aalto University, Finland

²NEC Laboratories, Germany

Abstract

Motivated by the great success and adoption of Bitcoin, a number of cryptocurrencies such as Litecoin, Dogecoin, and Ethereum are becoming increasingly popular. Although existing blockchain-based cryptocurrency schemes can ensure reasonable security for transactions, they do not consider any notion of fairness. Fair exchange allows two players to exchange digital “items”, such as digital signatures, over insecure networks *fairly*, so that either each player gets the other’s item, or neither player does. Given that blockchain participants typically do not trust each other, enabling fairness in existing cryptocurrencies is an essential but insufficiently explored problem.

In this paper, we explore the solution space for enabling the fair exchange of a cryptocurrency payment for a receipt. We identify the timeliness of an exchange as an important property especially when one of the parties involved in the exchange is resource-constrained. We introduce the notion of *strong timeliness* for a fair exchange protocol and propose two fair payment-for-receipt protocol instantiations that leverage functionality of the blockchain to achieve strong timeliness. We implement both and compare their security and efficiency.

Keywords

Cryptocurrency; Bitcoin; Ethereum; Fair Exchange

I. INTRODUCTION

The phenomenal success of Bitcoin [1] has fueled innovation in a number of application domains such as financial payments, smart contracts, and identity management. Recently, a number of blockchain-based solutions have also been proposed in the Internet of Things (IoT) domain. Device-to-device transactions in IoT infrastructures may well make use of cryptocurrency payments in exchange for information or other items of value, without requiring human participation. In scenarios involving resource-constrained devices, it is important to minimize the computation, communication and energy costs of participating in blockchain systems. For example, many blockchain platforms offer the so-called simplified payment verification (SPV) to interface with such devices.

Although existing blockchain-based cryptocurrency schemes can ensure the security of payments reasonably well, they do not however provide any notion of *fairness*. Given that blockchain participants do not necessarily trust each other, fairness is an especially important property in

existing cryptocurrency exchanges [2]. For instance, consider the example where a payer \mathcal{A} makes a payment to a payee \mathcal{B} in return for an expected good (digital or physical) or a service. It is unfair towards \mathcal{A} if her expectation is not met after \mathcal{B} receives the payment. On the other hand, it is unfair towards \mathcal{B} if he does provide the service but \mathcal{A} later cancels or double-spends the payment. Namely, a fair payment scheme should ensure that \mathcal{B} receives the payment if and only if \mathcal{A} 's expectations are met and vice versa. We can model this as a fair exchange of *payment for receipt* where the receipt is a digital signature, and it can act as a proxy for a physical/digital good or real-world service.

While there is a wealth of literature on fair exchange in a general setting [3], [4], little attention has been paid to the problem of fair exchange involving cryptocurrencies [2]. In this paper, we explore the solution space to achieve fair payment-for-receipt for cryptocurrencies. More specifically, we analyze how well known fair exchange techniques can be adapted for use with existing cryptocurrencies, in particular by leveraging functionality from the blockchain. We propose two such protocols that and analyze/compare their provisions.

In summary, our contributions in this paper are as follows:

- We introduce the notion of *strong timeliness* for a fair exchange protocol (cf. Section II-B). Completing the exchanges in a timely, but fair, manner is an important consideration for fair exchange, especially for resource-constrained devices.
- We propose *two fair payment-for-receipt protocols for cryptocurrency payments* (cf. Sections IV-B and IV-C) that leverage functionality from the blockchain to meet both fairness and strong timeliness. Additionally, we compare these protocols to the proposal by Heilmann *et al.* [5].
- Finally, we implement and evaluate the performance and costs of our constructions (cf. Section V).

II. BACKGROUND

A. Blockchain and Ethereum

The notion of blockchain was originally introduced by the well-known proof-of-work hash-based mechanism that *confirms* cryptocurrency payments in Bitcoin [1]. Bitcoin payments are performed by issuing transactions that transfer Bitcoin coins from the payer to the payee. These entities are called “peers”, and are referenced in each transaction by means of pseudonyms, denoted by Bitcoin addresses. Each address maps to a unique public/private key pair; these keys are used to transfer the ownership of coins among addresses. Miners are entities that participate in the generation of Bitcoin blocks. These blocks are generated by solving a hash-based proof-of-work (PoW) scheme; more specifically, miners must find a nonce value that, when hashed with additional fields (e.g., the Merkle hash of all valid transactions, the hash of the previous block), the result is below a given target value. If such a nonce is found, miners then include it in a new block thus allowing any entity to verify the PoW. Since each block links to the previously generated block, the Bitcoin *blockchain* grows upon the generation of a new block in the network.

As such, the PoW-based blockchain ensures that all transactions and their order of execution are available to all blockchain nodes, and can be verified by these entities. Consensus by the majority of participating miners is required for every transaction exchanged in the system. This inherently prevents double-spending attacks (where the payer attempts to spend the same coin twice), and ensures the correctness of all transactions confirmed in the blockchain as long as the majority of the network is honest.

To ensure that a payment in a cryptocurrency transaction is definitive, a payee needs to wait until a sufficient number of new blocks have been appended to the block that contains the particular transaction so as to minimize the probability that the block is not part of the eventual consensus. In Bitcoin this may take up to an hour. In some situations (e.g., low-value transaction), a payee may be willing to accept a transaction as soon as it is broadcast to the network. This is referred to as *zero confirmation* transaction, which are fast but carry a risk of payment reversal.

Bitcoin's blockchain fueled innovation, and a number of innovative applications have already been devised by exploiting the secure and distributed provisions of the underlying blockchain. Prominent applications include secure timestamping [6], timed commitment schemes [7], secure multiparty computations [8], and smart contracts [9].

Smart contracts refer to binding contracts between two or more parties that are enforced in a decentralized manner by the blockchain nodes without the need for a centralized enforcer. Smart contracts typically consist of self-contained code that is executed by all blockchain nodes. For example, Ethereum [9] is a decentralized platform that enables the execution of arbitrary applications (or contracts) on its blockchain. Owing to its support for a Turing-complete language, Ethereum offers an easy means for developers to deploy their distributed applications in the form of smart contracts. Ethereum additionally offers its own cryptocurrency *Ether* which is also used as the main fuel to execute the contracts and send transactions. Ether payments are commonly used to cover the costs related to contract execution; these costs are measured by the amount of *gas* they consume.

B. Properties of Fair Exchange

A two-party exchange usually involves two players who exchange items between themselves. Each player holds an *item* that it wants to contribute to the exchange and an *expectation* about the other player's item it wants to receive in exchange. Fair exchange is executed between parties that do not trust each other; examples include commercial scenarios such as payment-for-receipt, online purchase, digital contract signing, certified mail, and electronic barter. A fair exchange protocol must ensure that a malicious player cannot gain any advantage over an honest player. More specifically, it should satisfy the following requirements [4]:

- **Effectiveness:** If both players behave correctly, the protocol will conclude with a successful exchange.
- **Fairness:** There are two possible notions of fairness, namely:
 - **Strong Fairness:** When the protocol has completed, either each player receives the item it expects or neither player receives any additional information about each other's item than they already knew.

- **Weak Fairness:** In situations where strong fairness cannot be achieved, a player can prove to an (external) arbiter that the other player has received (or can still receive) the item the latter expects, without any further intervention from the prover.
- **Timeliness:** Both players can be certain that the protocol will be completed at a certain finite point in time. At completion of the protocol, the state of the exchange is final from that player’s perspective.
- **Non-invasiveness:** The protocol should allow the exchange of arbitrary items without making any demands on its structure, i.e., the fair exchange protocol does not itself impose any requirement on the form of the items being exchanged.¹
- **Transaction duration:** The time taken for the exchange should be short.

The timeliness requirement defines a fixed point in time at which the protocol will be completed. This property aims to avoid the case where one player in the exchange has to wait indefinitely for the other player to take an action that will determine how the exchange will be concluded (successfully or otherwise). Timeliness is particularly important for resource-constrained IoT devices which cannot afford to be online for long stretches of time or poll indefinitely. One way to achieve the timeliness requirement as stated in [4] is to agree on a pre-defined timeout. This is typically a challenging task, since it is difficult to predefine an optimal time point at which the protocol should be completed: a short timeout will result in the exchange failing even when both parties are honest (thus harming the effectiveness requirement), whereas a long timeout is unacceptable for resource-constrained devices with limited battery or bandwidth. Ideally, the notion of timeliness should capture the possibility that either player can decide to conclude the exchange *at any point* during the exchange without having to depend on the actions of the other player. To remedy this, we therefore define a new notion of timeliness, dubbed *strong timeliness*, as follows:

- **Strong timeliness:** An honest player can, *any point* in time, choose to complete the protocol. At completion, the state of the exchange is final from that player’s perspective.

C. Fair Cryptocurrency Payments

In this paper, we consider the “payment-for-receipt” scenario $\rho \rightarrow \sigma$, where an entity, \mathcal{A} , makes a digital payment ρ to another entity, \mathcal{B} , in order to get a receipt for the payment in the form of a digital signature σ . *Our goal is to explore the solution space for integrating a fair exchange of payment-for-receipt into existing cryptocurrency payment schemes.* (Hereafter referred to as “fair payments” for the sake of brevity.)

III. FAIR PAYMENTS WITH FIXED TIMEOUTS

We start by describing a solution adapted from Heilman et al. [5] that enables timelock-based fair payments and relies on an intermediary to improve the anonymity of cyptocurrency

¹For example, a fair exchange scheme to exchange signatures in any standard digital signature algorithm, such as RSA or ECDSA, is considered non-invasive. However, if a scheme requires anyone who wants to verify the exchanged signatures to access and perform some check on the blockchain, then that scheme is *invasive*.

payments. Here, a user \mathcal{A} wants to fairly exchange a cryptocurrency payment for a voucher from an intermediary \mathcal{B} . The voucher is in fact a blind signature; \mathcal{A} will unblind the voucher and send it privately to an anonymous payee who can exchange it with \mathcal{B} in such a way that \mathcal{B} cannot link \mathcal{A} and the payee [5].

This is achieved without relying on any external entity through the use of smart contracts and using a fixed, predefined, timeout to implement a timely $\rho \rightarrow \sigma$ fair exchange. First, \mathcal{A} generates a transaction contract T_ρ that enables her to pay a pre-defined amount to \mathcal{B} under the condition that \mathcal{B} must publish a valid signature on a message m within a time limit tw ; The output of T_ρ will become an input in one of the following two blockchain transactions:

- 1) A transaction which is signed by \mathcal{B} and contains a valid signature σ on m ; (i.e., exchange is successful and fair).
- 2) A transaction which is signed by \mathcal{A} and the time window tw has expired. (i.e., exchange fails and the money reverts to \mathcal{A}).

The contract is fulfilled if \mathcal{B} posts a transaction T_σ that contains a valid signature σ on m , and the promised payment is transferred from \mathcal{A} to \mathcal{B} . If \mathcal{B} does not publish a receipt within the time window tw , then \mathcal{A} can sign and post a transaction that returns the promised payment amount back to herself. All transactions are broadcast to the blockchain network, thus allowing all blockchain nodes to verify whether the payment conditions have been met, and reach consensus on the state of the exchange.

Analysis: This protocol ensures a fair exchange between \mathcal{A} and \mathcal{B} , i.e., prevents \mathcal{A} from double-spending her payment, and enables \mathcal{B} to spend \mathcal{A} 's payment only if \mathcal{B} has published his signature. We now analyze this protocol in relation to the requirements from Section II-B:

- **Effectiveness:** If the the timeout is too short, there may be not enough time for T_σ to be confirmed in the blockchain. Namely, the miners will refuse to confirm T_σ after the timeout has passed. In this case, the effectiveness of the fair exchange cannot be guaranteed since the exchange fails because of the timeout even when both parties are honest.
- **Fairness:** The protocol does not ensure strong fairness since it is possible that the timeout is reached after \mathcal{B} broadcasts T_σ , but before T_σ is confirmed in the blockchain. For example, the adversary may mount a denial-of-service attack against \mathcal{B} to throttle its network connectivity. In this case, \mathcal{A} might receive σ without \mathcal{B} receiving ρ . However, the protocol satisfies weak fairness, since \mathcal{B} can prove to an (external) arbiter that he did indeed broadcast T_σ , and as a result, σ has been revealed to the public.²
- **Timeliness:** This protocol satisfies weak timeliness but not strong timeliness, since once the transaction T_σ is confirmed, \mathcal{A} cannot decide to complete the exchange any sooner than the specific timeout.
- **Non-invasiveness:** The protocol is non-invasive because it does not impose any specific structure on σ .
- **Transaction duration:** \mathcal{B} needs to wait for a period t long enough blocks to be appended to the blockchain after its own transaction appears there. \mathcal{A} has to wait at most till the timeout tw (t depends on the blockchain system; $tw > t$ for ensuring effectiveness).

²Note that this renders the anonymous payment scheme of [5] insecure.

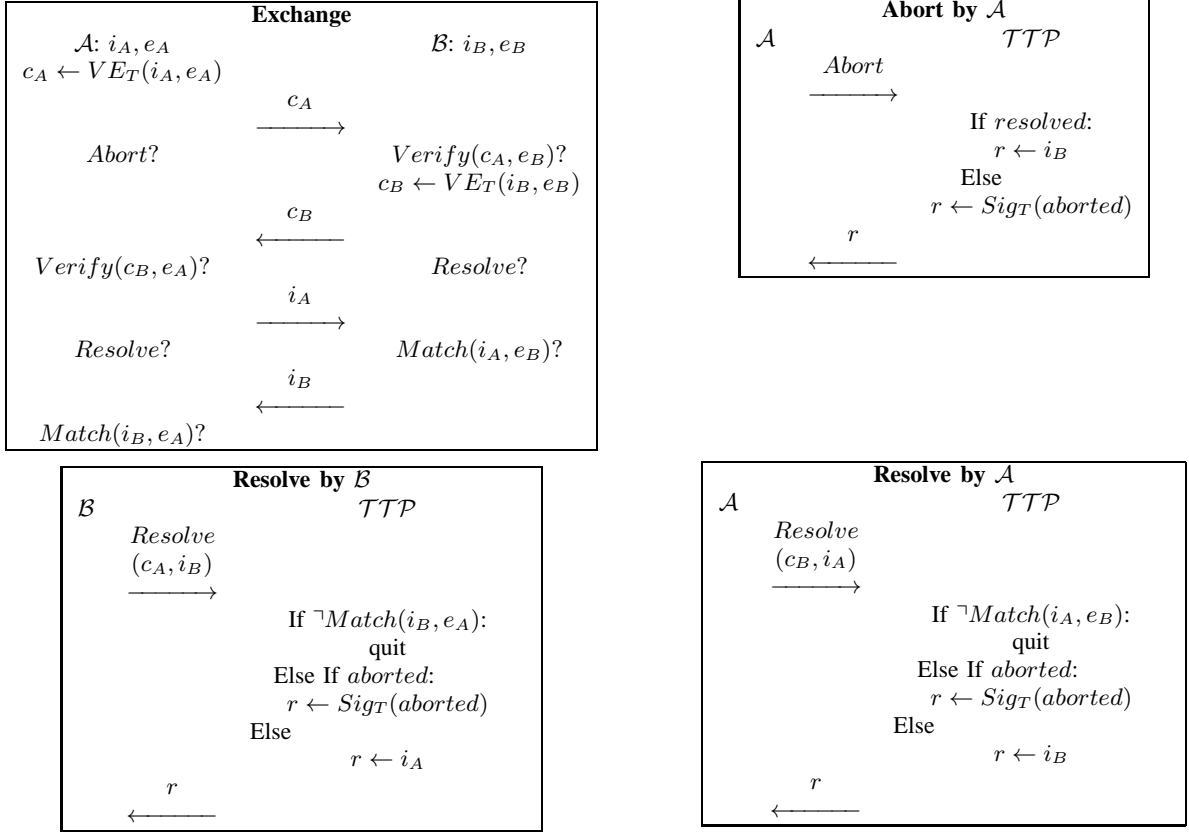


Fig. 1: Optimistic fair exchange.

IV. OPTIMISTIC FAIR PAYMENTS

We now introduce an alternative scheme for fair payments based on optimistic fair exchange.

A. Optimistic Fair Exchange

Optimistic fair exchange (OFE) protocols were first proposed by Asokan et al. [3], [4]; their protocol relies on the presence of a trusted third party (TTP) but only in an *optimistic* fashion: TTP is only required when one player attempts to cheat or simply crashes. In the common case where A and B are honest TTP need not be involved.

Optimistic fair exchange consists of an exchange protocol (protocol exchange) and two recovery protocols (protocol abort and protocol resolve). Figure 1 details the protocols. First, both players agree on what needs to be exchanged and which third party to use in case of an exception. Such an “agreement” is informal: it has no validity outside the context of the protocol. Then, one player (e.g., A) sends her item (i_A) and her expectation about B ’s item (e_A) encrypted under a verifiable encryption scheme (c_A), which enables any entity to verify whether i_A is a valid

signature by \mathcal{A} (without the need for decrypting the message); and can be decrypted only by the trusted third party \mathcal{TTP} . \mathcal{B} first verifies the validity of i_A , constructs an encryption c_B of (i_B, e_B) , and decides similarly whether to send it to \mathcal{A} .

If \mathcal{B} does not send c_B , \mathcal{A} can abort the protocol at any point in time by initiating protocol abort with \mathcal{TTP} which issues an abort token. In this case, the exchange is unsuccessful but fair: neither player receives any additional information about each other’s item. If \mathcal{B} sends c_B , and \mathcal{A} has not decided to abort, she verifies the validity of i_B and decides whether to send i_A to \mathcal{B} . While waiting for i_A , \mathcal{B} can initiate protocol resolve at any time by sending (c_A, i_B) to \mathcal{TTP} . \mathcal{TTP} will decrypt c_A to get (i_A, e_A) and return i_A to \mathcal{B} if i_B meets e_A and it has not previously issued an *abort* token for this particular exchange. If a transaction was previously aborted, \mathcal{TTP} will not agree to resolve it. Similarly it will not agree to abort a transaction that had already been resolved. \mathcal{A} can run resolve in the same way while waiting for i_B from \mathcal{B} . This is a general fair exchange protocol that can support “items” in the form signatures in standard signature schemes. It requires \mathcal{TTP} to keep state for every aborted or resolved transaction.

B. Blockchain-based OFE with a Stateless \mathcal{TTP}

We now extend the above OFE protocol by making use of a blockchain to avoid the need for \mathcal{TTP} to maintain state. Let i_A be a signature ρ corresponding to a payment message in a cryptocurrency scheme. First, we consider zero-confirm payments: the two players exchange ρ for a receipt σ but do not wait for $\rho(i_A)$ to be confirmed in the blockchain. \mathcal{A} can abort the exchange by publishing an abort transaction T_{abort} to the blockchain instead of sending an abort message to \mathcal{TTP} . Thus \mathcal{TTP} only needs to support protocol resolve. It does not need to keep any state w.r.t. the protocol execution since all needed state information is recorded in the blockchain. We implemented this variant of OFE using Ethereum’s smart contracts as shown in Figure 2. Note that when \mathcal{TTP} recovers item i_A in response to a resolve request from \mathcal{B} , it needs to save \mathcal{B} ’s item i_B so that any subsequent abort from \mathcal{A} can be answered correctly by the smart contract without violating \mathcal{A} ’s fairness. Therefore, \mathcal{TTP} will ask the smart contract to save i_B during \mathcal{B} ’s invocation of resolve.

We now analyze this protocol w.r.t. the properties listed in Section II-B:

- **Effectiveness:** Effectiveness is guaranteed in this case since i_A will be eventually confirmed if both players behave correctly.
- **Fairness:** Once the blockchain-based OFE protocol completes, \mathcal{A} receives $\sigma(i_B)$ and \mathcal{B} receives ρ . However, \mathcal{A} can double-spend the money associated with ρ after the completion of the optimistic fair exchange, thus invalidating strong fairness. \mathcal{B} can however prove this misbehavior to an arbiter by showing ρ . So, this scheme only satisfies the weak fairness property.
- **Timeliness:** Strong timeliness is inherited from classical OFE: either player can invoke protocol resolve at any point if they have received the other player’s verifiable encryption (c_A or c_B). \mathcal{A} can attempt to publish T_{abort} at any time. In all cases, the protocol is guaranteed to terminate fairly.

```

function abort(exchange id  $id_{ex}$ )
  if entry of  $id_{ex}$  exists then
    if sender is the originator and the retrieved entry is a resolved item  $i_B$  then
      return  $i_B$  to the sender
    end if
  else
    add an entry of  $id_{ex}$  with an abort token
  end if
end function

function resolve(exchange id  $id_{ex}$ , optional resolved item  $i_B$ )
  if sender is  $\mathcal{TTP}$  then
    if entry of  $id_{ex}$  exists and the retrieved entry is an abort token then
      return aborted
    else
      add an entry of  $id_{ex}$  with the optional resolved item  $i_B$ 
      return  $\neg$  aborted
    end if
  end if
end function

```

Fig. 2: Smart contract for Blockchain-based OFE to assist abort and resolve procedures in order to keep \mathcal{TTP} stateless.

- **Non-invasiveness:** The signature σ is non-invasive since it can be any signature in any form.
- **Transaction duration:** Since the transactions are zero confirmation, they can complete fast during optimistic execution (no need to wait any blocks confirmed on the blockchain).

This variant can be upgraded from zero confirmation to full confirmation. After getting i_A from \mathcal{A} , \mathcal{B} broadcasts it and waiting for it to be confirmed on the blockchain, then sends i_B to Alice. When \mathcal{TTP} resolves for \mathcal{B} , it just broadcasts i_A . When it resolves for \mathcal{A} , it return i_B after i_A being confirmed on the blockchain. If i_A being double spent before being confirmed, it means \mathcal{A} aborts the protocol. This full confirmation variant achieves strong fairness but at the expense of longer transaction duration.

C. Invasiveness vs. Avoiding Trusted Third Parties

We now describe a variant that dispenses with the need for \mathcal{TTP} altogether but at the expense of making σ invasive. Our proposal unfolds as follows. \mathcal{A} first constructs the message m to be signed, and uses m to create a transaction contract T_ρ with an output of some amount of digital money that is spendable in one of the following two transactions:


```

function initExchange(payment  $T_{pay}[v]$ , expected item  $m$  and recipient)
  if state is UNINITIALIZED and contract has received the payment with value  $v$  then
    record originator, recipient and  $m$ 
    switch state to INITIALIZED
  end if
end function
function abort
  if state is INITIALIZED and the message is sent by the originator then
    refund  $v$  to the originator
    clear up storage and switch state to UNINITIALIZED
  end if
end function
function resolve(signature on  $m$ )
  if state is INITIALIZED and the message is sent by the recipient then
    if signature on  $m$  is valid then
      send  $v$  to recipient and the signature to originator
      clear up storage and switch state to UNINITIALIZED
    end if
  end if
end function

```

Fig. 3: Smart contract for fair payment of blockchain-based signature.

- 1) T_σ which is signed by \mathcal{B} and contains a valid signature σ on m ;
- 2) T_{abort} which is signed by \mathcal{A} .

Notice that there are no time window/constraints in T_ρ , and it can trigger either T_σ or T_{abort} , depending on which one is confirmed in the blockchain. Recall that if both T_σ and T_{abort} are broadcast, only one of them will eventually be confirmed (since they conflict with one another). This protocol is invasive since \mathcal{B} 's signature σ is only valid if it is stored on the blockchain. Namely, a verifier must not only check that σ is (cryptographically) valid, but also that it is confirmed in the blockchain.

This protocol can be fully deployed as an Ethereum smart contract without the need for \mathcal{TTP} . In this case, \mathcal{A} will first send a deposit to the contract via T_ρ ; the contract will forward the deposit either to \mathcal{B} or back to \mathcal{A} , depending on whether it receives T_σ or T_{abort} first (respectively). An example of such contract functions is sketched in Figure 3.

Our extension ensures the following properties:

- **Effectiveness:** If both players behave correctly, \mathcal{B} will receive the payment by publishing a T_σ transaction, and \mathcal{A} will obtain her desired receipt when T_σ has been confirmed.
- **Fairness:** \mathcal{B} can only receive the payment when T_σ is confirmed. Similarly, once T_σ is confirmed, \mathcal{A} has a valid receipt. At least one of T_σ or T_{abort} will be confirmed.

- **Timeliness:** The protocol terminates after either T_σ or T_{abort} is confirmed. \mathcal{A} can choose to wait for T_σ to be confirmed or issue T_σ . In the former case, the exchange is successfully completed. In the latter case, \mathcal{A} cannot gain any advantage since the signature σ is valid only if it is confirmed in the blockchain. Similarly \mathcal{B} can either issue T_σ and wait for it to be confirmed or simply walk away. In either case, the state of the exchange is final.
- **Non-invasiveness:** Clearly, σ is invasive since it is only valid when it is confirmed in the blockchain.
- **Transaction duration:** The transaction duration is long because both parties need to wait for either T_σ or T_{abort} to be confirmed in the blockchain.

V. EXPERIMENTAL EVALUATION

We now describe and evaluate our Ethereum-based implementation of the contracts depicted in Figure 2 (Stateless \mathcal{TTP}) and Figure 3 (Invasive Signature).

A. Implementation Setup

To implement the contracts outlined in Figures 2 and 3, we assigned an Ethereum node to each entity (e.g., \mathcal{A} , \mathcal{B} , \mathcal{TTP}). These nodes are connected to a private Ethereum network (that is equipped with private mining functionality) with a bandwidth limit of 100Mbps. We deployed the mining node and \mathcal{TTP} on two servers both with 24-core Intel Xeon E5-2640 and 32GB of RAM. In our testbed, \mathcal{A} and \mathcal{B} reside on two machines equipped with 4-core Intel i5-6500 with 8GB of RAM and 8-core Intel Xeon E3-1230 with 16GB of RAM, respectively. In our implementation, these entities prepare and send the transactions to the blockchain using the Javascript library web3.js. This library interfaces the Ethereum nodes through its RPC calls. In the Stateless \mathcal{TTP} instantiation, we implement OFE computation and communication using GoLang and C. We used as the ECDSA signature scheme, which is directly supported by Ethereum contracts. To implement verifiable encryption, we use on the scheme in [4] implemented with cryptographic library GMP [10] in C. We preset and fix the difficulty of our private Ethereum testnet in the code and the genesis block so that the block generation time is around 5 seconds.

In our experiments, we measured the gas and time consumption for the following procedures: deploy, optimistic completion, abort, and \mathcal{TTP} resolve. Deploy refers to deploying the smart contract into the blockchain. In Stateless \mathcal{TTP} , the smart contract should be deployed by \mathcal{TTP} as the contract is managing its state. Optimistic completion refers to the successful completion of the exchange without invoking resolve or abort. We measured the cost for \mathcal{A} and \mathcal{B} separately in the case of the Invasive Signature variant (Section IV-C), as the protocol is asynchronous. Finally, the contract is triggered by \mathcal{A} for abort and by \mathcal{TTP} for resolve. We only consider the resolve protocol under the assumption that the exchange has not been aborted.

To measure gas consumption, we observe the difference in the account balance before and after invoking the contract, and we convert this amount to the amount of gas according to our fixed gas price. We measure the execution time starting from the initial contract invocation until the entities involved in the fair exchange protocol receive the relevant notifications from the

protocols		Stateless TTP (Section IV-B)	Invasive Signature (Section IV-C)
actions			
deploy		537,783	645,900
optimistic completion	\mathcal{A}	0	126,457
	\mathcal{B}		27,935
abort		67,574	33,746
TTP resolve		132,600	–

TABLE I: Gas consumption in Ethereum contracts for fair payment protocols.

Ethereum network through its event mechanism. In our evaluation, each time measurement is averaged over 10 independent executions of the fair exchange protocol; where appropriate, we also report the corresponding 95% confidence intervals.

B. Evaluation Results

1) *Gas consumption:* Our evaluation results are shown in Table I. We first observe that for both contracts, contract deployment consumes the most amount of gas since the process of creating contracts and storing data in the blockchain are one of the most expensive operations in Ethereum [11]. As described earlier, the Stateless TTP variant of OFE does not need to involve the blockchain at all during exchanges completed optimistically. Therefore the gas consumption for an optimistically concluded fair payment is zero. We contrast this with the Invasive Signature OFE variant (Section IV-C), which requires 126,457 gas from \mathcal{A} in order to initiate the exchange protocol and 27,935 gas from \mathcal{B} to complete it. The large overhead incurred on \mathcal{A} here is mainly caused by storing exchange contract parameters in the blockchain during contract initialization. Notice that abort requires considerably more gas in the Stateless TTP protocol when compared to the Invasive Signature (Section IV-C). This is due to the fact that the contract may spend more gas in order to transmit the previous resolved item (if any). Similarly, TTP resolve potentially needs to store resolved items in the contract (cf. Section IV-B) – which incurs additional gas consumption.

2) *Time consumption:* Recall that this refers to the time elapsed between the initial contract invocation until the entities involved in the fair exchange protocol receive the relevant notifications from the Ethereum network. Our findings are illustrated in Table II.

We observe that the contract invocation process is rather time-consuming; for instance, the protocol initialization procedure by \mathcal{A} in Invasive Signature (Section IV-C) consumes around 4 seconds. We contrast this with 277 milliseconds which is required for the completion of the Stateless TTP protocol. The latter is almost 14 times faster in spite of the reliance on verifiable encryption, due to the fact that the blockchain needs to generate a block in order to include the transactions. Recall that the average block generation time in our private Ethereum network is tuned to be around 5 seconds.

The time execution of the remaining operations is comparable in both protocols, which is around 4 seconds. This value largely depends on block generation time of the blockchain network. Notice that the width of the confidence interval corresponds to the variation of block generation times exhibited in Ethereum.

protocols		Stateless TTP (Section IV-B)	Invasive Signature (Section IV-C)
optimistic completion	A	277.0 \pm 9.2	3,831.0 \pm 973.2
	B		4,195.8 \pm 1,077.3
abort		3,432.4 \pm 1,000.3	4,301.5 \pm 1,305.6
TTP resolve		3,773 \pm 902.5	-

TABLE II: Time in milliseconds with 95% confidence interval to perform each protocol action.

Requirements	Types	Heilman et al. [5] (Section III)	Stateless TTP (Section IV-B)		Invasive signature (Section IV-C)
			zero-confirm	full-confirm	
Fairness		weak	weak	strong	strong
Timeliness		weak	strong	strong	sting
Effectiveness		?	✓	✓	✓
Non-invasiveness		✓	✓	✓	X
No TTP		✓	X	X	✓
Duration of transaction		long	short	long	long

TABLE III: Comparisons of fair payments protocols.

3) *Summary:* Given our findings, we conclude that the fair payment protocol based on Stateless TTP is more cost- and time-effective than its counterpart based on Invasive Signatures (Section IV-C) when the protocol is executed without exceptions.

In the case where an exception occurs, both protocols incur comparable costs and time overhead. Namely, since any transaction can only take effect once they are confirmed in the blockchain (i.e., every 12 seconds in the Ethereum public blockchain), the reliance on the blockchain in abort and resolve protocol incurs considerable time delays.

VI. COMPARISON AND OUTLOOK

In this paper, we explored the solution space to realize fair exchange within cryptocurrencies. To this end, we proposed two fair payment-for-receipt protocols for cryptocurrency payments (cf. Sections IV-B and IV-C) that leverage functionality from the blockchain to meet both fairness and strong timeliness. A systematic comparison between our proposals is shown in Table III.

Our findings suggest that the first scheme cannot satisfy the strong timeliness property, and as such can only guarantee weak fairness. Furthermore, choosing a short-timeout here can harm the effectiveness of this construct. In this respect, the constructs based on the Stateless TTP option (Section IV-B) and the blockchain-based invasive signature (Section IV-C) establish the strongest tradeoffs between performance and provisions. Namely, our performance evaluation shows that the Stateless TTP option is more efficient when the exchange concludes optimistically. As such, the Stateless TTP option seems to be ideal in those scenarios where only weak fairness is sufficient or when non-invasiveness is required. Otherwise, we recommend the reliance on the blockchain-based Invasive Signature option.

None of the variants we examined can simultaneously possess all desirable properties (strong fairness, strong timeliness, effectiveness, non-invasiveness, no use of TTP and short transaction

duration). Designing a fair payment mechanism that does possess all of these properties remains an open question. A potential option is to leverage the recent development of *witness encryption* [12], which is associated with a specific NP-relation $(x, w) \in R$, where x is a “statement” and w is a “witness”. It allows the sender to encrypt a message m with respect to statement x as $c \leftarrow Enc_R(x, m)$. Any receiver who knows the witness w that satisfies $(x, w) \in R$ can decrypt this ciphertext c as $m = Dec_R(c, w)$. So, we can improve the Invasive Signature (Section IV-C) approach by having \mathcal{B} broadcast a non-invasive signature encrypted under the witness encryption, and the witness for decryption is that the blockchain contains T_{σ_i} but does not contain T_{abort_i} . We did not implement this construction since all known constructions of witness encryption are penalizing in terms of performance.

Privacy has not been considered as a requirement for fair exchange protocols. However, there are a number of scenarios where privacy considerations play a paramount role. For example, m may contain some important information about \mathcal{A} that cannot be revealed. In all of the four constructions, the contents of the signature can be seen by the public; there are, however, a number of techniques that can be used to protect the contents of signatures. For instance, one can improve the stateless TTP (Section IV-B) based protocol (Figure II) by having TTP send a verifiable encryption of i_B to the resolve contract—thus preserving the privacy of \mathcal{A} .

In conclusion, while there are no “bullet-proof” solutions that simultaneously achieve all desirable properties discussed above we observe that a number of trade-offs exist within the solution space to sacrifice one property in order to achieve the rest. Depending on the application scope, this might already offer a differentiator and stronger value proposition for existing cryptocurrencies. We therefore hope that our findings motivate further research in this largely unexplored area.

REFERENCES

- [1] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system, 2009. <http://www.bitcoin.org/bitcoin.pdf>.
- [2] D. Jayasinghe, K. Markantonakis, and K. Mayes. Optimistic fair-exchange with anonymity for bitcoin users. In *e-Business Engineering (ICEBE), 2014 IEEE 11th International Conference on*, pages 44–51, Nov 2014.
- [3] N. Asokan, Victor Shoup, and Michael Waidner. Optimistic fair exchange of digital signatures. *IEEE Journal on Selected Areas in Communications*, 18(4):593–610, 2000. <http://dx.doi.org/10.1109/49.839935>.
- [4] N. Asokan. *Fairness in Electronic Commerce*. PhD thesis, University of Waterloo, 1998. <https://uwspace.uwaterloo.ca/bitstream/handle/10012/292/NQ32811.pdf>.
- [5] Ethan Heilman, Foteini Baldimtsi, and Sharon Goldberg. Blindly signed contracts: Anonymous on-blockchain and off-blockchain bitcoin transactions. In Jeremy Clark et al., editors, *Financial Cryptography and Data Security - FC 2016 International Workshops, BITCOIN, VOTING, and WAHC, Christ Church, Barbados, February 26, 2016, Revised Selected Papers*, pages 43–60, 2016. <http://eprint.iacr.org/2016/056>.
- [6] Bela Gipp, Norman Meuschke, and André Gernandt. Decentralized trusted timestamping using the crypto currency bitcoin. *CoRR*, abs/1502.04015, 2015.
- [7] Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and Lukasz Mazurek. Secure multiparty computations on bitcoin. In *Proceedings of the 2014 IEEE Symposium on Security and Privacy*, SP ’14, pages 443–458, Washington, DC, USA, 2014. IEEE Computer Society.
- [8] Iddo Bentov and Ranjit Kumaresan. How to use bitcoin to design fair protocols. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology – CRYPTO 2014: 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part II*, pages 421–439, Berlin, Heidelberg, 2014. Springer Berlin Heidelberg.

- [9] Vitalik. Buterin. A next-generation smart contract and decentralized application platform, 2014. <https://github.com/ethereum/wiki/wiki/White-Paper>.
- [10] GMP. The gnu multiple precision arithmetic library. <https://gmplib.org/>. Accessed: 2016-07-26.
- [11] Wood Gavin. Ethereum: A secure decentralised generalised transaction ledger., 2014. <http://bitcoinaffiliatelist.com/wp-content/uploads/ethereum.pdf>.
- [12] Sanjam Garg, Craig Gentry, Amit Sahai, and Brent Waters. Witness encryption and its applications. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing, STOC '13*, pages 467–476, New York, NY, USA, 2013. ACM.