

On Payment Channels in Asynchronous Money Transfer Systems

Oded Naor

Technion, Haifa, Israel

Idit Keidar

Technion, Haifa, Israel

Abstract

Money transfer is an abstraction that realizes the core of cryptocurrencies. It has been shown that, contrary to common belief, money transfer in the presence of Byzantine faults can be implemented in asynchronous networks and does not require consensus. Nonetheless, existing implementations of money transfer still require a quadratic message complexity per payment, making attempts to scale hard. In common blockchains, such as Bitcoin and Ethereum, this cost is mitigated by *payment channels* implemented as a second layer on top of the blockchain allowing to make many off-chain payments between two users who share a channel. Such channels require only on-chain transactions for channel opening and closing, while the intermediate payments are done off-chain with constant message complexity. But payment channels in-use today require synchrony; therefore, they are inadequate for asynchronous money transfer systems.

In this paper, we provide a series of possibility and impossibility results for payment channels in asynchronous money transfer systems. We first prove a quadratic lower bound on the message complexity of on-chain transfers. Then, we explore two types of payment channels, unidirectional and bidirectional. We define them as shared memory abstractions and prove that in certain cases they can be implemented as a second layer on top of an asynchronous money transfer system whereas in other cases it is impossible.

2012 ACM Subject Classification Computing methodologies → Distributed algorithms

Keywords and phrases Blockchains, Asynchrony, Byzantine faults, Payment channels

Digital Object Identifier 10.4230/LIPIcs.DISC.2022.29

Related Version *Full Version*: <https://arxiv.org/abs/2202.06693>

Funding *Oded Naor*: Oded Naor is grateful to the Azrieli Foundation for the award of an Azrieli Fellowship, and to the Technion Hiroshi Fujiwara Cyber-Security Research Center for providing a research grant.

1 Introduction

The rise of *cryptocurrencies*, such as Bitcoin [33], Ethereum [42], Ripple [2], and many more, has revolutionized the possibility of using decentralized money systems.

In 2019, Guerraoui et al. [20] defined the abstraction of *asset transfer* or *money transfer* capturing the original motivation of Bitcoin. This abstraction is based on a set of known users owning accounts, each account has some initial money, and the users can transfer money between the accounts. It is well-known that deterministic consensus cannot be solved in an asynchronous network [18], meaning that blockchains that rely on consensus require synchrony to work properly. Nonetheless, Guerraoui et al. showed that the asset transfer problem is weaker than consensus [28, 15], i.e., in the Byzantine message-passing model the problem can be solved in an asynchronous network. They provide a concrete implementation of the abstraction in this model using an asynchronous broadcast service.



© Oded Naor and Idit Keidar;

licensed under Creative Commons License CC-BY 4.0

36th International Symposium on Distributed Computing (DISC 2022).

Editor: Christian Scheideler; Article No. 29; pp. 29:1–29:20

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

■ **Table 1** Summary of the results.

Abstraction	Operations	Upper bound message complexity	Lower bound message complexity
Asset transfer	transfer read	$O(n^2)$ [20, 4] $O(1)$	$\Omega(n^2)$ [Theorem 3]
Bidirectional payment channel	open transfer close	Impossible in asynchronous networks [Theorem 5]	
Unidirectional payment channel with source close	open transfer source close target close	Impossible in asynchronous networks [Theorem 7]	
Unidirectional payment channel	open transfer target close	$O(n^2)$ $O(1)$ [Alg. 4] $O(n^2)$	$\Omega(n^2)$ [Lemma 8]

Yet, in this solution, each payment requires a message complexity of $O(n^2)$, where n is the number of processes in the system. If the number of processes grows, this per-payment quadratic message complexity can pose a real challenge in scaling the asset transfer network. In fact, *scalability* is one of the major limiting factors of blockchains and consensus protocols, and extensive research was done to reduce the message complexity [34, 17, 27, 43, 11] in various settings.

A promising approach to scale blockchain payments is by using *payment channels* [36, 29, 38] as a second off-chain layer on top of the blockchain. A payment channel can be *opened* between two blockchain account owners via an on-chain deposit made to fund the channel, after which the two users can *transfer* payments on the channel itself off the blockchain. At any time, one of the users can decide to *close* the channel, after which the current balance in the channel of each of the users is transferred to their on-chain accounts. In this scheme, opening or closing a channel requires a blockchain transaction, which incurs a large message complexity, but all the intermediate payments on the channel require message exchange only between the channel users, freeing the blockchain from these payments and messages.

Payment channels have been actively deployed in central blockchains. For example, Bitcoin’s *Lightning Network (LN)* [36] is a highly used payment network, with active channels holding a Bitcoin amount equivalent to hundreds of millions of dollars [40].

One of the downsides of using payment channels such as LN is that the payment channels themselves rely on network synchrony. For example, in LN, suppose Alice and Bob have an open payment channel between them and Alice acts maliciously and tries to steal money by closing the channel at a stale state. LN provides a way for Bob to penalize Alice and confiscate all the money in the channel. But to do so, Bob has to detect Alice’s misbehavior on-chain and act within a predetermined time frame, making this method inappropriate with asynchronous users.

In this paper, we explore the possibility of implementing payment channels in asynchronous asset transfer systems. The results of the paper and prior art are summarized in Table 1. We study four abstractions: asset transfer and different types of payment channels.

First, we prove a lower bound on the complexity of asset transfer without channels. We show that no matter what implementation is provided for the asset transfer abstraction, it still requires $\Omega(n^2)$ messages for a single payment to be made by one party and observed by others. This means that the upper bound is tight for the algorithm provided in [20], in which

transferring money has $O(n^2)$ message complexity and reading the balance of an account costs $O(1)$. This fundamental result shows that while the synchrony requirement can be relaxed for asset transfer, each payment still requires a rather large number of messages. This means that second-layer solutions such as payment channels are required for scalability.

Next, we consider payment channel abstractions with operations for opening a channel, transferring money in it, and closing the channel. We first consider a *bidirectional payment channel* as a second layer atop an asset transfer system, where each side of the channel can make payments to the other. This is similar to LN [36], Teechain [29], Sprites [31], and more payment channel proposals. Once we formalize the problem, it is easy to show that synchrony is required, and therefore a bidirectional channel cannot be implemented on top of an asynchronous asset transfer system.

We next look at *unidirectional payment channels*, in which only one user, the channel *source*, can make payments to the other user, the *target*. We differentiate between two types of unidirectional payment channels: If we allow both the source and the target to close the channel (the source and target close operations, respectively), we again show that synchrony is required. Indeed, a previous design of similar channels [39] still requires synchrony.

On the other hand, if we allow only the target to close the channel, we provide a concrete implementation that works in an asynchronous network. In this implementation, the opening and closure of the channel require a payment using the asset transfer system, incurring $O(n^2)$ message complexity, whereas every payment on the channel itself requires a single message from the source user to the target user. We note that once the channel supports a target close the target can claim its transferred funds in the channel, which is not the case when only the source can close the channel.

Finally, we outline an extension of payment channels to payment chains, in which payments are made across multiple channels atomically. Like their 2-party counterparts, a chain payment over unidirectional channels with only target close can be implemented using a technique used in LN. As for other channel types, k -hop chains are equivalent to k -process consensus, i.e., have a consensus number [23] of k .

To conclude, our contributions in this paper are as follows:

- We prove a quadratic message complexity lower bound for asynchronous asset transfer systems.
- We explore payment channels as a key scaling solution for asynchronous asset transfer systems and provide a series of impossibility results for different channel types.
- We provide a concrete possibility result and an implementation for an asynchronous payment channel.

Structure. The rest of the paper is structured as follows: §2 describes the model and preliminaries; §3 details the asset transfer abstraction and proves a lower bound on message complexity; §4 discusses bidirectional payment channels and §5 discusses unidirectional channels; §6 extends the discussion to chain payments; §7 discusses related work; and finally, §8 concludes the paper. Some of the full proofs are deferred to Appendix A.

2 Model and preliminaries

We study a message-passing distributed system that consists of a set $\Pi = \{p_1, \dots, p_n\}$ of n processes. The processes can interact among themselves by sending messages. An adversary can *corrupt* up to $f < n/3$ processes, where $f \in \Theta(n)$. If not mentioned explicitly, we assume a corrupt process is *Byzantine*, i.e., it can deviate from the prescribed algorithm and act

arbitrarily. Any non-corrupt process is *correct*. A *crash-fail* fault is when a process stops participating in the algorithm. Every two processes share an asynchronous reliable link between them, such that if one correct process sends a message to another correct process, it eventually arrives, and the target can ascertain its source.

We assume the existence of a *Public Key Infrastructure (PKI)*, whereby processes that know a private key can use it to sign messages such that all other processes can verify the signature. The adversary cannot forge a signature if the private key is owned by a correct process. We assume each private key is owned by one process. We further assume *multisignatures* [6], whereby in order to produce a valid signature matching a constant-sized public key, more than one private key is used to sign the message. Note that signing can be sequential.

We study algorithms in the message-passing model that implement abstractions that are defined as *shared-memory objects*. A shared memory object has a set of *operations*, and processes access the object via these operations. Each operation starts with an *invocation* event by a process and ends with a subsequent *response* event. Invocations and responses are discrete events.

An *implementation* or an *algorithm* π of a shared-memory object abstraction is a distributed protocol that defines the behaviors of processes as deterministic state machines, where state transitions are associated with *actions*: sending or receiving messages, and operation invocations or responses. A *global state* of the system is a mapping to states from systems components, i.e., processes and links. An *initial global state* is when all processes are in initial states and there are no messages on the links between the processes. A *run* or an *execution* of an implementation is an alternating series of global states and actions, beginning with some initial global state, such that state transitions occur according to π . We assume that the first action of each process is an invocation of an operation and that it does not invoke another operation before receiving a response for its last invoked operation.

Each execution creates a *history* H that consists of a sequence of matching invocations and responses, each with the assigned process that invoked the operation and the matching responses. A *sub-history* H' of H is a subset of H 's events. Let $H|_p$ denote the sub-history of H with process p 's events.

A history defines a *partial ordering*: operation op_1 precedes op_2 in history H , labeled $op_1 \prec_H op_2$, if op_1 's response event happens before op_2 's invocation event in H . History H is *sequential* if each invocation, except perhaps the last, is immediately followed by a matching response. An operation is *pending* in history H if it has an invocation event in H but does not have a matching response. A history H' is a *completion of history* H if it is identical to H except for removing zero or more pending operations in H and by adding matching responses for the remaining ones. A shared-memory abstraction is usually defined in terms of a *sequential specification*. A *legal* sequential history is a sequential history that preserves the sequential specification, i.e., the sequential specification is the set of all legal histories.

The correctness criteria we consider is *Byzantine sequential consistency (BSC)*. This allows to extend sequential consistency [3] to runs with Byzantine processes and not only crash-fail and is an adaptation of the definition of Byzantine linearizability [10].

First, we formally define a sequentially consistent history for runs with crash-fail faults.

► **Definition 1** (sequentially consistent history). *Let E be an execution of an algorithm, and H its matching history. Then H is sequentially consistent if there exists a completion \tilde{H} of H and a legal sequential history S such that for every process p , $S|_p = \tilde{H}|_p$.*

An algorithm is sequentially consistent if all its histories are sequentially consistent.

For Byzantine sequential consistency (BSC), let $H|_c$ denote the sub-history of H with all of the operations of correct processes. We say a history is BSC if $H|_c$ can be augmented with operations of Byzantine processes such that the completed history is sequentially consistent. Formally:

► **Definition 2** (Byzantine sequential consistency (BSC)). *A history H is BSC if there exists a history H' such that $H'|_c = H|_c$, and H' is sequentially consistent.*

Similar to sequential consistency, an algorithm is BSC if all its histories are BSC. We choose BSC as the correctness criterion and not Byzantine linearizability because it simplifies the implementations we provide below. We explain in §5.3 how to change the implementation that we provide to satisfy Byzantine linearizability. The difference between sequential consistency and from linearizability [24] is that linearizability also preserves real-time order, i.e., for any operations op_1, op_2 s.t. $op_1 \prec_{\tilde{H}} op_2$, then $op_1 \prec_S op_2$.

3 Asset transfer

3.1 Asset transfer abstraction

Let A be an asset transfer abstraction, which is based on the one defined in [20]. A holds a set of accounts. Each account $a \in A$ is defined by some public key, and there is a mapping $owner(a): A \mapsto \Pi$, that matches for each account a the process that can produce a signature corresponding to the public key associated with a . In case of an account b associated with a multisignature, $owner(b)$ is the set of processes whose private keys can produce a signature matching b 's public key. The state of each account a is of the form $A(a) \in \mathbb{R}_{\geq 0}$ and represents the balance of the account a . Each account initially holds its initial balance.

A has two operations: The first, $A.read(a)$, returns the balance of account a , i.e., it returns $A(a)$ and can be called by any process. The second is $A.transfer(a, [(b_1, amt_1), \dots, (b_k, amt_k)])$, which, for every $1 \leq i \leq k$, transfers from account a 's balance amt_i and deposits it in b_i . This call succeeds, and returns *success* if it is called by $owner(a)$ and if the account has enough balance to make the transfer, i.e., $A(a) \geq \sum_{i=1}^k amt_i$. Otherwise, it returns *fail* and does nothing.

The *transfer* operation is an extension of the one defined in [20] in that it allows to transfer money from one account to multiple accounts, whereas the original work only allows transferring money from one account to another each time.

In this paper, we consider implementations of the asset transfer abstraction that are BSC. We note that the message-passing asset transfer implementation in [20] is based on a reliable broadcast that preserves source order. Their correctness criteria is neither Byzantine linearizability nor BSC, but it ensures that for every transfer operation, there exists a time t such that if a correct process invokes the read operation after t , then it observes the changes made by the transfer. This is a property we use throughout the proofs which are detailed below.

3.2 Message complexity of asset transfer

We begin by proving a quadratic message complexity lower bound on any asset transfer implementation in §3.2.1, and §3.2.2 shows this lower bound is tight by discussing a concrete implementation of an asset transfer that has a quadratic message complexity for a transfer operation and a constant message complexity for a read operation.

3.2.1 Lower bound

We show that if $f \in \Theta(n)$, then there is a quadratic message complexity lower bound in runs in which money is transferred to some account b , and multiple processes read the balance of b . The proof we use for the lower bound follows the technique used in Dolev-Reischuk's lower bound for Byzantine Broadcast [14]. Formally, we prove the following theorem:

► **Theorem 3.** *Consider an algorithm that implements the asset transfer abstraction. Then there exists a run with a single transfer invocation and multiple read invocations in which the correct processes send at least $(f/2)^2$ messages.*

Proof. Let π be an algorithm that implements asset transfer, and assume by contradiction that in all its runs with a single transfer correct processes send less than $(f/2)^2$ messages. We look at all executions of π with two accounts $a, b \in A$ s.t. $owner(a) = p$ for some process $p \in \Pi$, and initially $A(a) = 1, A(b) = 0$.

Consider first an execution σ_0 in which the adversary, denoted adv_0 , corrupts a set V of processes, not including p , such that $|V| = \lceil f/2 \rceil$. Denote the set of remaining correct processes as U . In σ_0 , process p calls $A.transfer(a, [(b, 1)])$. By the correctness definition of A , and since p is correct, there exists a time t_0 during the run after which any correct process that invokes $A.read(b)$ returns 1.

The adversary adv_0 causes the corrupt processes in V to simulate the behavior of correct processes that call $A.read(b)$ after t_0 , and follow the algorithm except for the following changes: they ignore the first $f/2$ messages they receive from processes in U , and they do not send any message to other processes in V . Note that while t_0 is not known to the processes, we construct the runs from the perspective of a global observer and may invoke read after t_0 .

Because correct processes send, in total, less than $(f/2)^2$ messages and corrupt processes do not send messages to other processes in V , then the processes in V together receive less than $(f/2)^2$ messages. Thus, by the pigeonhole principle, there exists at least one process $q \in V$ that receives less than $f/2$ messages. Denote the set of processes that send messages to q as U' , and denote $U'' = U \setminus U'$. Note that U' may include process p , and that $|U'| < f/2$.

Next, we construct a run σ_1 with an adversary adv_1 that are the same as σ_0 and adv_0 , respectively, except for the following changes: adv_1 corrupts all the processes in $V \setminus \{q\}$, and all the processes in U' . Since $|U'| < f/2$ and $|V| \leq \lceil f/2 \rceil$, the adversary adv_1 corrupts at most f processes in σ_1 . adv_1 prevents the corrupt processes from sending any message to q , but causes them to behave correctly towards all other correct processes in U'' .

By definition, the behavior of the corrupt processes in σ_1 , i.e., the processes in $U' \cup (V \setminus \{q\})$, towards the correct processes in U'' is the same as in σ_0 . Since process q simulates a correct process that ignores the first $f/2$ messages in σ_0 , its behavior towards the processes in U is identical in both runs as well. Thus, runs σ_0 and σ_1 are indistinguishable for the correct processes in U'' , ensuring that they behave the same. Since process q acts in σ_0 like a correct process that does not receive any message, both runs are indistinguishable to it as well.

Nonetheless, process q still has to return a value for its $A.read(b)$ call. Denote the time when the call returns as t_1 . If it returns a value different from 1, then we conclude the proof, since it is a violation of the read call specification. Otherwise, we construct a run σ_2 with an adversary adv_2 that are the same as σ_1 and adv_1 , respectively, except that there is no transfer invocation and all messages to q are delayed until after t_1 . For process q , runs σ_1 and σ_2 are indistinguishable until t_1 , therefore it returns 1 for the $A.read(b)$ call, violating the read call specification which should return 0, concluding the proof. ◀

We proved that the lower bound for the message complexity of an asset transfer object is $\Omega(n^2)$, assuming $f \in \Theta(n)$.

3.2.2 Upper bound

In [20], an implementation in the message-passing model for the asset transfer abstraction is provided. It uses a broadcast service defined in [30] that tolerates up to $f < n/3$ Byzantine failures. This broadcast has all the guarantees of reliable broadcast [9] (integrity, agreement, and validity), and also preserves source order, i.e., any two correct processes p_1 and p_2 that deliver messages m and m' broadcast from the same process p_3 , do so in the same order. The read operation is computed locally, and the transfer operation consists of a broadcast of a single message.

Several protocols can be used to implement such a source-order broadcast service, including a protocol in [30]. Bracha's reliable broadcast [9] can also be used to implement such a service if each correct process adds a sequence number to each message it broadcasts, and each correct process delivers messages from the same process in the order of the sequence numbers. These protocols have a message complexity of $O(n^2)$ per broadcast, proving that the lower bound message complexity we prove above is tight.

Note that our definition for the *transfer* call of the asset transfer abstraction allows transferring in each invocation money from one account to multiple accounts, while in [20] the transfer call allows a transfer to a single account for each invocation. The implementation in [20] can easily be adjusted to support this change by including in each broadcast message the multiple accounts to which money is transferred.

4 Bidirectional payment channel

We seek to analyze if payment channels that are used in common blockchains as a second layer can also be used similarly on top of an asynchronous asset transfer system. We discuss bidirectional payment channels, in which a channel is opened between two processes by making a transfer on the asset transfer system. After the channel is opened, both processes can make bidirectional payments on the same channel. Either process can close the channel at any time, after which their accounts in the asset transfer system reflect the state of the channel. This abstraction is similar to payment channels in the Lightning Network [36] in Bitcoin [33] and Raiden [38] in Ethereum [42]. We compare our payment channel abstractions to the currently available implementations in the related work in §7.

First, we formally define this abstraction, and then provide an impossibility result, proving it cannot be implemented in asynchronous networks.

4.1 Definition

We define a bidirectional payment channel abstraction as a shared memory object BC . The formal definition is in Specification 1. BC is defined based on the existence of an asset transfer object A . A channel in BC is of the form (a, b) , where a, b are accounts in A .

The state of a payment channel $BC(a, b)$ is $\{\mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0}\} \cup \{\perp\}$. The channel $BC(a, b)$ can be *open*, and then $BC(a, b) = (bal_a, bal_b)$, which represents the balances bal_a, bal_b of accounts a, b in the channel (a, b) , respectively. If the channel is *closed*, then its state is $BC(a, b) = \perp$.

The set of operations is the following:

- *open*. We do not provide a detailed specification for this call, as we do not require the full specification to prove that there is no implementation for a bidirectional payment channel in the asynchronous message-passing model. Instead, we assume that all channels are open at the beginning of the run with some initial balances.

■ **Specification 1 Bidirectional payment channel abstraction.** Operations for process p .

Shared Objects:	
A - asset transfer object	
BC - Bidirectional payment channel object	
1 Procedure $BC.transfer((a, b), amt)$:	9 Function $execute_payment((a, b), amt_a, amt_b)$:
2 if $BC(a, b) = \perp$	10 $(bal_a, bal_b) \leftarrow BC(a, b)$
3 return	11 $new_bal_a \leftarrow bal_a + amt_a$
4 $(bal_a, bal_b) \leftarrow BC(a, b)$	12 $new_bal_b \leftarrow bal_b + amt_b$
5 if $owner(a) = p \wedge bal_a \geq amt$	13 $BC(a, b) \leftarrow (new_bal_a, new_bal_b)$
6 $execute_payment((a, b), -amt, amt)$	
7 if $owner(b) = p \wedge bal_b \geq amt$	
8 $execute_payment((a, b), amt, -amt)$	
14 Procedure $BC.close((a, b), bal)$:	28 Function $execute_close((a, b), amt_a, amt_b)$:
15 if $B(a, b) = \perp$	29 $A(a) \leftarrow A(a) + amt_a$
16 return fail	30 $A(b) \leftarrow A(b) + amt_b$
17 $(curr_bal_a, curr_bal_b) \leftarrow BC(a, b)$	31 $BC(a, b) \leftarrow \perp$
18 $other_bal = curr_bal_a + curr_bal_b - bal$	32 return success
19 if $owner(a) = p$	
20 if $bal \neq curr_bal_a$	
21 return fail	
22 return $execute_close((a, b), bal, other_bal)$	
23 if $owner(b) = p$	
24 if $bal \neq curr_bal_b$	
25 return fail	
26 return $execute_close((a, b), other_bal, bal)$	
27 return fail	

- $transfer((a, b), amt)$. A payment in channel (a, b) is possible if the channel is open, if the caller process is a valid owner of either a or b , and the caller's balance in the channel is enough to make the payment. Otherwise, it does nothing. After the call ends, the state $BC(a, b)$ is changed to reflect the payment.
- $close((a, b), bal)$. This call can be invoked if the caller process is a valid owner of either a or b . The close of the channel is successful if the process that invokes the call does not try to close it with balance bal that is not the amount it has in the channel. Otherwise, the call fails and does nothing. The call transfers to accounts a, b their balances from the channel, and then changes its status to \perp . This call returns *success* or *fail* to indicate the outcome of the call.

In this paper, we consider sequentially consistent bidirectional channels.

4.2 Impossibility of a bidirectional payment channel object

We show that implementing a sequentially consistent bidirectional payment channel in the message-passing model requires synchrony. To this end, we solve wait-free consensus among 2 processes with shared registers and an instance of BC . The consensus abstraction has one call, $propose(v)$, which is called with some proposal v , and returns a value. The returned value for any process making the call has to be an input of the call from one of the processes, and it has to be the same value for all invocations, regardless of the caller process. We assume in this proof *crash-fail* faults, i.e., a process corrupted by the adversary stops participating in the protocol but does not deviate from it. Since this is an impossibility result and crash-fail faults are weaker than Byzantine faults, it also applies to runs with Byzantine processes.

► **Lemma 4.** *Consensus has a wait-free implementation for 2 processes in the read-write shared memory model with an instance of a bidirectional payment channel shared-memory object and shared registers.*

■ **Algorithm 2** Wait-free implementation of consensus among 2 processes using a bidirectional payment channel. Operations for processes $p_1 = owner(a), p_2 = owner(b)$.

Shared Objects:
 A - asset transfer object, initially two accounts $a, b \in A$ s.t. $A(a) = A(b) = 0$
 BC - Bidirectional payment channel object, initially $BC(a, b) = (1, 1)$
 R_1, R_2 - shared registers with read write calls, initially $R_1 = R_2 = \perp$

// We assume that at the beginning of the run there exists an open payment channel (a, b) s.t.
 $BC(a, b) = (1, 1)$

<pre> // Algorithm for process p1: 1 Procedure propose(v): 2 $R_1.write(v)$ 3 $BC.transfer((a, b), 1)$ 4 $BC.close((a, b), 0)$ 5 return make_decision() // Algorithm for process p2: 6 Procedure propose(v): 7 $R_2.write(v)$ 8 $BC.close((a, b), 1)$ 9 return make_decision() </pre>	<pre> // Algorithm for processes p1 and p2: 10 Procedure make_decision(): 11 wait until $A.read(b) \neq 0$ 12 if $A.read(b) = 2$ 13 return $R_1.read()$ 14 else 15 return $R_2.read()$ </pre>
---	---

Proof. The algorithm for solving consensus among two processes using a bidirectional payment channel object is detailed in Alg. 2. We assume that there are two processes p_1, p_2 with ownership of accounts a, b , respectively, and an open payment channel (a, b) at the beginning of the run with balances $BC(a, b) = (1, 1)$.

Before either of the processes invokes an operation on the payment channel, they write their proposal v in a shared register (Lines 2 and 7). Then, p_1 attempts to make a payment on the channel and then close it, and p_2 tries to close the channel without making or accepting any payment.

Because BC is sequentially consistent, the algorithm ensures that eventually after the channel is closed either the payment from a to b on the channel succeeds or not, and the balance in b 's account reflects it, i.e., there exists a time t after which invoking $A.read(b)$ returns either 1 or 2.

If the read call in Alg. 2 returns 2, then the channel was closed by p_1 after it successfully made the payment on the channel. Since before p_1 makes the payment on the channel it writes its proposal to register R_1 , then its value is already written by the time the channel is closed, and it is returned by the propose call. If the return value of the read call is 1, then process p_2 closed the channel. Since p_2 closes the channel after it writes its proposal in R_2 , then its proposed value is returned.

In either case, when the channel is closed, there is already a proposal written in either R_1 or R_2 , i.e., the returned value is an input to the propose call by either process, and both processes return the same value. ◀

Based on the above theorem and FLP [18], we get the following result:

► **Theorem 5.** *There does not exist an implementation of the bidirectional payment channel abstraction in the asynchronous message-passing model.*

5 Unidirectional payment channel

After proving that a bidirectional payment channel cannot be implemented in asynchronous networks, we explore another type of payment channel, *unidirectional*. The main difference from bidirectional channels is that unidirectional channels are asymmetric. There is only one

■ **Specification 3 Unidirectional payment channel abstraction.** Operations for process p .

Shared Objects:	
A - asset transfer object, with initial accounts	
B - unidirectional payment channel object	
1 Procedure $B.open((a, b), amt)$:	24 Procedure $B.source_close((a, b), bal_a)$:
2 if $p \neq owner(a) \vee B(a, b) \neq \perp \vee A(a) < amt$	25 if $p \neq owner(a) \vee B(a, b) = \perp$
3 return <i>fail</i>	26 return <i>fail</i>
4 $A(a) \leftarrow A(a) - amt$	27 $(curr_bal_a, curr_bal_b) \leftarrow B(a, b)$
5 $B(a, b) \leftarrow (amt, 0)$	28 if $bal_a \neq curr_bal_a$
6 return <i>success</i>	29 return <i>fail</i>
7 Procedure $B.transfer((a, b), amt)$:	30 $bal_b \leftarrow curr_bal_a + curr_bal_b - bal_a$
8 if $p \neq owner(a) \vee B(a, b) = \perp$	31 return $execute_close((a, b), (bal_a, bal_b))$
9 return	32 Function $execute_close((a, b), (bal_a, bal_b))$:
10 $(bal_a, bal_b) \leftarrow B(a, b)$	33 $A(a) \leftarrow A(a) + bal_a$
11 if $bal_a < amt$	34 $A(b) \leftarrow A(b) + bal_b$
12 return	35 $B(a, b) \leftarrow \perp$
13 $new_bal_a \leftarrow bal_a - amt$	36 return <i>success</i>
14 $new_bal_b \leftarrow bal_b + amt$	
15 $B(a, b) \leftarrow (new_bal_a, new_bal_b)$	
16 Procedure $B.target_close((a, b), bal_b)$:	
17 if $p \neq owner(b) \vee B(a, b) = \perp$	
18 return <i>fail</i>	
19 $(curr_bal_a, curr_bal_b) \leftarrow B(a, b)$	
20 if $bal_b \neq curr_bal_b$	
21 return <i>fail</i>	
22 $bal_a \leftarrow curr_bal_a + curr_bal_b - bal_b$	
23 return $execute_close((a, b), (bal_a, bal_b))$	

user, the source, who can open and transfer money in the channel. We show in which cases unidirectional payment channels can be implemented in an asynchronous message-passing network and in which cases they cannot. We begin by formally defining the unidirectional payment channel abstraction.

5.1 Definition

We define a unidirectional payment channel abstraction as a shared memory object B . The formal definition is in Specification 3. B is defined based on the existence of an asset transfer object A . A payment channel in B is of the form (a, b) where a, b are accounts in A .

The state of a payment channel $B(a, b)$ is $\{\mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0}\} \cup \{\perp\}$. Intuitively, a unidirectional payment channel $B(a, b)$ can either be open, and then $B(a, b) = (bal_a, bal_b)$, or closed, and then $B(a, b) = \perp$. The initial state is that all unidirectional payment channels are closed, e.g., for payment channel (a, b) , the state is $B(a, b) = \perp$ at the beginning of the run.

The set of operations is the following:

- $open((a, b), amt)$. A process that owns account a can open a unidirectional payment channel with any other account b with amount amt , as long as it has enough balance in A and does not already have an open payment channel with b . The call returns *success* if the channel is opened successfully and *fail* otherwise.
- $transfer((a, b), amt)$. A payment in the payment channel (a, b) is possible if the channel is open, if the caller of the operation a is $owner(a)$ and a has enough balance in the channel to make the payment. This call does not return a response.
- $source_close((a, b), bal_a)$. A source closing of a payment channel (a, b) can be called by $owner(a)$. The call succeeds if the process that invokes the call does not try to close it with balance bal_a that is not the amount it has in the channel. After the call ends, the balances in the channel are transferred to accounts a, b in the asset transfer system. The call returns *success* if the channel is closed successfully and *fail* otherwise.

- *target_close*((a, b), bal_b). A target closure of a payment channel (a, b) is symmetrical to the *source_close* call, but is invoked by *owner*(b).

We differentiate between two types of unidirectional payment channels, depending on whether the source close call is included in the allowed set of operations of the shared object or not. Note that without source close, the source depends on the target to close the channel to receive its deposit back after the channel is opened. However, for the target to receive its balance from the channel in the asset transfer system, it has to eventually close the channel. When the target closes the channel, the source also receives its respective balance.

We do not consider a channel with source close and without target close, as in this case only the source has operations it can call, and the target relies on the source for all its operations regarding the channel, and cannot receive the funds transferred to it in the channel on-chain unless the source closes the channel. Also, since in this case, only the source has operations it can invoke, this abstraction can be implemented easily in an asynchronous network as it does not require consensus or any interaction at all between the source and target. We also believe this case does not correspond correctly to existing implementations of payment channels as discussed in §7.

5.2 Impossibility of a unidirectional payment channel with source close

We show that a unidirectional payment channel that has the source close operation has a consensus number of at least 2, and therefore cannot be implemented in an asynchronous message passing network. Formally, we prove:

► **Lemma 6.** *Consensus has a wait-free implementation for 2 processes in the read-write shared memory model with an instance of a unidirectional payment channel with source close shared memory object and shared registers.*

The proof is similar to the proof of Lemma 4 and is deferred to Appendix A. Based on the above lemma and FLP [18], we get the following result:

► **Theorem 7.** *There does not exist an implementation of the unidirectional payment channel abstraction with source close in the asynchronous message-passing model.*

5.3 Unidirectional payment channel without source close

Next, we discuss unidirectional payment channel without the source close operation. We first prove a lower bound on the message complexity of an implementation and then prove that this lower bound is tight by providing an implementation for the abstraction in the asynchronous message-passing model.

Lower bound. We prove that any algorithm that implements the unidirectional payment channel specification incurs a combined message complexity for open, transfer, and close of $\Omega(n^2)$. To this end, we prove the following lemma:

► **Lemma 8.** *Consider an algorithm that implements the unidirectional payment channel abstraction B , and an asset transfer A . Then there exists a run with $B.open$, $B.transfer$, $B.target_close$, and $A.read$ calls, in which correct processes send at least $(f/2)^2$ messages.*

The full proof is deferred to Appendix A as it shares similar concepts to the lower bound proof of Theorem 3. The intuition is that an $A.transfer$ call can be simulated by opening a channel, transferring money in it, and then closing it.

■ **Algorithm 4** Unidirectional payment channel without `source_close` implementation in the asynchronous message-passing model. Operations for process p .

<p>Shared Objects: A - asset transfer object</p> <p>Local variables: $source[]$ - a dictionary with the balances of all channels that p is the source, initially \perp $target[]$ - a dictionary with A multisig invocations for all channels p is the target, initially \perp // This call can be invoked by $owner(a)$</p> <pre> 1 Procedure <i>open</i>((a, b), amt) : 2 if $owner(a) \neq p \vee source[ab] \neq \perp \vee A.read(a) < amt$ 3 return fail 4 invoke $A.transfer(a, [(ab, amt)])$ // ab is multisig 5 create $A.transfer(ab, [(a, amt), (b, 0)])$ 6 invocation tx 7 add p's signature to tx // tx is not a valid 8 transaction without $owner(b)$'s signature 9 send ("open", tx, amt) to $owner(b)$ 10 $source[ab] \leftarrow (amt, 0)$ 11 return success </pre> <p>// This call can be invoked by $owner(a)$</p> <pre> 15 Procedure <i>transfer</i>((a, b), amt): 16 if $owner(a) \neq p \vee source[ab] = \perp$ 17 return 18 $(bal_a, bal_b) \leftarrow source[ab]$ 19 if $bal_a < amt$ 20 return 21 $(new_bal_a, new_bal_b) \leftarrow (bal_a - amt, bal_b + amt)$ 22 create $A.transfer(ab, [(a, new_bal_a), (b, new_bal_b)])$ 23 invocation tx 24 add p's signature to tx 25 send ("transfer", tx, amt) to $owner(b)$ 26 $source[ab] \leftarrow (new_bal_a, new_bal_b)$ </pre> <p>// This call can be invoked by $owner(b)$</p> <pre> 34 Procedure <i>target_close</i>((a, b), bal_b): 35 if $owner(b) \neq p \vee target[ab] = \perp$ 36 return fail 37 get $A.transfer(ab, [(a, curr_bal_a), (b, curr_bal_b)])$ 38 transaction tx from $target[ab]$ 39 if $bal_b \neq curr_bal_b$ 40 return fail 41 add p's signature to tx // complete the multisig 42 invoke tx // invoke A with closing transaction 43 $target[ab] \leftarrow \perp$ 44 send ("close", (a, b)) to $owner(a)$ 45 return success </pre>	<p>// This message is received by $owner(b)$</p> <pre> 10 Upon receiving ("open", tx, amt) and $A.read(ab) = amt$: 11 let tx be $A.transfer(ab, [(a, amt), (b, 0)])$ invocation 12 if $owner(b) \neq p \vee \neg validate(tx, a) \vee target[ab] \neq \perp$ 13 return 14 $target[ab] \leftarrow tx$ </pre> <p>// This message is received by $owner(b)$</p> <pre> 26 Upon receiving ("transfer", tx, amt): 27 let tx be $A.transfer(ab, [(a, bal_a), (b, bal_b)])$ invocation 28 if $owner(b) \neq p \vee \neg validate(tx, a) \vee target[ab] = \perp$ 29 return 30 get $A.transfer(ab, [(a, c_bal_a), (b, c_bal_b)])$ invocation 31 from $target[ab]$ // The currently stored transaction 32 if $c_bal_a \neq bal_a - amt \vee c_bal_b \neq bal_b + amt$ 33 return 34 $target[ab] \leftarrow tx$ // store new transaction </pre> <p>// This message is received by $owner(a)$</p> <pre> 45 Upon receiving ("close", (a, b)) and $A.read(ab) = 0$: 46 $source[ab] \leftarrow \perp$ </pre> <p>Function <i>validate</i>(tx, a): 47 return tx is a valid invocation of A and it contains 48 $owner(a)$'s signature</p>
---	--

Upper bound. We provide an algorithm in the asynchronous message-passing model that implements a unidirectional payment channel without source close. The algorithm assumes an asset transfer system A , implemented as in [20], and discussed in §3.2.2. The algorithm is detailed in Alg. 4. We denote an account name with a string $c_1c_2 \cdots c_k \in A$ to refer to an account with a public key that is a k -of- k multisignature of $\{owner(c_1), \dots, owner(c_k)\}$. For example, to sign an invocation of the asset transfer object A of account ab , like transferring money from ab to another account, both $owner(a)$ and $owner(b)$ need to sign the message with their respective private keys before the call can be invoked. The transfer call with an appropriate multisignature can be invoked by any process, in particular, the last process to sign the invocation and complete the signature. When the algorithm mentions that a process creates an $A.transfer$ invocation, e.g., in lines 5 and 22, it does not mean the process invokes the transfer operation, but rather that it adds its signature to a multisignature message allowing an invocation of A 's operation. Any invocation of $A.transfer$ call is explicitly mentioned (lines 4, 40). We further assume FIFO order on messages sent between every two processes. This can be easily implemented with sequence numbers.

We explain below the implementation details of the algorithm for each of the operations:

- **Open.** The open procedure of a channel (a, b) (Alg. 4) requires $p_1 = \text{owner}(a)$ to make an initial deposit by invoking the transfer method of A from account a to a multisignature account ab (Alg. 4). After the transfer is completed, p_1 creates a transaction tx that transfers the deposit back to its account and 0 to $p_2 = \text{owner}(b)$ and sends it to p_2 (Alg. 4). Note that at this stage, process p_1 cannot invoke A with tx since it requires a multisignature, but when p_2 receives it, it can add its signature as well and then invoke A with the tx .

When p_2 receives tx , this transaction message it also waits for the balance in account ab to reflect the deposit (Alg. 4) to ensure the money was deposited in account ab using the asset transfer system, after which it considers the account as open.

- **Transfer.** When p_1 wants to transfer money in an open channel (a, b) (Alg. 4) it creates a transaction tx which is an $A.\text{transfer}$ invocation transferring money from the multisignature account ab to accounts a and b with the last balance of the channel after the payment. E.g., if the balance of the channel is $(10, 1)$, and p_1 wants to make a payment of 1 on the channel, it creates transaction tx required to invoke $A.\text{transfer}(ab, [(a, 9), (b, 2)])$, which transfers 9 money units to p_1 and 2 to p_2 . Then p_1 adds its signature to tx (Alg. 4), and sends it to p_2 , which stores it. Note that p_1 cannot invoke A with tx since it is still missing p_2 's signature. Thus, making a payment on the channel simply requires one message from the source user to the target user containing tx , and multiple payments can be made on the channel without invoking A 's transfer call.
- **Close.** When p_2 wants to close the channel (a, b) (Alg. 4), it takes the last transaction of account ab it received from p_1 and adds its signature to it (Alg. 4). p_2 's signature completes the multisignature, making it a valid transaction, and allowing p_2 to use it to invoke A 's transfer operation (Alg. 4). Process p_2 notifies p_1 that it closed the channel, after which p_1 considers the channel closed. After the channel is closed, p_1 can reopen it with a new call of the open operation.

Thus, opening and closing of the channel requires invoking a single $A.\text{transfer}$ operation, which incurs $O(n^2)$ messages because of the broadcast, but transferring money on the channel itself requires only one message per transfer.

Correctness. We prove below that the implementation (Alg. 4) is Byzantine sequentially consistent (BSC) with respect to the sequential specification (Specification 3).

► **Definition 9.** Let E be an execution of Alg. 4 and H its matching history. Let \tilde{H} be a completion of H by removing any pending open and close calls that did not reach A 's transfer call invocation (Lines 4 and 41, respectively), and let $\tilde{H}|_c$ be \tilde{H} 's history with the operations of correct processes.

Define H' as an augmentation of $\tilde{H}|_c$ as follows: For any correct process $q = \text{owner}(b)$ that invokes a successful $B.\text{target_close}((a, b), \text{bal}_b)$ s.t. process $p = \text{owner}(a)$ is a Byzantine process, we add before the target close call the following two invocations to H' by p :

- $B.\text{open}((a, b), \text{bal}_b)$ with an account a s.t. $A(a) \geq \text{bal}_b$. Since account a has enough money to open the channel, the open call succeeds.
- $B.\text{transfer}((a, b), \text{bal}_b)$ which is invoked immediately after the previous open call returns.

The two added Byzantine calls ensure that when q invokes the close operation, it succeeds. Next, we construct a linearization of H' .

► **Definition 10.** Let H' be the augmented history of Alg. 4 as defined in Definition 9. Let E' be a linearization of H' by defining the following linearization points:

- Any open or close call that fails is linearized immediately after its invocation.
- A transfer call that returns because of the *if* statements (Lines 16, 19) is linearized immediately after its invocation.

29:14 On Payment Channels in Asynchronous Money Transfer Systems

- Any successful $\text{open}((a, b), \text{amt})$ s.t. $q = \text{owner}(b)$ is a correct process, then it linearizes after q reaches Alg. 4. If q is Byzantine, the call linearizes when it ends.
- Any $\text{transfer}((a, b), \text{amt})$ that reaches Alg. 4 s.t. $q = \text{owner}(b)$ is a correct process, then it linearizes after q reaches Alg. 4. If q is Byzantine, the call linearizes when it ends.
- Any successful $\text{target_close}((a, b), \text{bal}_b)$ s.t. $p = \text{owner}(a)$ is a correct process, then it linearizes after p reaches Alg. 4. If p is Byzantine, the call linearizes when it ends.

The open and transfer calls change the state of the channel. By the sequential specification, the target can call target close with this new state. Therefore, the linearization point of these calls occur after the target receives the message informing it of the new state. Regarding the close call linearization point: the source can only reopen a channel after it learns that the channel has been closed, and therefore the linearization point is when the source receives the information of the closure and verifies it on-chain.

Next, we provide below the lemmas showing that the linearization E' satisfies the sequential specification. We defer the proofs of these lemmas to Appendix A.

► **Lemma 11.** *A B.open call for channel (a, b) succeeds only if the channel is closed when the call is invoked.*

► **Lemma 12.** *For any B.transfer call for channel (a, b) there is a preceding open call for the channel in H' .*

► **Lemma 13.** *For any successful B.target_close $((a, b), \text{bal}_b)$ call in H' there is a preceding B.open $((a, b), \text{amt})$ call for channel (a, b) s.t. $\text{amt} \geq \text{bal}_b$, followed by a B.transfer call that changes the state of the channel to $(\text{bal}_a, \text{bal}_b)$ for $\text{bal}_a = \text{amt} - \text{bal}_b$.*

► **Lemma 14.** *In an infinite execution of Alg. 4 every call invoked by a correct process eventually returns.*

Thus, we can conclude the following result from the lemmas above:

► **Theorem 15.** *Alg. 4 implements a Byzantine sequentially consistent unidirectional payment channel without source close abstraction.*

Changing the algorithm to be Byzantine linearizable. The algorithm can be changed to be Byzantine linearizable by having each message answered with an ack message. E.g., after the open message is received in Alg. 4, the process sends an ack message back to the original sender. The linearization point then is when the ack message is received. This change requires sending more messages as part of the algorithm and also extends the latency, but this change does not affect the overall asymptotic message complexity. This is also the reason why we choose BSC as the correctness criterion, not Byzantine linearizability.

6 Chain payments

We can extend the discussion of payment channels to chain payments. A chain payment system allows making payments off-chain between users who do not share a direct payment channel between them but do share a route through intermediate users. For example, suppose that Alice wants to make a payment to Bob, but she does not share a direct payment channel with him. Rather, she has a channel with Charlie and Charlie has a channel with Bob. A chain payment allows using the route from Alice to Bob via Charlie to make the payment on all channels atomically. Chain payments are used extensively in Lightning Network [36].

The intuitive way to define a chain payment abstraction is with a single operation that makes a payment through a chain $((a_1, a_2), (a_2, a_3), \dots, (a_{k-1}, a_k))$, and the outcome of the payment affects all channels of the chain in an atomic manner. For example, suppose that the balances of the above channels of the chain are $(bal_1, bal_2), \dots, (bal_{k-1}, bal_k)$, respectively, and a payment of amt is made via the chain. Then, after the linearization point, the balances of the channels are $(bal_1 - amt, bal_2 + amt), \dots, (bal_{k-1} - amt, bal_k + amt)$, respectively. In this case, assuming that the channels are bidirectional or unidirectional with source close, it can be proven that the consensus number [23] of such chain payment object is k , meaning this object can be used to solve consensus between k processes, in a similar manner to the 2-consensus we prove in this paper for these channel types.

We also note that even if the channels of the chain are unidirectional without source close, implementing a chain payment is not intuitive and straightforward in asynchronous networks. A possible solution is to adopt the use of *Hash Time-Locked Contracts (HTLCs)* [13, 25] which are in use in the Lightning Network for chain payments. An HTLC is a special conditioned payment between two users Alice and Bob. An HTLC allows Alice to make a conditioned payment to Bob that includes some timeout Δ and a hash value y . Bob can receive the payment if he exposes on-chain a value x whose hash value is y before Δ elapses. We can equip the asset transfer system with a similar Hash Time Contract operation, without the timeout component. That is, Bob receives the payment if he exposes x without any timeout assumptions. With this, we can implement chain payments in asynchronous networks based on the unidirectional channels without source close we presented in Alg. 4. We leave as open future work to formalize these ideas and explore other possible asynchronous implementations of chain payments.

7 Related work

A few works predating the blockchain era [35, 22], identified that a transfer-like system can be implemented under asynchrony. The first definition of the asset transfer abstraction is due to Guerraoui et al. [20]. Subsequently, Auvolat et al. [4] provide a weaker specification for asset transfer and detail an implementation that uses a broadcast service that guarantees FIFO order between every two processes. Astro [12] implements and empirically evaluates an asynchronous payment system.

Solving Byzantine consensus in an asynchronous network is possible by using randomization to circumvent the FLP [18] result. Earlier protocols such as [7, 37] have exponential message complexity. Later protocols such as [32, 1, 21, 27] improve the message complexity in various settings, but they are not deterministic, and therefore their performance can only be measured in the expected case.

There is also extensive research done on scaling cryptocurrencies using payment channels, including the Lightning Network [36], Teechain [29], Bolt [19], Sprites [31], Perun [16], Duplex Micropayment Channels [13], Raiden [38], and more. In Ethereum [42], there are multiple second-layer networks like Arbitrum [26], StarkNet [8], and Optimism [41]. All these works assume an underlying synchronous network to function correctly. E.g., Lightning Network uses a penalizing mechanism where one side of the channel can confiscate the other side's balance in the channel if it misbehaves and tries to close the channel at a stale state. But for this mechanism to work correctly, the penalizing party has to place a transaction on-chain within a certain time period, thus requiring a synchronous network. Brick [5] is a payment channel that preserves safety and liveness in asynchrony, but to do so requires a rather complex third-party entities (referred to as wardens) which validate channel transactions before they can be placed on-chain. This work also assumes rational (and not Byzantine) behavior of the wardens.

Spilman [39] proposed in 2013 a unidirectional payment channel implementation that shares similar concepts to Alg. 4. Spilman’s design has a timeout for each channel. The target of the channel has to close the channel before the timeout passes, otherwise, the source side can refund its initial deposit in the channel. Because of the timeout, this design relies on network synchrony. To the best of our knowledge, our implementation of a unidirectional channel without source close is the first that works on an underlying asynchronous network without requiring any third-party assistance to operate the channel, and in the presence of Byzantine processes.

We are also the first to show a quadratic lower bound for payments in asset transfer systems, as well as to explore second-layer payment channels as a scaling solution in asynchrony.

8 Conclusion

In this paper we presented the possibility of using payment channels in asynchronous asset transfer systems as a scaling solution. We showed that an asset transfer system requires a quadratic message complexity per payment. Then, we showed a series of possibility and impossibility results regarding payment channels as a scaling solution.

References

- 1 Ittai Abraham, Dahlia Malkhi, and Alexander Spiegelman. Asymptotically optimal validated asynchronous byzantine agreement. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, pages 337–346, 2019.
- 2 Frederik Armknecht, Ghassan O Karame, Avikarsha Mandal, Franck Youssef, and Erik Zenner. Ripple: Overview and outlook. In *International Conference on Trust and Trustworthy Computing*, pages 163–180. Springer, 2015.
- 3 Hagit Attiya and Jennifer L Welch. Sequential consistency versus linearizability. *ACM Transactions on Computer Systems (TOCS)*, 12(2):91–122, 1994.
- 4 Alex Auvolat, Davide Frey, Michel Raynal, and François Taïani. Money transfer made simple: a specification, a generic algorithm, and its proof. *Bull. EATCS*, 132, 2020. URL: <http://eatcs.org/beatcs/index.php/beatcs/article/view/629>.
- 5 Zeta Avarikioti, Eleftherios Kokoris-Kogias, Roger Wattenhofer, and Dionysis Zindros. Brick: Asynchronous incentive-compatible payment channels. In *International Conference on Financial Cryptography and Data Security*, pages 209–230. Springer, 2021.
- 6 Mihir Bellare and Gregory Neven. Identity-based multi-signatures from rsa. In *Cryptographers’ Track at the RSA Conference*, pages 145–162. Springer, 2007.
- 7 Shai Ben-David, Allan Borodin, Richard Karp, Gabor Tardos, and Avi Wigderson. On the power of randomization in on-line algorithms. *Algorithmica*, 11(1):2–14, 1994.
- 8 Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Scalable, transparent, and post-quantum secure computational integrity. *IACR Cryptol. ePrint Arch.*, 2018:46, 2018.
- 9 Gabriel Bracha. Asynchronous Byzantine agreement protocols. *Information and Computation*, 75(2):130–143, 1987.
- 10 Shir Cohen and Idit Keidar. Tame the Wild with Byzantine Linearizability: Reliable Broadcast, Snapshots, and Asset Transfer. In Seth Gilbert, editor, *35th International Symposium on Distributed Computing (DISC 2021)*, volume 209 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 18:1–18:18, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.DISC.2021.18.
- 11 Shir Cohen, Idit Keidar, and Alexander Spiegelman. Not a coincidence: Sub-quadratic asynchronous byzantine agreement whp. In *34th International Symposium on Distributed Computing (DISC 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.

- 12 Daniel Collins, Rachid Guerraoui, Jovan Komatovic, Petr Kuznetsov, Matteo Monti, Matej Pavlovic, Yvonne-Anne Pignolet, Dragos-Adrian Seredinschi, Andrei Tonkikh, and Athanasios Xytkis. Online payments by merely broadcasting messages. In *2020 50th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, pages 26–38. IEEE, 2020.
- 13 Christian Decker and Roger Wattenhofer. A fast and scalable payment network with bitcoin duplex micropayment channels. In *Symposium on Self-Stabilizing Systems*, pages 3–18. Springer, 2015.
- 14 Danny Dolev and Ruediger Reischuk. Bounds on information exchange for byzantine agreement. In *Proceedings of the First ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing*, PODC '82, pages 132–140, New York, NY, USA, 1982. Association for Computing Machinery.
- 15 Cynthia Dwork, Nancy Lynch, and Larry Stockmeyer. Consensus in the presence of partial synchrony. *Journal of the ACM (JACM)*, 35(2):288–323, 1988.
- 16 Stefan Dziembowski, Lisa Ekey, Sebastian Faust, and Daniel Malinowski. Perun: Virtual payment channels over cryptographic currencies. *IACR Cryptol. ePrint Arch.*, 2017:635, 2017.
- 17 Ittay Eyal, Adem Efe Gencer, Emin Gün Sirer, and Robbert Van Renesse. Bitcoin-ng: A scalable blockchain protocol. In *13th USENIX symposium on networked systems design and implementation (NSDI 16)*, pages 45–59, 2016.
- 18 Michael J Fischer, Nancy A Lynch, and Michael S Paterson. Impossibility of distributed consensus with one faulty process. *Journal of the ACM (JACM)*, 32(2):374–382, 1985.
- 19 Matthew Green and Ian Miers. Bolt: Anonymous payment channels for decentralized currencies. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 473–489, 2017.
- 20 Rachid Guerraoui, Petr Kuznetsov, Matteo Monti, Matej Pavlovič, and Dragos-Adrian Seredinschi. The consensus number of a cryptocurrency. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, pages 307–316, 2019.
- 21 Bingyong Guo, Zhenliang Lu, Qiang Tang, Jing Xu, and Zhenfeng Zhang. Dumbo: Faster asynchronous bft protocols. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, pages 803–818, 2020.
- 22 Saurabh Gupta. *A non-consensus based decentralized financial transaction processing model with support for efficient auditing*. Arizona State University, 2016.
- 23 Maurice Herlihy. Wait-free synchronization. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 13(1):124–149, 1991.
- 24 Maurice P Herlihy and Jeannette M Wing. Linearizability: A correctness condition for concurrent objects. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 12(3):463–492, 1990.
- 25 HTLC. Hash time locked contracts. Accessed: 2022-02-05. URL: https://en.bitcoin.it/wiki/Hash_Time_Locked_Contracts.
- 26 Harry Kalodner, Steven Goldfeder, Xiaoqi Chen, S Matthew Weinberg, and Edward W Felten. Arbitrum: Scalable, private smart contracts. In *27th USENIX Security Symposium 18*, pages 1353–1370, 2018.
- 27 Idit Keidar, Eleftherios Kokoris-Kogias, Oded Naor, and Alexander Spiegelman. All you need is dag. In *Proceedings of the 2021 ACM Symposium on Principles of Distributed Computing*, pages 165–175, 2021.
- 28 Leslie Lamport, Robert Shostak, and Marshall Pease. The byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4(3):382–401, 1982.
- 29 Joshua Lind, Oded Naor, Ittay Eyal, Florian Kelbert, Emin Gün Sirer, and Peter Pietzuch. Teechain: a secure payment network with asynchronous blockchain access. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles*, pages 63–79, 2019.
- 30 Dahlia Malkhi and Michael Reiter. A high-throughput secure reliable multicast protocol. *Journal of Computer Security*, 5(2):113–127, 1997.

- 31 Andrew Miller, Iddo Bentov, Ranjit Kumaresan, and Patrick McCorry. Sprites: Payment channels that go faster than lightning. *CoRR*, *abs/1702.05812*, 2017. arXiv:1702.05812.
- 32 Andrew Miller, Yu Xia, Kyle Croman, Elaine Shi, and Dawn Song. The honey badger of bft protocols. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pages 31–42, 2016.
- 33 Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. *Decentralized Business Review*, page 21260, 2008.
- 34 Oded Naor and Idit Keidar. Expected linear round synchronization: The missing link for linear byzantine smr. In *34th International Symposium on Distributed Computing (DISC 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.
- 35 Fernando Pedone and André Schiper. Handling message semantics with generic broadcast protocols. *Distributed Computing*, 15(2):97–107, 2002.
- 36 Joseph Poon and Thaddeus Dryja. The bitcoin lightning network: Scalable off-chain instant payments, 2016.
- 37 Michael O Rabin. Randomized byzantine generals. In *24th annual symposium on foundations of computer science (sfcs 1983)*, pages 403–409. IEEE, 1983.
- 38 Raiden. Raiden network, 2020. Accessed: 2022-02-05. URL: <https://raiden.network/>.
- 39 Jeremy Spilman. Anti dos for tx replacement, April 2013. URL: <https://lists.linuxfoundation.org/pipermail/bitcoin-dev/2013-April/002433.html>.
- 40 Lightning Network Statistics. Real-time lightning network statistics. Accessed: 2022-02-01. URL: <https://lml.com/statistics>.
- 41 Optimism website. Accessed: 2022-02-01. URL: <https://www.optimism.io/>.
- 42 Gavin Wood et al. Ethereum: A secure decentralised generalised transaction ledger. *Ethereum project yellow paper*, 151(2014):1–32, 2014.
- 43 Maofan Yin, Dahlia Malkhi, MK Reiter and, Guy Golan Gueta, and Ittai Abraham. HotStuff: BFT consensus with linearity and responsiveness. In *38th ACM symposium on Principles of Distributed Computing (PODC'19)*, 2019.

A Correctness proofs

In this section we provide the full deferred proofs from the paper.

► **Lemma 6.** *Consensus has a wait-free implementation for 2 processes in the read-write shared memory model with an instance of a unidirectional payment channel with source close shared memory object and shared registers.*

Proof. In Alg. 2, only process p_1 transfers money to p_2 using the channel, and they both attempt to close the channel: p_1 after the payment is made, and p_2 without accepting the payment. Thus, changing the close call in Alg. 2 to $B.source_close$ and the call in Alg. 2 to $B.target_close$ yields a consensus algorithm among 2 processes using a unidirectional payment channel with source and target close operations. ◀

► **Lemma 8.** *Consider an algorithm that implements the unidirectional payment channel abstraction B , and an asset transfer A . Then there exists a run with $B.open$, $B.transfer$, $B.target_close$, and $A.read$ calls, in which correct processes send at least $(f/2)^2$ messages.*

Proof. We can simulate an $A.transfer$ call between two accounts a, b with initial balances 1, 0, respectively. This is done by having $owner(a)$ call $B.open((a, b), 1)$, followed $B.transfer((a, b), 1)$, and lastly, having $owner(b)$ call $B.target_close((a, b), 1)$. If accounts a, b are owned by the same process p , we can construct exactly the the same runs used in the proof of Theorem 3 in order to prove this lemma. The only change is that we replace the $A.transfer$ call invoked by p in the original proof with the three calls mentioned above. ◀

Next, we provide the full proofs of the correctness of the unidirectional channel without source close implementation in §5.3. We prove that E' , the linearization of H' , satisfies the sequential specification of a unidirectional payment channel without source close. We assume that for channel (a, b) , $p = \text{owner}(a)$, $q = \text{owner}(b)$.

► **Lemma 11.** *A B .open call for channel (a, b) succeeds only if the channel is closed when the call is invoked.*

Proof. Immediate from the algorithm. A channel (a, b) opening fails if $\text{source}[ab] = \perp$ (Alg. 4). This is the case for all channels at the beginning of the run, or if the channel was previously closed successfully (Alg. 4). ◀

► **Lemma 12.** *For any B .transfer call for channel (a, b) there is a preceding open call for the channel in H' .*

Proof. If the transfer call in H' is invoked by a correct process, then from the algorithm $\text{source}[ab] \neq \perp$. This is only possible by the algorithm if p invokes an open call for the channel before the transfer invocation. If q is correct, then the open and transfer calls linearize when q receives the messages for the corresponding calls (Lines 10, 24). Since we assume FIFO order on the links between any two processes, then the open call linearizes before the transfer call. If q is Byzantine, then the open and transfer calls linearize immediately after they successfully return.

If p is Byzantine, then the transfer invocation is in H' because there is some close invocation by a correct process q . Before that, there is also a matching open call by p . The open call linearizes before the transfer call. ◀

► **Lemma 13.** *For any successful B .target_close($(a, b), \text{bal}_b$) call in H' there is a preceding B .open($(a, b), \text{amt}$) call for channel (a, b) s.t. $\text{amt} \geq \text{bal}_b$, followed by a B .transfer call that changes the state of the channel to $(\text{bal}_a, \text{bal}_b)$ for $\text{bal}_a = \text{amt} - \text{bal}_b$.*

Proof. If the close call ends successfully, then q has in $\text{target}[ab]$ a valid transaction A .transfer($ab, [(a, \text{bal}_a), (b, \text{bal}_b)]$), otherwise, the call fails. Therefore, if p is a correct process, it opens the channel (a, b) with some deposit amt , and transfers in the channel s.t. the balances in the channel change to $(\text{bal}_a, \text{bal}_b)$. Both the open and transfer are linearized before the close invocation, otherwise, the close call fails. If p is Byzantine, then we add the matching open and transfer invocations to H' which are linearized before the close invocation. ◀

► **Lemma 14.** *In an infinite execution of Alg. 4 every call invoked by a correct process eventually returns.*

Proof. In all cases where an open or close calls return *fail* it does so immediately, since it is done prior to any invocation of A . Transfer calls that return due to the if statements (Lines 16, 19) also return immediately.

For the open call, the if condition in the channel open call (Alg. 4) ensures that the process that invokes the call owns account a and that it has enough balance to open the channel. We also assume that a correct process does not invoke a new call before a previous call has a response event. Therefore, the conditions checked during the if statement hold when A 's transfer call is invoked, and by the asset transfer specification the call succeeds.

A transfer call that does not invoke any of A 's calls, nor does it wait for a reply after it sends the transaction in Alg. 4. Therefore, this call also returns immediately.

29:20 On Payment Channels in Asynchronous Money Transfer Systems

A `target_close` call that returns success invokes A with a transfer call that transfers money from account ab (Alg. 4). To reach this line, the process has to first check if the channel is open, and it has the matching transaction in $target[ab]$ during the if statement of the call. Therefore, invoking A will eventually succeed by the asset transfer specification, and the `target_close` call returns successfully. ◀