

# Cross-Chain Payment Protocols with Success Guarantees

Rob van Glabbeek

Data61, CSIRO, Sydney, Australia  
UNSW, Sydney, Australia

Vincent Gramoli

University of Sydney, Australia  
Data61, CSIRO, Sydney, Australia

Pierre Tholoniati

University of Sydney, Australia  
École Polytechnique, France

## Abstract

In this paper, we consider the problem of cross-chain payment whereby customers of different escrows—implemented by a bank or a blockchain smart contract—successfully transfer digital assets without trusting each other. Prior to this work, cross-chain payment problems did not require this success or any form of progress. We introduce a new specification formalism called *Asynchronous Networks of Timed Automata (ANTA)* to formalise such protocols. We present the first cross-chain payment protocol that ensures termination in a bounded amount of time and works correctly in the presence of clock skew. We then demonstrate that it is impossible to solve this problem without assuming synchrony, in the sense that each message is guaranteed to arrive within a known amount of time. We also offer a protocol that solves an eventually terminating variant of this cross-chain payment problem without synchrony, and even in the presence of Byzantine failures.

## 1 Introduction

With the advent of various payment protocols comes the problem of interoperability between them. A simple way for users of different protocols to interact is to do a *cross-chain payment* whereby intermediaries can help customer Alice transfer digital assets to Bob even though Alice and Bob own accounts in different banks or blockchains.

A payment between two customers of the same bank is simple. Alice just informs the bank that she wants to transfer a certain amount from her account to the account of the receiving party Bob; and then the bank carries out this request. Alice and Bob do not need to trust each other but need to trust the bank to not withdraw the money from Alice’s account and never deposit it on Bob’s. As Bob trusts the bank, he can issue a signed certificate assuring Alice that if the bank says that he has been paid, then Alice is off the hook, and any further disputes about possible non-payment to Bob are between Bob and the bank. To prevent litigation against her, Alice simply needs this statement from Bob as well as a statement that Bob has been paid. As Alice does not trust Bob this second statement must come from the bank.

To generalise this protocol to payments between customers of different banks, it helps if the two banks have ways to transfer assets to each other, and moreover trust each other. A sound protocol becomes:

- (i) Bob provides Alice with a signed statement that all he requires for her to have satisfied her payment obligation, is a statement from his own bank saying that he has been paid.
- (ii) Alice’s bank promises Alice that if she transfers money to Bob, she will get a statement from Bob’s bank that the transfer has been carried out.
  1. Alice orders her bank to initiate the transfer to Bob.
  2. Alice’s bank withdraws the money from her account, and sends it to Bob’s bank.
    - 3a. Bob’s bank places the money in Bob’s account
    - 3b. and notifies Alice’s bank of this.
  4. Alice’s bank forwards to Alice the statement by Bob’s bank saying that Bob has been paid.

This is roughly how payments between customers of different banks happen in the world of banking. Step (ii) is part of general banking agreements, not specific to Alice and Bob. Step (i) is typically left implicit in negotiations between Alice and Bob. All that Alice needs is the combination of steps (i), (ii) and 4 above. Once step (i) and (ii) have been made, she confidently takes step 1, knowing that this will be followed by Step 4. Alice’s bank is willing to take step (ii) because it trusts Bob’s bank, in the sense that step 2 *will* be followed by Step 3b. In fact, when the two banks trust each other, and have ways to transfer assets to each

other, they can abstractly be seen as one bigger bank, and the problem becomes similar to the problem of payments between customers of the same bank.




Figure 1: Trust relations

The problem becomes more interesting when the banks cannot transfer assets to each other and the only trust is the one of customers to their own bank. Typical solutions consist of considering banks as *escrows* and having intermediaries, like Chloe, that play the role of connectors between these escrows. Figure 1 depicts the relations of trust between three customers and two escrows and where the flow of money is from left to right. Thomas and Schwartz [19] proposed two interledger protocols: (i) the universal protocol requires synchrony [5] in that every message between participants is received within a known upper bound and the clock skew between participants is bounded; (ii) the atomic protocol coordinates transfers using an ad-hoc group of notaries, selected by the participants. Herlihy, Liskov and Shrira [11] represent a cross-chain payment as a deal matrix  $M$  where  $M_{i,j}$  characterises a transfer of some asset from participant  $i$  to participant  $j$ . They offer a timelock<sup>1</sup> commit protocol that requires synchrony, and a certified blockchain commit protocol that requires partial synchrony and a certified blockchain. The correctness proofs of these algorithms consist of observing reasonable properties that sum up to a correctness argument. But, as far as we know, there is no formal treatment of a cross-chain payment protocol in the presence of (multiplicative) clock skew.

In this paper, we introduce a new specification formalism for cross-chain payment protocols, called *Asynchronous Networks of Timed Automata (ANTA)*. ANTA simplify the representation of cross-chain payment to a customer automaton and an escrow automaton that describe states from which outgoing transitions are immediately enabled and states from which outgoing transitions are conditional upon some predicates. These automata allow us to reason formally about the liveness and safety of cross-chain payment protocols. ANTA differ from Alur and Dill’s *timed automata* [1], their networks [3] and I/O automata [13] in subtle ways, tuned to the problem at hand. We illustrate ANTA by specifying the interledger universal protocol and proving that it solves the time-bounded variant of the cross-chain payment problem. Moreover, we fine-tune the protocol to work correctly even in the presence of clock skew.

We also show that there exist no algorithm that can solve the time-bounded cross-chain payment problem without synchrony even if all participants either behave correctly or simply crash. This impossibility result relies on classic indistinguishability arguments from the distributed computing literature and highlights an interesting relation between the the cross-chain payment problem and the well-known transaction commit problem [7, 9, 8]. Inspired by this earlier work on the transaction commit, we define a weaker variant of our problem called the eventual cross-chain payment problem that relaxes the liveness guarantees to be solvable with partial synchrony. This new problem differs from the transaction commit problem and its variants like the non-blocking weak atomic commit problem [8] by tolerating Byzantine failures. It is also different from the interledger problem solved in [19], and from the problem solved by the certified blockchain commit protocol of [11] in a partially synchronous setting, by requiring some liveness. In particular, a protocol where all participants always abort is not permitted by our problem specification. We propose an algorithm that solves the eventual cross-chain payment problem only assuming partial synchrony, and in the presence of Byzantine failures, using the ANTA formalism.

Interestingly, the classical notion of *atomicity*, meaning that the entire transaction goes through, or is rolled back completely, is not appropriate for this kind of protocols. In the words of [11], “This notion of atomicity cannot be guaranteed when parties are potentially malicious: the best one can do is to ensure that honest parties cannot be cheated.”

<sup>1</sup>[https://en.bitcoin.it/wiki/Hashed\\_Timelock\\_Contracts](https://en.bitcoin.it/wiki/Hashed_Timelock_Contracts).

## 2 Related work

**The Interledger protocol.** In [19] a protocol is presented for payments across payment systems. Here a *payment system* can be thought of as an independent bank, where people can have accounts. The intended application is in digital payment systems, such as Bitcoin [16]. In [19] the payment systems, or rather the functionality they offer, are loosely referred to as *ledgers*, and payments between customers of different ledgers as *interledger payments*. Following popular terminology, we here speak of *escrows* and *cross-chain payments*.

The protocol from [19] generalises to the situation of a longer zigzag than indicated in Figure 1, involving  $n$  escrows and  $n-1$  intermediaries  $Chloe_i$ . It comes in two variants: an *atomic mode*, in which transfers are coordinated using an ad-hoc group of notaries, selected by the participants, and a *universal mode*, not requiring external coordination. The atomic mode relies on more than two-third of the notaries to be reliable, although it may be unknown which ones.

The correctness proof by Thomas and Schwartz [19] of the universal mode of the protocol requires the assumption of *bounded synchrony* on the communication between the parties in the protocol, simply called *synchrony* by Dwork et al. [5]. This assumption says that any message sent by one of the parties, is guaranteed to be received by the intended recipient(s). Moreover, there is a known upper bound on the time a message is underway. Additionally, there needs to be a known upper bound on the clock skew between participants. Here we point out that this assumption is necessary, in the sense that without this assumption, and without the help of mostly reliable notaries, no correct cross-chain payment protocol exists.

To make such a statement, we need to define when we consider a cross-chain payment protocol correct. The correctness proof by Thomas and Schwartz [19] consists of reasonable properties that are shown to hold for the chosen protocol. These properties are stated in terms of that specific protocol, and the reader can infer that they sum up to a correctness argument. But no formal statement occurs of what it means for a general cross-chain payment protocol to be correct, and this is what we need to make any negative statement about it.

**Other cross-chain technologies.** In the lightning network, Poon and Dryja [17] can relay payment outside the blockchain or “offchain” through connected intermediaries, but they require synchrony and do not consider clock skew. Gazi, Kiayias and Zindros [6] propose a rigorous formalisation of ledger and cross-chain transfers, but focus on proof-of-stake (in particular Cardano Ouroboros). The assumption of “semi-synchronous communications”, which in practice corresponds to synchrony as defined by Dwork, Lynch and Stockmeyer [5], make their results not extensible to partial synchrony. They do not consider clock skew. Herlihy proposes atomic cross-chain swaps [10] to exchange assets between distinct blockchains, but only in a synchronous environment without clock skew. Ron van der Meyden [14] verifies a cross-chain swap protocol by modelling a timelock predicate as a Boolean variable indicating whether the asset is transferred. This approach also requires synchrony, and does not consider clock skew. In XCLAIM [21] Zamyatin et al. propose a solution to swap blockchain-backed assets. Their protocol assumes that adversaries are behaving rationally, and requires synchrony.

Lind et al. [12] relax the synchrony assumption but require a trusted execution environment (TEE). Such a solution cannot be used to solve our problem as it would require trusting a third-party, often represented as the manufacturer of this TEE.

Other approaches [20, 18, 11] rely on a separate blockchain that plays the same role as our transaction manager (cf. Section 6.2). However, [20] and [18] do not aim at ensuring liveness, and [11] aims at ensuring liveness only in periods where communication proceeds synchronously. Wood [20] proposes a multi-chain solution that aims at combining heterogeneous blockchains together without trust. As far as we know, it has not been proved that the protocol terminates. Ranchal-Pedrosa and Gramoli [18] relax the synchrony assumption using an alternative ‘child’ blockchain to the so-called ‘parent’ blockchain in order to execute a series of transfers outside the parent blockchain. This protocol does not guarantee that the set of intermediary transfers on the child blockchain eventually take effect. Herlihy, Liskov and Shrira [11] model cross-chain deals as a matrix  $M$  where  $M_{i,j}$  characterises a transfer of some asset from participant  $i$  to participant  $j$ , and offer a timelock-based solution under the synchrony assumption, without clock skew, and a certified blockchain commit protocol that requires partial synchrony and a certified blockchain. As remarked in [11], a strong liveness guarantee is not feasible when merely assuming partial synchrony. In this context the targeted cross-chain deal problem admits solutions where all correct processes simply abort. Our corresponding problem differs by requiring a weaker liveness guarantee in that it formulates conditions under which a successful

transfer is ensured. We present a more detailed comparison between our work and that of [11] in Section 7.

**Crash fault tolerant solutions.** The transaction commit problem is a classical problem from the database literature, tackled for instance by Gray & Lamport [7], Guerraoui [8] and Hadzilacos [9]. It consists of ensuring that either all the processes commit a given transaction or all the processes abort this transaction.

The Non-Blocking Atomic Commitment problem has been formally defined by Guerraoui, for asynchronous systems with unreliable failure detectors (encapsulating partial synchrony) in a crash (fail-stop) model [8] but not in a Byzantine model. Guerraoui proved an impossibility result when the problem requires that “If all participants vote *yes*, and there is no failure, then every correct participant eventually decides *commit*”. He then proposes a weaker variant of the problem where the above requirement is replaced by “If all participants vote *yes*, and no participant is ever suspected, then every correct participant eventually decides *commit*”. The two problems we consider in this paper present a similar distinction; however, our problems target Byzantine fault tolerance, not crash fault tolerance.

Anta, Georgiou and Nicolaou [2] propose a general and rigorous definition of a ledger in which they consider the atomic append problem in an asynchronous model. Similarly, their model considers exclusively crash failures.

### 3 Model and definitions

#### 3.1 Participants and money

We assume  $n$  banks or *escrows*  $e_0, \dots, e_{n-1}$  and  $n+1$  *customers*  $c_0, \dots, c_n$ . These  $2n+1$  processes are called *participants*.

**Escrows.** An escrow is a specific type of process that can handle values for other parties in a predefined manner.

**Customers.** Customer  $c_0$  is Alice and  $c_n$  is Bob. The customers  $c_1, \dots, c_{n-1}$  are intermediaries in the interaction between Alice and Bob; we call them *connectors*, named  $Chloe_i$ .

**Trust.** Customers  $c_{i-1}$  and  $c_i$  have accounts at escrow  $e_{i-1}$ , and trust this escrow ( $i = 1, \dots, n$ ). We do not assume any other relations of trust.




Figure 2: Customers and escrows.

**Topology.** Not every customer can send value to any other. Here we assume that value can be transferred directly only between customers of the same escrow. Moreover, any transfer between two customers of an escrow can be modelled as two transfers: one from the originating customer to the escrow, and one from the escrow to the receiving customer. Thus, the connections from Figure 2 describe both the relations of trust and the possible transfers of value. The case  $n = 2$  was depicted in Figure 1.

We can talk of the *customers of an escrow* to designate the customers that are connected to this escrow. We also talk of the *escrows of a customer* to designate the escrows that are connected to this customer—only one escrow for Alice and Bob, and two escrows for a connector.

**Placing value in escrow.** Customers can send a specific type of message to ask their escrow to put money aside for them. In particular, two customers may make a deal with an escrow to place value originating from the first customer “in escrow”, and, after a predefined period, depending on which conditions are met, either complete the transfer to the second customer, or return the value to the first one.

**Abstracting the transfer of value.** There are many ways of transferring value from one party to another: one could give a physical object, such as cash or gold, to someone else. One could also send a transaction to transfer cryptocurrency or tokens on a distributed ledger, or send a specific message on some banking application. We do not care of how this process is implemented, and we suppose that the participants have already agreed upon the value they expect to be transferred.

We use therefore a unified notation:  $s(p, \$)$  to say “send a message to trigger the transfer of some previously agreed-upon value to participant  $p$ ”.

**Chloe’s fee.** As Chloe helps out transferring value from Alice to Bob, it is only reasonable that she is paid a small commission. Hence the value transferred from Alice to Chloe might be larger than the value transferred from Chloe to Bob. Additionally, these values may be expressed in different currencies, with possibly fluctuating exchange rates, or they may be objects such as bags of flour that have a quality-dependent value. Deciding which values to transfer may thus be an interesting problem. However, it is entirely orthogonal to the matter discussed in this paper, and hence we shall not consider it any further.

### 3.2 Communication and computation model

The following model holds for the rest of the paper. When necessary, we will mention explicitly if we have to add some assumptions, such as synchrony of communication.

**Communication.** We assume that the network does not lose, duplicate, modify or create messages. However messages can be delayed for arbitrarily long: by default we assume asynchronous communication.

**Authentication.** We assume that each customer can sign a message with his unique identifier, thanks to an idealised public key infrastructure. No other process can forge its signature, and any process (including escrows) can verify it. An escrow is not necessarily able to sign a message.

**Certificates.** As a consequence, any customer can issue a certificate by signing an appropriate message. For example, Bob can issue a receipt certificate to Alice by signing a RECEIVED message. By combining several signatures, one can define threshold certificates, for instance requiring the signature of a COMMIT message by strictly more than  $n/3$  customers or notaries. A correct implementation of certificates should take care of preventing replay attacks.

**Faults.** We do not make any assumption on the behaviour of the processes *a priori*. Later we will define a protocol, and processes will either follow the protocol or deviate from the protocol.

### 3.3 Synchronous versus asynchronous communication

In the literature five levels of synchrony in communication can be distinguished. As indicated in the table below, terminology is not uniform between the concurrency and the distributed systems communities.

<i>concurrency</i>	<i>distributed systems</i>
synchronous communication	rendezvous
I/O automata	
asynchronous communication	synchronous communication
	partially synchronous communication
	asynchronous communication

In the concurrency community, communication is called *synchronous* if sending and receiving occur simultane-

ously, and the sender cannot proceed before receipt of the message is complete. This is the typical paradigm in process algebras such as CCS [15]. Communication is called *asynchronous* if sending occurs strictly before receipt, and the sender can proceed after sending regardless of the state of the recipient(s). An intermediate form is modelled by I/O automata [13]; here sending and receipt is assumed to occur simultaneously, yet the sender proceeds after sending regardless of the state of possible receivers.

In the distributed systems community all communication is by default assumed to be asynchronous in the sense above; synchronous communication as defined above is sometimes called a *rendezvous*. Following [5], communication is called *synchronous* when there is a known upperbound on the time messages can be in transit, and moreover there is a known upperbound on the relative clock skew between parallel processes. It is *partially synchronous* when these upperbounds exist but are not known, or when it is known that after a finite but unknown amount of time these upperbounds will come into effect. If these conditions are not met, communication is deemed *asynchronous*.

The present paper follows the latter terminology; we speak of *fully synchronised* communication when referring to synchronous communication from concurrency theory.

### 3.4 Informal description of the syntax and semantics of ANTA

Although strongly inspired by the timed automata of Alur and Dill [1] and their networks [3], our ANTA differs from those models in subtle ways, tuned to the problem at hand. The first A refers to asynchronous communication in the concurrency-theoretic sense, and contrasts with the fully synchronised communication that is assumed in NTA [3].

In Figure 3 our time-bounded cross-chain protocol—essentially the universal mode of the interledger protocol from [19]—is depicted as an ANTA. There is one automaton for each participant in the protocol, that is, for each escrow  $e_i$  ( $i = 1, \dots, n-1$ ) and each customer  $c_i$  ( $i = 1, \dots, n$ ). Each automaton is equipped with a unique identifier, in this case  $e_i$  and  $c_i$ . Each automaton has a finite number of *states*, depicted as circles, one of which is marked as the *initial state*, indicated by a short incoming arrow. The states are partitioned into *termination states*, indicated by a double circle, *input states*, coloured white, and *output states*, coloured grey or black. Furthermore there are finitely many *transitions*, indicated as arrows between states. The transitions are partitioned into *input transitions*, labelled  $r(id, m)$ , *output transitions*, labelled  $s(id, m)$ , and *time-out transitions*, labelled by arithmetical formulas  $\psi$  featuring the variable *now*. Here  $id$  must be the identifier of another automaton in the network, and  $m$  a message, taken from a set `MSG` of allowed messages. Whereas each input and time-out transition has a unique label  $r(id, m)$  and  $\psi$ , respectively, an output transition may have multiple labels  $s(id, m)$ . All transitions may have additional labels of the form  $x := now$  for some variable  $x$ . A termination state has no outgoing transitions. An output state has exactly one outgoing transition, which must be an output transition. An input state may have any number of outgoing input and time-out transitions, and no outgoing output transitions.

Each automaton keeps an internal clock, whose value, a real number, is stored in the variable *now*. The value of *now* increases monotonically as time goes on. All variables maintained by an automaton are local to that automata, and not accessible by other automata in the network. Each transition is assumed to occur instantaneously, at a particular point in time. In case a transition occurs that is labelled by an assignment  $x := now$ , the variable  $x$  will remember the point in time when the transition took place. Such a variable may be used later in time-out formulas. When (or shortly before, see below) an output transition with label  $s(id, m)$  occurs, the automaton sends the message  $m$  to the automaton with identifier  $id$ . A time-out transition labelled  $\psi$  is *enabled* at a time *now* when the formula  $\psi$  evaluates to `true`. An input transition labelled  $r(id, m)$  is *enabled* only at a time when the automaton receives the message  $m$  from the automaton  $id$  in the network. Whereas an output transition may be scheduled to occur by the automata at any time, an input or time-out transition can occur only when enabled.

When an automaton is not performing a transition, it must be in exactly one of its states. It starts at the initial state, where its clock is initialised with an arbitrary value. When the automaton is in an input state, it stays there (possibly forever) until one of its outgoing transitions becomes enabled; in that case that transition will be taken immediately. In case multiple transitions become enabled simultaneously, the choice is non-deterministic. When the automaton reaches a termination state, it halts.

In general, an output state is labelled with a positive *time-out* value  $to \in \mathbb{R} \cup \{\infty\}$ . It constitutes a strict upperbound on the time the automaton will stay in that state. In case the automaton enters an output state at time *now*, it will take its outgoing transition between times *now* and  $now + to$ . In case its output

transition has multiple labels  $s(id, m)$ , the corresponding transmissions need not occur simultaneously; they can occur in any order between *now* and *now + to*. The output transition is considered to be taken when the last of these actions occurs. In this paper time-out values are indicated by colouring: for the grey states it is the constant  $\varepsilon$  from Section 4.1, while for the black state it is  $\infty$ .

### 3.5 Cross-chain payment protocol

A cross-chain payment protocol prescribes a behaviour for each of the participants in the protocol, the escrows and the customers. Let  $\chi$  be a certificate signed by Bob saying that Alice’s obligation to pay him has been met.

**Definition 1** (Time-bounded cross-chain payment protocol). A cross-chain payment protocol is a *time-bounded cross-chain payment protocol* if it satisfies the following properties:

**C Consistency.** For each participant in the protocol it is possible to abide by the protocol.

**T Time-bounded termination.** Each customer that abides by the protocol, and either makes a payment or issues a certificate, terminates within an a priori known period, provided her escrows abide by the protocol.

**ES Escrow security.** Each escrow that abides by the protocol does not lose money.

**CS Customer security.**

CS1 Upon termination, if Alice and her escrow abide by the protocol, Alice has either got her money back or received the statement  $\chi$ .

CS2 Upon termination, if Bob and his escrow abide by the protocol, Bob has either received the money or not issued certificate  $\chi$ .

CS3 Upon termination, each connector that abides by the protocol has got her money back, provided her escrows abide by the protocol.

**L Strong liveness.** If all parties abide by the protocol, Bob is paid eventually.

Requirement C (*consistency* of the protocol) is essential. In the absence of this requirement, any protocol that prescribes an impossible task for each participant would be a correct cross-chain payment protocol (since it trivially meets T, ES, CS and L).

Requirements ES and CS (the *safety* properties) say that if a participant abides by the protocol, nothing really bad can happen to her. These requirements do not assume that any other participant abides by the protocol, and should hold no matter how malicious the other participants turn out to be. The only exception to that is that the safety properties for a customer (CS) are guaranteed only when the escrow(s) of this customer abide by the protocol.

Property L, saying that the protocol serves its intended purpose, is the only one that is contingent on *all* parties abiding by the protocol.

## 4 A time-bounded cross-chain payment protocol

### 4.1 Assumptions

**Synchrony.** The assumption of *synchrony* considered e.g. by [5], and called *bounded synchrony* by [19], says ‘that there is a fixed upper bound  $\Delta$  on the time for messages to be delivered (*communication is synchronous*) and a fixed upper bound  $\Phi$  on the rate at which one processor’s clock can run faster than another’s (*processors are synchronous*), and that these bounds are known a priori and can be “built into” the protocol.’ [5] A consequence of this assumption is that if participant  $p_1$  sends at its local time  $t_0$  a message to participant  $p_2$ , and participant  $p_2$  takes  $t$  units of its local time to send an answer back to  $p_2$ , then  $p_1$  can count on arrival of that reply no later than time  $t_0 + \Phi \cdot t + 2 \cdot \Delta$ .

**Bounded reaction speed.** When a participant in the protocol receives a message, it will take some time to calculate the right response and then to transmit that response. Here it will be essential that that amount of time is bounded. So we assume a reaction time  $\varepsilon > 0$  such that any received message can be answered within time  $\varepsilon$ .

## 4.2 A cross-chain payment protocol formalised as an ANTA

In this section we formally model the *universal mode* of the interledger protocol from [19] as an ANTA. Moreover, we replace the timing constants employed in [19] by parameters, and then calculate the optimal value of these parameters to unsure correctness of the protocol in the presence of clock skew.

To interpret Figure 3, all that is left to do is specify the messages that are exchanged between escrows and their customers. We consider 4 kinds of messages. One is the certificate  $\chi$ , signed by Bob, saying that Alice’s obligation to pay him has been met. Another is the value  $\$$  that is transmitted from one participant to another. The remaining messages are promises made by escrow  $e_i$  to its customers  $c_i$  and  $c_{i+1}$ , respectively:

$G(d) :=$  “I guarantee that if I receive  $\$$  from you at my local time  $w$ , then I will send you either  $\$$  or  $\chi$  by my local time  $w + d$ .”

$P(a) :=$  “I promise that if I receive  $\chi$  from you at my local time  $v$ , with  $v < \text{now} + a$ , then I will send you  $\$$  by my local time  $v + \varepsilon$ .”

The automata of Figure 3 can be informally described as follows: A escrow  $e_i$  first send promise  $G(d_i)$  to its (upstream) customer  $c_i$ . Here “upstream” refers to the flow of money. The precise values of  $d_i$  will be determined later; here they are simply parameters in the design of the protocol. Then it awaits receipt of the money/value from customer  $c_i$ . If the money does arrive, the escrow issues promise  $P(a_i)$  to its downstream customer  $c_{i+1}$  as soon as it can. It remembers the time this promise was issued as  $u$ . Then it awaits receipt of the certificate  $\chi$  from customer  $c_{i+1}$ . If the certificate does not arrive by time  $u + a_i$ , a time-out occurs, and the escrow refunds the money to customer  $c_i$ . If the money does arrive in time, the escrow reacts by forwarding the certificate to customer  $c_i$ , and forwarding the money to customer  $c_{i+1}$ .

A connector  $\text{Chloe}_i$  starts by awaiting promises  $G(d_i)$  from her downstream escrow  $e_i$ , and  $P(a_{i-1})$  from her upstream escrow  $e_{i-1}$ . Then she proceeds by sending the money to escrow  $e_i$ . After sending the money,  $\text{Chloe}_i$  waits for escrow  $e_i$  to send her either the certificate  $\chi$  or the money back. In the latter case, her work is done; in the former, she forwards the certificate to escrow  $e_{i-1}$  and awaits for the money to be send by escrow  $e_{i-1}$ .

The automata for Alice and Bob are both simplifications of the one for  $\text{Chloe}_i$ . Alice awaits promise  $G(d_0)$  from her escrow, and then sends the escrow the money. The protocol allows her to wait arbitrary long before taking that step. Subsequently, she patiently await for either the return of her money, or certificate  $\chi$ . Bob awaits promise  $P(a_{n-1})$  from his escrow, and then issues certificate  $\chi$  and sends it to his escrow. He then awaits the money.


## 4.3 Running the protocol

The protocol consists of two parts. The *set-up* involves the sending and receiving of the promises  $G(d_i)$ . As these promises are not time-sensitive, they can be exchanged months before the *active part* of the protocol is ran, consisting of all other actions. Here an *action* is an entity  $\text{act}@p$ , with  $\text{act}$  a transition label, and  $p$  the identifier of the participant taking that transition. The active part has essentially only one successful run, i.e., when never taking a time-out transition, consisting of the following actions, executed in the following order. Actions separated by commas can be executed in either order.


$s(e_0, \$)@c_0$	$r(c_0, \$)@e_0$	$s(c_1, P(a_0))@e_0$	$r(e_0, P(a_0))@c_1$
$s(e_1, \$)@c_1$	$r(c_1, \$)@e_1$	$s(c_2, P(a_1))@e_1$	$r(e_1, P(a_1))@c_2$
...			
$s(e_{n-1}, \$)@c_{n-1}$	$r(c_{n-1}, \$)@e_{n-1}$	$s(c_n, P(a_{n-1}))@e_{n-1}$	$r(e_{n-1}, P(a_{n-1}))@c_n$
$s(e_{n-1}, \chi)@c_n$	$r(c_n, \chi)@e_{n-1}$	$s(c_n, \$), s(c_{n-1}, \chi)@e_{n-1}$	$r(e_{n-1}, \$)@c_n, r(e_{n-1}, \chi)@c_{n-1}$
...			
$s(e_1, \chi)@c_2$	$r(c_2, \chi)@e_1$	$s(c_2, \$), s(c_1, \chi)@e_1$	$r(e_1, \$)@c_2, r(e_1, \chi)@c_1$
$s(e_0, \chi)@c_1$	$r(c_1, \chi)@e_0$	$s(c_1, \$), s(c_0, \chi)@e_0$	$r(e_0, \$)@c_1, r(e_0, \chi)@c_0$




$e_i$ : Escrow  $i$  ( $i = 0, \dots, n-1$ )



$c_i$ : Customer  $i$  ( $i = 1, \dots, n-1$ );  $Chloe_i$



$c_0$ : Customer 0; *Alice*



$c_n$ : Customer  $n$ ; *Bob*




Figure 3: Automata representing escrows and customers

## 4.4 Initialisation

There is a scenario where  $\text{Chloe}_i$  will never send money to escrow  $i$ , namely when she receives promise  $P(a_{i-1})$  from escrow  $e_{i-1}$  before she receives promise  $G(d_i)$  from escrow  $e_i$ . In that case the receipt of  $P(a_{i-1})$  does not trigger a transition, and  $\text{Chloe}_i$  will remain stuck in her second state. We now modify the protocol in such a way that that scenario can not occur. This can be done in three ways; it does not matter which of the three modifications we take.

1. One solution is to make Alice wait before starting the active part of the protocol (by leaving the black state) until a point in time when she is sure that all parties  $\text{Chloe}_i$  already have received promise  $G(d_i)$ . If we assume that all parties start at the same time, using the reasoning of Section 4.1, Alice has to wait at most  $\Phi \cdot \varepsilon + \Delta$  before this point has been reached. The only drawback of this solution is that it may be hard to realise that all parties start at the same time.
2. Another approach is to assume that the set-up phase occurred long before Alice actually wants to send money to Bob. It may be part of a general banking agreement. Possibly each escrow always offers promises  $G(d)$  for different values of  $d$ , and when sending money to escrow  $i$ ,  $\text{Chloe}_i$  simply tags it as taking advantage of promise  $G(d_i)$ . In this approach, the protocol does not feature the transitions labelled  $s(c_i, G(d_i))$  and  $r(e_i, G(d_i))$ , with the initial states shifted accordingly. Still, the promise  $G(d_i)$  counts as having been made to customer  $c_i$  by escrow  $e_i$ .
3. A final solution is to introduce to the protocol a message “We are ready”, sent by  $\text{Chloe}_{n-1}$  to escrow  $e_{n-2}$ , and forwarded, via  $\text{Chloe}_i$  and escrow  $e_{i-1}$  all the way to Alice. Each  $\text{Chloe}_i$  forwards the “We are ready” message only after receiving promise  $G(d_i)$  from escrow  $e_i$ , so when Alice receives the “We are ready” message she can safely initiate the transfer.

## 4.5 Correctness of the protocol

Now we show that the protocol from Figure 3 is correct, in the sense that it satisfies the properties of Definition 1, when making the assumptions of Section 4.1. In doing so, we also calculate the values of the parameters  $d_i$  and  $a_i$ .

**Consistency.** To check that the protocol is consistent, in the sense that each participant can abide by it, we first of all invoke the assumption of bounded reaction speed, described in Section 4.1, and use that the constant  $\varepsilon$  assumed to exist in Section 4.1 is in fact the time-out value associated to most output states. This ensures that it is always possible to send messages in a timely manner.

The only remaining potential failure of consistency is when the protocol prescribes the transmission of a resource that it is not available. Assuming that the sending of promises and money is not an obstacle ( $\text{Chloe}$  has been selected, in part, for having this kind of money available), the only issue could be the sending of the certificate signed by Bob. For anyone but Bob this can only be done after receiving it first. However, a simple inspection of the automata of the escrows,  $\text{Chloe}_i$  and Alice shows that any transition sending the certificate is preceded by a transition receiving it. This establishes requirement C.

**Escrow security.** That escrows cannot lose money (requirement ES) follows immediately from the observation that an escrow spends the money only after receiving it. This follows from the order of the transitions in the automaton for the escrows.

**Honesty.** Although not part of Definition 1, we show that an escrow that issues a promise always keeps that promise, when abiding by the protocol. This property (H) will be used below to establish CS.

To show H, suppose the escrow  $e_i$  issues promise  $P(a_i)$ , and subsequently receives the certificate  $\chi$  from customer  $c_{i+i}$  at a time  $v < u + a_i$ , where  $u$  refers to the time promise  $P(a_i, \varepsilon)$  was issued. Then it is too soon for the time-out transition, so the transition labelled  $r(c_{i+1}, \chi)$  in the automaton of  $e_i$  will be taken, at time  $v$ . The automaton shows that  $s(c_{i+1}, \chi)$  will occur by time  $v + \varepsilon$ , thus fulfilling the promise.

To show that an escrow that issues promise  $G$  always keeps it, when abiding by the protocol, suppose the escrow  $e_i$  receives the money at a time  $w$ . Then the transition labelled  $r(c_i, \$)$  in the automaton of  $e_i$  will be taken, at time  $w$ . The automaton shows that either  $s(c_i, \$)$  will occur by time  $w + \varepsilon + a_i + \varepsilon$ , or  $s(c_i, \chi)$

will occur by time  $w + \varepsilon + a_i + \varepsilon$ . Thus, to guarantee that promise  $G$  is met, we need to choose  $d_i$  and  $a_i$  in such a way that

$$d_i \geq a_i + 2\varepsilon \quad (1)$$

for  $i = 0, \dots, n-1$ . In fact, making the promise as strong as possible yields  $d_i := a_i + 2\varepsilon$ . When this condition is met, we have established requirement H.

**Customer security and time-bounded termination.** We will check time-bounded termination (T) together with customer security (CS). To check requirement CS1, suppose that Alice will make the payment  $s(e_0, \$)$ , at time  $t$ . Then earlier she has received promise  $G(d_0)$  from escrow  $e_0$ . This promise ensures Alice that escrow  $e_0$  will send her either  $\$$  or  $\chi$  by its local time  $w + d_0$ , where  $w$  is the time Alice's payment is received. Consequently, by the reasoning of Section 4.1, using the assumption of bounded synchrony, Alice will receive either certificate  $\chi$  or her money back by time  $t + \Phi \cdot d_0 + 2 \cdot \Delta$ .

To check requirement CS2, suppose that Bob issues certificate  $\chi$ , at time  $x$ . Then earlier, at time  $t$ , he has received promise  $P(a_{n-1})$  from escrow  $e_{n-1}$ . Moreover,  $x < t + \varepsilon$ . For the promise to be meaningful, his certificate needs to arrive at  $e_{n-1}$  before time  $u + a_{n-1}$ , where  $u$  refers to the local time at  $e_{n-1}$  when the promise was issued. By the reasoning of Section 4.1, using the assumption of bounded synchrony, Bob's certificate will arrive at escrow  $e_{n-1}$  before time  $u + \Phi \cdot \varepsilon + 2 \cdot \Delta$ . Hence, we need to choose  $a_{n-1}$  in such a way that

$$a_{n-1} \geq \Phi \cdot \varepsilon + 2 \cdot \Delta. \quad (2)$$

When this requirement is met, the promise ensures Bob that escrow  $e_{n-1}$  will send him the money by its local time  $v + \varepsilon$ , where  $v$  is the time Bob's certificate is received by  $e_{n-1}$ . Consequently, Bob will receive payment by time  $x + \Phi \cdot \varepsilon + 2 \cdot \Delta$ .

To check requirement CS3, suppose that Chloe<sub>*i*</sub> will make the payment  $s(e_i, \$)$ , at time  $t_0$ . Then earlier, she has received promise  $G(d_i)$  from escrow  $e_i$  and promise  $P(a_{i-1})$  from escrow  $e_{i-1}$ , the latter at time  $t$ . Moreover,  $t_0 < t + \varepsilon$ . Promise  $G(d_i)$  ensures Chloe<sub>*i*</sub> that escrow  $e_i$  will send her either  $\$$  or  $\chi$  by its local time  $w + d_i$ , where  $w$  is the time Chloe<sub>*i*</sub>'s payment is received. Consequently,  $c_i$  will receive either certificate  $\chi$  or her money back by time  $t_0 + \Phi \cdot d_i + 2 \cdot \Delta$ .


Continuing with the case that she receives  $\chi$  rather than her money back, she will forward  $\chi$  to escrow  $e_{i-1}$  by time  $t_0 + \Phi \cdot d_i + 2 \cdot \Delta + \varepsilon$ , which is before  $t + \varepsilon + \Phi \cdot d_i + 2 \cdot \Delta + \varepsilon$ . Hence it arrives at  $e_{i-1}$  by its local time  $u + 2 \cdot \Phi \cdot \varepsilon + \Phi \cdot d_i + 4 \cdot \Delta$ , where  $u$  is the time promise  $P(a_{i-1})$  was issued. Here we use that  $\Delta$  is a valid upperbound on transition times by anyone's clock, and that there is no need to square  $\Phi$  in  $\Phi \cdot d_i$ , as also the clock skew between escrows  $e_i$  and  $e_{i-1}$  is bounded by  $\Phi$ . This calculation is illustrated by the message sequence diagram above. Since  $\chi$  needs to arrive at  $e_{i-1}$  before time  $u + a_{i-1}$  in order for promise  $P(a_{i-1})$  to be meaningful, we need to pick

$$a_{i-1} \geq 2 \cdot \Phi \cdot \varepsilon + \Phi \cdot d_i + 4 \cdot \Delta \quad (3)$$

for  $i = 1, \dots, n-1$ . When (3) holds, promise  $P(a_{i-1})$  ensures Chloe<sub>*i*</sub> that escrow  $e_{i-1}$  will send her the money by its local time  $v + \varepsilon$ , where  $v$  is the time the certificate is received by  $e_{i-1}$ . Consequently,  $c_i$  will receive  $\$$  by time  $t_0 + \Phi \cdot d_i + 4 \cdot \Delta + \varepsilon + \Phi \cdot \varepsilon$ . Hence, assuming (3), CS3 is guaranteed.

Choosing  $=$  for  $\geq$  in (1)–(3), we ensure requirement CS by solving these equations. In particular, for  $i = 0, \dots, n-1$ ,

$$a_i := \Phi^{n-1-i} \cdot (\Phi \cdot \varepsilon + 2 \cdot \Delta) + \sum_{j=i+1}^{n-1} 4 \cdot \Phi^{j-i-1} \cdot (\Phi \cdot \varepsilon + \Delta).$$



For  $i = n-1$  this follows by (2). Assume we have it for  $i+1$ . Then

$$a_{i+1} = \Phi^{n-2-i} \cdot (\Phi \cdot \varepsilon + 2 \cdot \Delta) + \sum_{j=i+2}^{n-1} 4 \cdot \Phi^{j-i-2} \cdot (\Phi \cdot \varepsilon + \Delta).$$

Now, applying (1) and (3), multiply by  $\Phi$  and add  $4 \cdot \Phi \cdot \varepsilon + 4 \cdot \Delta$ .

**Liveness.** It remains to check property L. Suppose that all parties abide by the protocol. By the reasoning in Section 4.4 we may assume that the action  $s(e_0, \$)@c_o$  (using the terminology of Section 4.3) of Alice sending money to her escrow will not take place until all actions  $s(c_i, G(d_i))@e_i$  and  $r(e_i, G(d_i))@c_i$  have occurred. In terms of Section 4.3 we show that in the remaining active phase of the protocol at least the prefix of the displayed sequence of actions until and including  $s(c_n, \$)@e_{n-1}$  will take place, in that order, and not interleaved with any other actions. When this happens, Bob must be in his third state, and the required action  $r(e_{n-1}, \$)@c_n$  will follow surely.

Towards a contradiction, let the initial behaviour of the active part of the protocol be a strict prefix of this sequence, where  $a$  is the first action in the sequence that does not occur as scheduled. A simple case analysis shows that when  $a$  is scheduled, in fact no action other than  $a$  is possible.

The action  $a$  cannot be of the form  $s(c_i, \$)@c_i$ , because when this action is due, customer  $c_i$  is in a state where this action must be taken within time  $\varepsilon$ . An exception is the case  $n = 0$ , but also here the action must be taken in within a finite amount of time (or there is nothing to verify).

The action  $a$  cannot be  $r(c_i, \$)@e_i$  either, because each message sent must arrive eventually, and the receiving party  $e_i$  is in its second state, and thus able to perform the receive action.

Similarly,  $a$  cannot be  $s(c_{i+1}, P(a_i))@e_i$ , as here the sender  $e_i$  must be in its third state.

Since all actions  $r(e_i, G(d_i))@c_i$  have already occurred,  $a$  cannot be of the form  $r(e_{i-1}, P(a_{i-1}))@c_i$ .

The case  $a = s(e_{n-1}, \chi)@c_n$  can be excluded, as here Bob must be in his second state.

If  $a = r(c_n, \chi)@e_{n-1}$ , then the recipient  $e_{n-1}$  must be in its fourth state and, due to the careful choice of  $a_{n-1}$  (see (2)), the time-out transition can not intervene. So that choice of  $a$  is excluded as well.

The argument against  $a = s(c_n, \$)@e_{n-1}$  is trivial.

## 5 Impossibility under partial synchrony of communications

When communications can experience arbitrarily long delays, it is not possible to expect from a protocol to terminate in an a priori known amount of time. If we want to perform cross-chain payments in a partially synchronous setting, it is therefore necessary to define a different class of cross-chain payment protocols. The obvious idea is to remove any time bound in the definition. But as we will show, this is not enough to make the problem solvable.

**Definition 2** (Eventually terminating cross-chain payment protocol with strong liveness guarantees). A cross-chain payment protocol is an *eventually terminating cross-chain payment protocol with strong liveness guarantees* it satisfies all the properties of Definition 1 except Property T which is replaced by the following:

T' **Eventual termination.** Each customer that abides by the protocol, and either makes a payment or issues a certificate, terminates *eventually*, provided her escrows abide by the protocol.

**Theorem 1.** *If communications are partially synchronous (assuming that customers may crash), then there is no eventually terminating cross-chain payment protocol with strong liveness guarantees.*

*Proof.* Assume an eventually terminating cross-chain payment protocol with strong liveness guarantees.

Consider a run  $r$  in which all participants abide by the protocol. It exists by property C. By property L Bob will be paid in this run, and by property T' all customers terminate. Since by properties ES and CS no participant makes a loss, Alice will not get her money back. Hence by property CS1 Alice will end up with the certificate  $\chi$ . Let  $c_i$  be the last customer that holds the certificate before it reaches Alice. This must be either Bob or one of the connectors. Let  $s$  be the state in which  $c_i$  is about to send on  $\chi$ .

The protocol may not prescribe that  $c_i$  has already received the money in state  $s$ . For then customer  $c_i$  could decide to keep the money as well as the certificate, while all other participants keep abiding by the protocol, which would violate properties T', ES or CS.

Now consider the following two runs of the system, that are the same until state  $s$ . In run  $r_1$  customer  $c_i$  never lets go of the certificate, nor sends out any other message past state  $s$ , while all other participants abide by the protocol; in run  $r_2$  all participants abide by the protocol, but  $c_i$ 's message with the certificate, and all subsequent messages from  $c_i$ , experience an extreme delay.

First assume that  $c_i$  is in fact Bob. By property T' run  $r_1$  reaches a state  $s'$  in which all customers other than Bob are terminated. Using properties ES and CS, in this state Alice ends up without the certificate, and thus with the money, and all customers  $Chloe_j$  play even. It follows that Bob never receives his money. Yet for all participants other than Bob, runs  $r_1$  and  $r_2$  are indistinguishable, so  $r_2$  will reach a similar state. This violates property CS2.

Now assume customer  $c_i$  is not Bob. So Bob has already issued the certificate. By property T' run  $r_1$  reaches a state  $s'$  in which all customers other than  $c_i$  are terminated. Using properties ES and CS, in this state Alice as well as Bob end up with the money, and all customers  $Chloe_j$  with  $j \neq i$  play even. It follows that  $Chloe_i$  loses her money. Yet for all participants other than  $Chloe_i$ , runs  $r_1$  and  $r_2$  are indistinguishable, and the delayed certificate sent by  $Chloe_i$  may arrive only after the system has reached state  $s'$ . This violates property CS3.  $\square$

## 6 Solution to a variant under partial synchrony

### 6.1 Eventually terminating cross-chain payment protocol with weak liveness guarantees

**Weak liveness guarantees.** Since it is impossible to design an *eventually terminating cross-chain payment protocol with strong liveness guarantees* under partial synchrony, we are going to define what is an *eventually terminating cross-chain payment protocol with weak liveness guarantees*, and show that it is possible to implement such a protocol.

In view of the impossibility proof given above, the *strong liveness* condition (L) is too strong. We would like to weaken it, and replace it by a (realistic and still desirable) property called *weak liveness* such that the problem becomes solvable. A similar situation exists in the atomic commit problem literature, for instance *weak non-triviality* defined by Guerraoui in [8] or condition *AC4* for atomic commit defined by Hadzilacos in [9]:

If all existing failures are repaired and no new failures occur for a sufficiently long period of time, then all processes will reach a decision.

**Abort certificate.** In the synchronous solution, we used timelocks to ensure that the money will not get stuck forever in escrow. Under partial synchrony, it is no longer pertinent to use timelocks. We need to replace them by a safe way to unlock the funds stored in an escrow.

We therefore modify the definition of the certificate  $\chi$ . Instead of having a single certificate simply signed by Bob, we have two more general certificates called *commit certificate*  $\chi_c$  and *abort certificate*  $\chi_a$ , that can never exist simultaneously.

**New definition of the problem.** When we take into account the two previous tweaks, we reach the following definition. We highlight in italics the difference with our previous definition of a *time-bounded cross-chain payment protocol*. In particular, we replace “Bob will not issue  $\chi$ ” by “Bob will receive  $\chi_a$ ”, and “Alice will receive  $\chi$ ” by “Alice will receive  $\chi_c$ ”.

**Definition 3** (Eventually terminating cross-chain payment protocol with weak liveness guarantees). A cross-chain payment protocol is an *eventually terminating cross-chain payment protocol with weak liveness guarantees* if it satisfies the following properties:

- C **Consistency.** For each participant in the protocol it is possible to abide by the protocol.
- CC **Certificate consistency.** *An abort and a commit certificate can never be issued both.*
- T' **Eventual termination.** Each customer that abides by the protocol terminates eventually, provided her escrows abide by the protocol.
- ES **Escrow security.** Each escrow that abides by the protocol does not lose money.
- CS' **Customer security.**
  - CS1' Upon termination, if Alice and her escrow abide by the protocol, Alice has either got her money back or received the *commit certificate*  $\chi_c$ .
  - CS2' Upon termination, if Bob and his escrow abide by the protocol, Bob has either received the money *or the abort certificate*  $\chi_a$ .
  - CS3' Upon termination, each connector that abides by the protocol has got her money back, provided her escrows abide by the protocol.
- L' **Weak liveness.** If all parties abide by the protocol, *and if the customers wait sufficiently long before and after sending money*, then Bob is eventually paid.

## 6.2 Transaction manager abstraction

The previous impossibility result shows that when there is no synchrony, it is hard for processes to agree on a uniform commitment decision (abort or commit). We are going to leverage the existing solutions to the classical consensus problem, and embed them in an abstraction called “transaction manager”, defined as follows:

**Definition 4** (Transaction manager). A *transaction manager*, called  $TM$ , is a process that can receive binary values from the set  $\{\text{COMMIT}, \text{ABORT}\}$  from a customer, and that can send certificates drawn from  $\{\chi_c, \chi_a\}$  to customers.

A transaction manager must verify the following properties:

- **TM-Termination.** If a customer proposes to ABORT, or if Bob proposes to COMMIT, then  $TM$  sends eventually a certificate  $\chi_c$  or  $\chi_a$  to every customer. If a customer proposes ABORT or COMMIT after  $TM$  has issued a certificate,  $TM$  will send a copy of that certificate to that customer.
- **TM-Consistency.**  $TM$  does not issue two different certificates.
- **TM-Commit-Validity.**  $TM$  can issue  $\chi_c$  only if Bob proposed COMMIT.
- **TM-Abort-Validity.**  $TM$  can issue  $\chi_a$  only if some customer proposed ABORT.

There are several ways of implementing a transaction manager.

- **Centralised transaction manager.** The transaction manager can be a centralised actor trusted by every customer.
- **Distributed transaction manager.** The transaction manager could be a collection of  $k$  parties (“validators”) appointed by the participants in the protocol. These validators could run the consensus algorithm for partial synchrony from Dwork, Lynch & Stockmeyer [5], or any equivalent algorithm.

This works when less than one third of these validators are unreliable. In this case, “sending a message to the TM” means sending it to each of these validators, and “the TM sending a decision” means strictly more than one third of the validators sending the jointly taken decision.

- **External decentralised transaction manager.** The transaction manager can be a decentralised data structure. For example, a smart contract running on a permissionless blockchain shared by every customer can be programmed to be a transaction manager.

In the following, we assume that we have such a transaction manager. Even if it is run by the customers, we do not specify the messages exchanged to run it: the transaction manager is a black-box embedding a consensus algorithm that is running off-protocol. Section 6.5 provides an example implementation.

### 6.3 A protocol responding to the problem




Figure 4: Automaton for escrow  $e_i$




Figure 5: Automaton for customer  $c_i, i \in \{1..n-1\}$




Figure 6: Automaton for customer  $c_0$  (Alice)

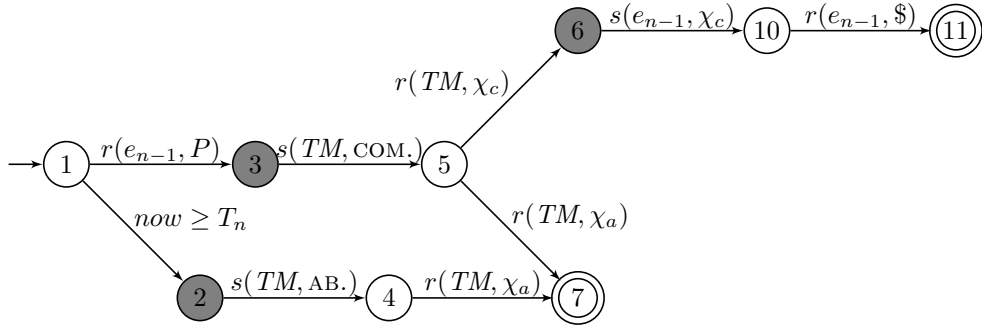


Figure 7: Automaton for customer  $c_n$  (Bob)

**Description of the protocol.** In this version of the protocol Chloe<sub>*i*</sub> awaits the two promises of her escrows, just like in the protocol of Figure 3, and then sends the money to her downstream escrow. Subsequently she awaits an abort or commit certificate from the transaction manager. If she gets a commit certificate, she cashes it in at her upstream escrow to obtain the money. If she gets an abort certificate instead, she cashes it in at her downstream escrow for a refund of the money she paid earlier. In case she loses patience before she gets the second promise, which happens at a time  $T_i$  specific for Chloe<sub>*i*</sub>, she simply quits. In case she loses patience after she has invested the money but before she gets any certificate, the timeout transition occurs, and she sends an abort proposal to the transaction manager. The latter replies on this with either an abort or a commit certificate, and she cashes those in as above.

The tags  $G$  and  $P$  can thus be understood as promises that would tell:

- $G$ : “I guarantee that if I receive \$ from you, then if you send me  $\chi_a$  I will send you \$”.
- $P$ : “I promise that if you send me  $\chi_c$  I will send \$ to you”.

The automaton for Alice is just a simplified version of the one for Chloe, and the one for the escrows is trivial. For Bob, the important modification is that he alerts the TM with a COMMIT message when the protocol is ready for this. Moreover, in case Bob loses patience before receiving any promise, he send a ABORT message to the TM, so that he receives the abort certificate in response.

## 6.4 Proof of correctness

Let us call  $P$  the protocol defined by the above ANTA. Section 4.4 (Initialisation) applies to  $P$  as well, and we assume that the appropriate modifications are made.

**Theorem 2.** *Protocol  $P$  is an eventually terminating cross-chain payment protocol with weak liveness guarantees.*

*Proof.* The properties to prove are close to the time-bounded cross-chain payment protocol, and the proof is similar. We split the proof in the following lemmas: Lemma 1, 2, 3, 4, 5 and 6.  $\square$

**Lemma 1** (Consistency). *For each participant in  $P$  it is possible to abide by  $P$ .*

*Proof.* We have to ensure that each participant will be able to follow the transitions after a grey or black state. The only way this would not be possible would be when the protocol asks to transmit a resource that is not available. It is always possible to send tags and money, but we need to verify that any sending of a certificate is preceded by the receipt of this certificate—except for the issuer of the certificate. Such a property is clear after inspection of the automata of escrows and customers.  $\square$

**Lemma 2** (Certificate consistency). *An abort and a commit certificate can never be issued both.*

*Proof.* This is a consequence of Property **TM-Consistency** of Definition 4.  $\square$

**Lemma 3** (Eventual termination). *Each customer that abides by  $P$  will terminate eventually, provided her escrows abide by the protocol.*



*Proof.* Thanks to Lemma 1, in order to show that a terminating state will be reached, it is sufficient to prove that every input state will be left eventually. We thus establish Lemma 3 by showing that Chloe<sub>*i*</sub> and Alice cannot get stuck in states 1, 2, 5, 9, 10 and 12, and Bob cannot get stuck in states 1, 4, 5 and 10.

Since her downstream escrow will surely leave state 1, and thus send the message  $G$ , it follows that Chloe<sub>*i*</sub>, and Alice, will receive this message, and thereby leave state 1.

Chloe<sub>*i*</sub> will leave state 2 at time  $T_i$  at the latest, or immediately when reaching this state after time  $T_i$ . By the same reasoning, Bob will not get stuck in state 1.

Chloe<sub>*i*</sub> and Alice will leave state 5 at time  $T_i$  at the latest, or immediately when reaching this state after time  $T_i$ . She will send an ABORT message to  $TM$  to reach state 9, and by the property **TM-Termination** of Definition 4,  $TM$  will eventually reply to her with  $\chi_a$  or  $\chi_a$ . Hence she will leave state 9. By the same reasoning, Bob will not get stuck in state 5.

To reach state 4, Bob sends an ABORT message to  $TM$ . By properties **TM-Termination** and **TM-Commit-Validity** of Definition 4,  $TM$  will eventually reply to him with  $\chi_a$ . Hence he will leave state 4.

Now assume Customer  $c_{i+1}$  (Chloe or Bob) reaches state 10. Then Customer  $c_{i+1}$  has already received promise  $P$  from Escrow  $i$ , and thus Escrow  $i$  must have send this promise, thereby reaching state 4. To reach state 10, Customer  $c_{i+1}$  sends certificate  $\chi_c$  to Escrow  $i$ . By Lemma 2, the TM never issues certificate  $\chi_a$ , so Customer  $c_i$  cannot send it to Escrow  $i$ . It follows that Escrow  $i$  will reach state 5, and send the money to Customer  $c_{i+1}$ . Hence Customer  $c_{i+1}$  will leave state 10 and reach the terminating state 11.

Finally assume Customer  $c_i$  (Alice or Chloe) reaches state 12. Then Customer  $c_i$  has already received promise  $G$  from Escrow  $i$ , and thus Escrow  $i$  must have send this promise, thereby reaching state 2. After receiving promise  $G$ , Customer  $c_i$  has send the money to Escrow  $i$ , so Escrow  $i$  will have reached state 3, and hence also state 4. To reach state 12, Customer  $c_i$  sends certificate  $\chi_a$  to Escrow  $i$ . By Lemma 2, the TM never issues certificate  $\chi_c$ , so Customer  $c_{i+1}$  cannot send it to Escrow  $i$ . It follows that Escrow  $i$  will reach state 6, and send the money to Customer  $c_i$ . Hence Customer  $c_i$  will leave state 12 and reach the terminating state 13.  $\square$

**Lemma 4** (Escrow-security). *Each escrow that abides by  $P$  will not loose money.*

*Proof.* The result is immediate: any transition where an escrow sends money has been preceded by a transition where it receives the money. An escrow's balance cannot become negative if it follows the protocol.  $\square$

**Lemma 5** (Customer-security).

1. Upon termination, if Alice and her escrow abide by  $P$ , Alice has either got her money back or received the commit certificate  $\chi_c$ .
2. Upon termination, if Bob and his escrow abide by  $P$ , Bob has either received the money or the abort certificate  $\chi_a$ .
3. Upon termination, each connector that abides by  $P$  has got her money back, provided her escrows abide by  $P$ .

*Proof.*

1. If Alice terminates in state 13, she has got her money back in the last transition. If she terminates in state 6, she has got the certificate  $\chi_c$  in the last transition.
2. The result is similar for Bob: if he terminates in state 7, he has received  $\chi_a$ . Otherwise, he terminates in state 11 and has been paid correctly.
3. To reach termination, Chloe<sub>*i*</sub>,  $i \in \{1..n-1\}$  has either never spend the money (termination state 3), or received \$ either from  $e_{i-1}$  (termination state 11) or from  $e_i$  (termination state 13). In both these cases, she has got her money back.  $\square$

**Lemma 6** (Weak liveness). *If all participants abide by  $P$ , and if the customers wait sufficiently long before and after sending money, then Bob will be paid.*

*Proof.* Let us suppose that all the participants abide by  $P$ . Suppose that for all  $i \in \llbracket 0, n-1 \rrbracket$ ,  $T_i$  is large enough for Alice, Bob and every connector to never take any time-out transition. Instead, Bob will be the first customer to call the transaction manager in his transition from state 3 to 5.

Using the notation of Section 4.3, the active part of the protocol—after the exchange of the tags  $G$ —must start with the following sequence of actions, executed in this order:

$$\begin{array}{ccccccc}
s(e_0, \$)@c_0 & r(c_0, \$)@e_0 & s(c_1, P)@e_0 & r(e_0, P)@c_1 & & & \\
s(e_1, \$)@c_1 & r(c_1, \$)@e_1 & s(c_2, P)@e_1 & r(e_1, P)@c_2 & & & \\
\dots & & & & & & \\
s(e_{n-1}, \$)@c_{n-1} & r(c_{n-1}, \$)@e_{n-1} & s(c_n, P)@e_{n-1} & r(e_{n-1}, P)@c_n & s(TM, COM.)@c_n & & 
\end{array}$$

By the **TM-Abort-Validity** property of Definition 4, since the only proposal was COMMIT,  $TM$  will issue the certificate  $\chi_c$ . In particular, Bob will give it to  $e_{n-1}$  and receive the payment in exchange.  $\square$

Interestingly, nothing in the above proofs depends in any way on the assumption of partially synchronous communication. The only place where this is needed is for the implementation of the transaction manager  $TM$ —see Section 6.5. In case one is content with a centralised transaction manager as described in Section 6.2, our protocol works correctly also when assuming communication to be asynchronous.

## 6.5 Implementation of a decentralised transaction manager

As an example, in this section we provide an explicit implementation of a transaction manager. It is a wrapper, expressed in pseudo-code, around a binary Byzantine consensus algorithm, such as the one from [5], which is treated as a black box. It follows that we do not need a shared blockchain to achieve our cross-chain payment protocol.

Suppose that we have  $m$  *validators*, which are agents running a consensus algorithm. The validators communicate with each other by exchanging messages. We suppose that a certain number  $f$  of validators can be *faulty*, in the sense that we allow arbitrary (Byzantine) behaviour. All other validators are assumed to abide by the protocol defining a consensus algorithm.

**Definition 5** (Binary Byzantine Consensus). A binary Byzantine consensus algorithm is an algorithm in which every validator can *propose* a binary value (*i.e.* in  $\{0, 1\}$ ) and *decide* a binary value. Assuming that every non-faulty validator proposes a binary value, the following properties must be ensured:

- **BBC-Termination.** Every non-faulty process eventually decides on a binary value.
- **BBC-Agreement.** No two non-faulty processes decide on different binary values.
- **BBC-Validity.** If all non-faulty processes propose the same value, no other value can be decided.

**Theorem 3** ([5]). *Assuming partially synchronous communication, a binary Byzantine consensus algorithm exists when  $f < m/3$ .*

Using such an algorithm as a black box, we now implement a  $TM$ . Our validators can either be customers, like Alice, Chloe and Bob, or external parties. If the set of validators is included in or equal to the set of customers, we can talk of an *internal decentralised transaction manager*. Our  $TM$  implementation is only valid when assuming partially synchronous communication, and  $f < m/3$ . Each customer can communicate with every validator.

**Reliable broadcast call.** To call the transaction manager, a customer reliably broadcasts a message to all validators. Here a *reliable broadcast* is a protocol described by Bracha in [4]. It is guaranteed to terminate, even in a setting with asynchronous communication, provided less than one-third of all the broadcast recipients is faulty — the rest abiding by the protocol. It guarantees that if the sender abides by the protocol, all recipients will receive the message sent. Moreover, even if the sender is faulty, either all correct recipients agree on the same value sent, or none of them accepts any value as having been sent [4].

If a faulty customer sends a call with different values to different validators, or sends something to some validators and nothing to others, the reliable broadcast primitive will filter these messages out. In particular, if a validator receives a call from a customer then eventually every validator will receive this call from this customer.

---

**Algorithm 1** *BFT-TM* algorithm for validator  $v_k$ ,  $k \in [1..m]$ .

---

- 1: When validator  $k$  has not proposed, nor decided, a value so far:
  - 2:   when receive ABORT from a customer  $c_i$ ,  $i \in [0..n]$ :
  - 3:     propose(0) to BBC
  - 4:   when receive COMMIT from  $c_n$ :
  - 5:     propose(1) to BBC
  
  - 6: When validator  $k$  decides value  $v$  (by running BBC):
  - 7:   broadcast( $\sigma_i(v)$ ) to all validators
  - 8:   await receipt of  $\sigma_j(v)$  for all  $j$  in a certain  $J \subseteq [1..m]$  such that  $|J| > f$
  - 9:    $\chi := \{\sigma_j(v), j \in J\}$
  - 10:   broadcast( $\chi$ ) to all customers
  
  - 11: When validator  $k$  has decided a value  $v$ :
  - 12:   when receive ABORT or COMMIT from a customer  $c_i$ ,  $i \in [0..n]$ :
  - 13:     send  $\chi$  to customer  $c_i$ .
- 

**Certificate implementation.** With  $\sigma_k(v)$  we denote the value  $v \in \{0, 1\}$  cryptographically signed by validator  $k$ . Such a signed message models the decision ABORT (if  $v = 0$ ) or COMMIT ( $v = 1$ ) taken by validator  $k$ . Since up to  $f$  validators may be unreliable, a valid certificate is a message that contains (for instance as attachments) more than  $f$  copies of the same decision, taken by different validators  $k$ . We model such certificates as sets. Hence:

- $\chi_c$  is any set of at least  $f+1$  messages  $\sigma_k(1)$  signed by at least  $f+1$  different validators  $k$ .
- $\chi_a$  is any set of at least  $f+1$  messages  $\sigma_k(0)$  signed by at least  $f+1$  different validators  $k$ .

Such a certificate is verifiable non-interactively by a third party such as any customer or escrow. Asking for  $f+1$  signatures ensures that at least one correct validator has issued this certificate, and in particular will guarantee property CC (certificate consistency). Of course a real implementation will not rely on simple signatures of the string "0" or "1" because of the possibility of replay attacks.

**TM implementation.** When our protocol prescribes  $s(TM, \text{AB.})$ , resp.  $s(TM, \text{COM.})$ , this is implemented as reliably broadcasting ABORT, resp. COMMIT, to all validators. The transition  $r(TM, \chi)$  denotes the receipt of certificate  $\chi$  from one of the validators. The behaviour of  $TM$  is described as Algorithm 1.

**Theorem 4** (BFT-TM correctness). *The BFT-TM algorithm implements a TM as defined in Definition 4.*

*Proof.* The correctness properties of a transaction manager derive almost immediately from the correctness properties of a binary Byzantine consensus algorithm.

1. **TM-Termination.** Let us suppose that a customer sends an abort proposal to  $TM$ , i.e. to each and every validator, or that Bob proposes to commit. Then every correct validator starts participating in BBC with the initially proposed binary value 0 or 1, respectively. By **BBC-Termination**, every correct validator eventually passes line 6 of the BFT-TM algorithm. By assumption on the number of correct validators, every validator eventually receives at least  $f+1$  signatures on this value, thereby forming a certificate that is sent to every customer.

If a customer proposes ABORT or COMMIT after  $TM$  has issued a certificate, each correct validator will send a copy of that certificate to that customer by lines 11–13 of the algorithm.

2. **TM-Consistency.** By contradiction, let us suppose that two customers receive different certificates. Since a certificate contains at least  $f+1$  signatures, a correct validator has broadcast  $\sigma_i(0)$  and another correct validator has broadcast  $\sigma_j(1)$ . Line 6 of BFT-TM shows that the value signed and broadcast has been decided by BBC. This is a contradiction with **BBC-Agreement**.

3. **TM-Commit-Validity.** TM-Commit-Validity says that if Bob does not propose commit, then  $TM$  cannot issue  $\chi_c$ . If Bob does not propose commit, then no correct validator will ever propose 1 because of lines 1–5 of BFT-TM. Now **BBC-Validity** implies that no correct validator can decide 1, and consequently the commit certificate cannot be issued by  $TM$ .
4. **TM-Abort-Validity.** TM-Abort-Validity says that if no customer proposes to abort, then  $TM$  cannot issue  $\chi_a$ . If no customer proposes to abort, then no correct validator will ever propose 0 because of lines 1 – 5 of BFT-TM. The final argument is the same as above.  $\square$

## 7 Cross-chain deals versus cross-chain payments

In a preprint by Herlihy, Liskov and Shrira [11], a *cross-chain deal* is given by a matrix  $M$  where  $M_{i,j}$  is listing an asset to be transferred from party  $i$  to party  $j$ . A cross-chain deal can also be represented as a directed graph, where each vertex represents a party, and each arc a transfer; there is an arc from  $i$  to  $j$  labelled  $v$  iff  $M_{i,j} = v$  and  $v \neq 0$ .

They present two solutions to the problem of implementing such a deal, while aiming to ensure:

- **Safety.** For every protocol execution, every compliant party ends up with an acceptable payoff.
- **Weak liveness.** No asset belonging to a compliant party is escrowed forever.
- **Strong liveness.** If all parties are compliant and willing to accept their proposed payoffs, then all transfers happen.

Here a payoff is deemed *acceptable* to a party  $i$  in the deal if party  $i$  either receives all assets  $M_{j,i}$  while parting with all assets  $M_{i,j}$ , or if party  $i$  loses nothing at all; moreover, any outcome where she loses less and/or gains more than an acceptable outcome is also acceptable.

Each entry  $M_{i,j}$  contains a type of asset and a magnitude—for instance “5 bitcoins”. For each type of asset a separate blockchain is assumed that acts as escrow. The programming of these blockchains is assumed to be open source, so that all parties can convince themselves that all escrows abide by the protocol. Unlike in our problem statements, this condition is taken for granted and not stated explicitly in the problem description. With this in mind, **Weak liveness** corresponds with our **Eventual termination** of Definition 3, while **Safety** is the counterpart of our **Customer security**. Our requirement of **Escrow security** is left implicit in [11]; since blockchains do not possess any assets to start with, they surely cannot lose them. Finally, **Strong liveness** is the counterpart of our **Strong liveness** property of Definition 1.

Herlihy, Liskov and Shrira [11] offer a timelock commit protocol that requires synchrony, and assures all three of the above correctness properties. They also offer a certified blockchain commit protocol that requires partial synchrony and a certified blockchain, and ensures **Safety** and **Weak liveness**; no protocol can offer **Strong liveness** in a partially synchronous environment. For both protocols the correctness is proven for so-called *well-formed* cross-chain deals: those whose associated directed graph is strongly connected.

The cross-chain payment cannot be seen as a special kind of cross-chain deal. In first approximation, a cross-chain payment looks like a non-well-formed deal of the form

$$\left[ \begin{array}{cccc} 0 & \$ & & \\ & 0 & \$ & (0) \\ & & 0 & \ddots \\ & & & \ddots & \$ \\ (0) & & & & 0 & \$ \\ & & & & & 0 \end{array} \right] \cong c_0 \xrightarrow{\$} c_1 \xrightarrow{\$} \dots \xrightarrow{\$} c_n .$$

However, this representation abstracts from the certificate  $\chi$  that plays an essential role in the statement of

the time-bounded cross-chain payment problem. Factoring in  $\chi$ , an alternative representation would be

$$\begin{bmatrix} 0 & \$ & & & \\ & 0 & \$ & & (0) \\ & & & 0 & \ddots \\ & & & & \ddots & \$ \\ & (0) & & & & 0 & \$ \\ \chi & & & & & & 0 \end{bmatrix} \quad \text{or} \quad \begin{bmatrix} 0 & \$ & & & \\ \chi & 0 & \$ & & (0) \\ & & \chi & 0 & \ddots \\ & & & & \ddots & \$ \\ & (0) & & \chi & 0 & \$ \\ & & & & \chi & 0 \end{bmatrix}.$$

However, these solutions presume a shared blockchain between Alice and Bob for the transfer of the certificate; this runs counter to the problem description of cross-chain payments.

Conversely, neither is there a reduction from the cross-chain deal problem to the cross-chain payment problem. A deal presented by a cyclic graph can be represented as a cross-chain payment where Alice and Bob are identified. For instance, an atomic swap between two customers A and C can be expressed as a cross-chain payment with three customers:

$$\begin{bmatrix} 0 & a \\ b & 0 \end{bmatrix} \cong A \xrightarrow{a} C \xrightarrow{b} B = A.$$

However, this idea does not generalise to well-formed cross-chain deals in general. Since every strongly connected graph can be represented as a single cycle with repeated elements, there is an obvious candidate reduction of such deals to cross-chain payments, simply by identifying suitable intermediaries  $\text{Chloe}_i$  and  $\text{Chloe}_j$ . However, this reduction does not preserve the safety property of cross-chain deals; for when the deal goes through for  $\text{Chloe}_j$  but is aborted for  $\text{Chloe}_i$ , the resulting outcome is not (necessarily) acceptable for the unified participant  $\text{Chloe}_{\{i,j\}}$ .

## 8 Conclusion

We formalised the problem of cross-chain payment with success guarantees. We show that there is no solution to the existing variant of this problem without assuming synchrony, and offer such a solution—one that works even in the presence of clock skew. We then relax the liveness guarantee of this problem in order to propose a solution that works in a partially synchronous setting. This new problem differs from existing ones in that it prevents all participants from always aborting, hence guaranteeing success when possible. Besides the new problem statements and our impossibility result, an interesting aspect of our work is to relate recent blockchain problems, like interledger payments, to the classic problem of transaction commit, and to offer Byzantine fault tolerant solutions to these.

## Acknowledgements

This research is supported under Australian Research Council Discovery Projects funding scheme (project number 180104030) entitled “Taipan: A Blockchain with Democratic Consensus and Validated Contracts” and Australian Research Council Future Fellowship funding scheme (project number 180100496) entitled “The Red Belly Blockchain: A Scalable Blockchain for Internet of Things”.

## References

- [1] Rajeev Alur & David L. Dill (1994): *A Theory of Timed Automata*. *Theor. Comput. Sci.* 126(2), pp. 183–235, doi:10.1016/0304-3975(94)90010-8.
- [2] Antonio Fernández Anta, Chryssis Georgiou & Nicolas C. Nicolaou (2018): *Atomic Appends: Selling Cars and Coordinating Armies with Multiple Distributed Ledgers*. *CoRR* abs/1812.08446. Available at <http://arxiv.org/abs/1812.08446>.
- [3] Johan Bengtsson, Kim Guldstrand Larsen, Fredrik Larsson, Paul Pettersson & Wang Yi (1996): *UP-PAAL - a Tool Suite for Automatic Verification of Real-Time Systems*. In Rajeev Alur, Thomas A.

- Henzinger & Eduardo D. Sontag, editors: *Hybrid Systems III: Verification and Control*, Proceedings of the DIMACS/SYCON Workshop on *Verification and Control of Hybrid Systems*, October 1995, LNCS 1066, Springer, pp. 232–243, doi:10.1007/BFb0020949.
- [4] Gabriel Bracha (1987): *Asynchronous Byzantine Agreement Protocols*. *Inf. Comput.* 75(2), pp. 130–143, doi:10.1016/0890-5401(87)90054-X.
- [5] Cynthia Dwork, Nancy A. Lynch & Larry J. Stockmeyer (1988): *Consensus in the presence of partial synchrony*. *J. ACM* 35(2), pp. 288–323, doi:10.1145/42282.42283.
- [6] Peter Gazi, Aggelos Kiayias & Dionysis Zindros (2018): *Proof-of-Stake Sidechains*. *IACR Cryptology ePrint Archive* 2018, p. 1239. Available at <https://eprint.iacr.org/2018/1239>.
- [7] Jim Gray & Leslie Lamport (2006): *Consensus on Transaction Commit*. *ACM Trans. Database Syst.* 31(1), pp. 133–160, doi:10.1145/1132863.1132867.
- [8] Rachid Guerraoui (1995): *Revisiting the relationship between non-blocking atomic commitment and consensus*. In Jean-Michel Hélary & Michel Raynal, editors: *Proceedings of the 9th International Workshop on Distributed Algorithms (WDAG)*, LNCS 972, Springer, pp. 87–100, doi:10.1007/BFb0022140.
- [9] Vassos Hadzilacos (1990): *On the Relationship Between the Atomic Commitment and Consensus Problems*. In: *Proceedings of the Asilomar Workshop on Fault-Tolerant Distributed Computing*, Springer, pp. 201–208. Available at <http://dl.acm.org/citation.cfm?id=645425.652317>.
- [10] Maurice Herlihy (2018): *Atomic Cross-Chain Swaps*. In: *Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing, PODC*, pp. 245–254, doi:10.1145/3212734.3212736.
- [11] Maurice Herlihy, Barbara Liskov & Liuba Shrira (2019): *Cross-chain Deals and Adversarial Commerce*. *CoRR* abs/1905.09743. Available at <http://arxiv.org/abs/1905.09743>.
- [12] Joshua Lind, Oded Naor, Ittay Eyal, Florian Kelbert, Peter Pietzuch & Emin Gün Sirer (2018): *Teechain: Reducing storage costs on the blockchain with offline payment channels*. In: *Proceedings of the 11th ACM International Systems and Storage Conference, ACM*, pp. 125–125, doi:10.1145/3211890.3211904.
- [13] Nancy Lynch (2003): *Input/Output Automata: Basic, Timed, Hybrid, Probabilistic, Dynamic,...* In Roberto M. Amadio & Denis Lugiez, editors: *Proceedings of the Conference on Concurrency Theory (CONCUR)*, LNCS 2761, pp. 191–192, doi:10.1007/978-3-540-45187-7\_12.
- [14] Ron van der Meyden (2019): *On the specification and verification of atomic swap smart contracts (extended abstract)*. In: *Proceedings of the IEEE International Conference on Blockchain and Cryptocurrency, ICBC*, pp. 176–179, doi:10.1109/BLOC.2019.8751250.
- [15] R. Milner (1989): *Communication and Concurrency*. Prentice Hall, Englewood Cliffs.
- [16] Satoshi Nakamoto (2008): *Bitcoin: a peer-to-peer electronic cash system*.
- [17] Joseph Poon & Thaddeus Dryja (2016): *The bitcoin lightning network: Scalable off-chain instant payments*. Available at <https://lightning.network/lightning-network-paper.pdf>.
- [18] Alejandro Ranchal-Pedrosa & Vincent Gramoli (2019): *Platypus: Offchain Protocol Without Synchrony*. In: *Proceeding of the 18th IEEE International Symposium on Network Computing and Applications, NCA*. Available at <https://arxiv.org/abs/1907.03730>.
- [19] Stefan Thomas & Evan Schwartz (2015): *A Protocol for Interledger Payments*. Available at <https://interledger.org/interledger.pdf>.
- [20] Gavin Wood (2016): *Polkadot: Vision for a heterogeneous multi-chain framework*. *White Paper*. Available at <https://polkadot.network/PolkaDotPaper.pdf>.
- [21] Alexei Zamyatin, Dominik Harz, Joshua Lind, Panayiotis Panayiotou, Arthur Gervais & William J. Knottenbelt (2019): *XCLAIM: Trustless, Interoperable, Cryptocurrency-Backed Assets*. In: *2019 IEEE Symposium on Security and Privacy, SP 2019*, pp. 193–210, doi:10.1109/SP.2019.00085.