

# SoK: Communication Across Distributed Ledgers

Alexei Zamyatin<sup>1,2</sup>, Mustafa Al-Bassam<sup>3</sup>, Dionysis Zindros<sup>4</sup>,  
Eleftherios Kokoris-Kogias<sup>5,9</sup>, Pedro Moreno-Sanchez<sup>6</sup>, and Aggelos Kiayias<sup>7,8</sup>,  
and William J. Knottenbelt<sup>1</sup>

<sup>1</sup> Imperial College London

<sup>2</sup> Interlay

<sup>3</sup> University College London

<sup>4</sup> University of Athens

<sup>5</sup> IST Austria

<sup>6</sup> Novi Research

<sup>7</sup> IMDEA Software Institute

<sup>8</sup> IOHK

<sup>9</sup> University of Edinburgh

**Abstract.** Since the inception of Bitcoin, a plethora of distributed ledgers differing in design and purpose has been created. While by design, blockchains provide no means to securely communicate with external systems, numerous attempts towards trustless cross-chain communication have been proposed over the years. Today, cross-chain communication (CCC) plays a fundamental role in cryptocurrency exchanges, scalability efforts via sharding, extension of existing systems through sidechains, and bootstrapping of new blockchains. Unfortunately, existing proposals are designed ad-hoc for specific use-cases, making it hard to gain confidence in their correctness and composability.

We provide the first systematic exposition of cross-chain communication protocols. We formalize the underlying research problem and show that CCC is *impossible without a trusted third party*, contrary to common beliefs in the blockchain community. With this result in mind, we develop a framework to design new and evaluate existing CCC protocols, focusing on the inherent trust assumptions thereof, and derive a classification covering the field of cross-chain communication to date. We conclude by discussing open challenges for CCC research and the implications of interoperability on the security and privacy of blockchains.

## 1 Introduction

Since the introduction of Bitcoin [152] as the first decentralized ledger currency in 2008, the topic of blockchains (or distributed ledgers) has evolved into a well-studied field both in industry and academia. Nevertheless, developments are still largely driven by community effort, resulting in a plethora of blockchain-based digital currencies being created. Taking into account the heterogeneous nature of

these systems in terms of design and purpose, it is unlikely there shall emerge a “coin to rule them all”, yielding interoperability an important research problem.

Today, cross-chain communication is found not only in research on cryptocurrency transfers and exchanges [20, 19, 108, 107, 187], but is a critical component of scalability solutions such as sharding [129, 34, 185, 33, 36], feature extensions via sidechains [47, 125, 97, 136], as well as bootstrapping of new systems [116, 169, 118]. In practice, over \$1bn worth of Bitcoin has been moved to other blockchains [27], and numerous competing interoperability projects, attempting to unite independent systems, have been deployed to practice [176, 180, 133, 167, 110, 178, 21], creating a multi-million dollar industry.

However, in spite of the vast number of use cases and solution attempts, the underlying problem of cross-chain communication has neither been clearly defined, nor have the associated challenges been studied or related to existing research. Early attempts to overview this field offer iterative summaries of mostly community-lead efforts [68, 165, 113], or focus on a subset of this space, such as atomic swaps [52, 150], and support our study. Belchior et al. [48] provide another, more recent, iterative overview of cross-chain projects, yet without clear taxonomy or classification.

**This work.** This Systematization of Knowledge (SoK) offers a comprehensive guide for designing protocols bridging the numerous distributed ledgers available today, aiming to facilitate clearer communication between academia, community, and industry. The contributions of this work are thereby twofold:

- We formalize the underlying problem of Correct Cross-Chain Communication (CCC) (Section 2), relating CCC to existing research and outlining a generic CCC protocol encompassing existing solutions. We then relate CCC to the Fair Exchange problem and show that contrary to common beliefs in the blockchain community, CCC is *impossible without a trusted third party* (Section 3).

- With the impossibility result in mind, we present a framework to design new and evaluate existing CCC protocols, focusing on the inherent trust assumptions thereof (Sections 4). We apply this framework to classify the field of CCC protocols to date (Section 5), highlighting similarities and key differences. Finally, we outline general observations on current developments, provide an outlook on the challenges of CCC research, and discuss the implications of interoperability on the security and privacy of blockchains (Section 6).

## 2 The Cross-Chain Communication Problem

In this section, we relate cross-chain communication to existing research, introduce the model for interconnected distributed ledgers, provide a formal definition of the Correct Cross-Chain Communication (CCC) problem, and sketch the main phases of a generic CCC protocol.

### 2.1 Historical Background: Distributed Databases

The need for communication among distributed processes is fundamental to any distributed computing algorithm. In databases, to ensure the atomicity of a dis-

tributed transaction, an agreement problem must be solved among the set of participating processes. Referred to as the Atomic Commit problem (AC) [56], it requires the processes to agree on a common outcome for the transaction: commit or abort. If there is a strong requirement that every correct process should eventually reach an outcome despite the failure of other processes, the problem is called Non-Blocking Atomic Commit (NB-AC) [46]. Solving this problem enables correct processes to relinquish locks without waiting for crashed processes to recover. As such, we can relate the core ideas of communication across distributed ledgers to NB-AC. The key difference hereby lies within the security model of the interconnected systems. While in classic distributed databases all processes are expected to *adhere to protocol rules* and, in the worst case, may crash, distributed ledgers, where consensus is maintained by a committee, must also consider and handle *Byzantine failures*.

## 2.2 Distributed Ledger Model

We use the terms *blockchain* and *distributed ledger* as synonyms and introduce some notation, based on [97] with minor alterations.

**Ledgers and State Evolution.** When speaking of CCC, we consider the interaction between two distributed systems  $X$  and  $Y$ , which can have distinct consensus participants and may employ different agreement protocols. Thereby, it is assumed the majority<sup>1</sup> of consensus participants in both  $X$  and  $Y$  are honest, namely, that they follow the designated protocol. The data structures underlying  $X$  and  $Y$  are *blockchains* (or *chains*), i.e., append-only sequences of blocks, where each block contains a reference to its predecessor(s). We denote a ledger as  $L$  ( $L_x$  and  $L_y$  respectively) and define its *state* as the dynamically evolving sequence of included transactions  $\langle TX_1, \dots, TX_n \rangle$ . We assume that the evolution of the ledger state progresses in discrete *rounds* indexed by natural numbers  $r \in \mathbb{N}$ . At each round  $r$ , a new set of transactions (included in a newly generated block) is written to the ledger  $L$ . We use  $L^P[r]$  to denote the state of ledger  $L$  at round  $r$ , i.e., after applying all transactions *written* to the ledger since round  $r - 1$ , according to the view of some party  $P$ . A transaction can be written to  $L$  only if it is consistent with the system’s consensus rules, given the current ledger state  $L^P[r]$ . This consistency is left for the particular system to define, and we describe it as a free predicate  $\text{valid}(\cdot)$  and we write  $\text{valid}(TX, L_x^P[r])$  to denote that  $TX$  is valid under the consensus rules of  $L_x$  at round  $r$  according to the view of party  $P$ . To denote that a transaction  $TX$  has been *included* in / successfully written to a ledger  $L$  as position  $r$  we write  $TX \in L^P[r]$ . While the ordering of transactions in a block is crucial for their validity, for simplicity, we omit the position of transactions in blocks and assume correct ordering implicitly.

**Notion of Time.** The state evolution of two distinct ledgers  $L_x$  and  $L_y$  may progress at different *time* intervals: In the time that  $L_x$  progresses *one* round,  $L_y$  may, for example, progress *forty* rounds (e.g., as in the case of Bitcoin [152] and

<sup>1</sup> In case of Proof-of-Work or Proof-of-Stake blockchains, the majority pertains to computational power [152] or stake [123] respectively.

Ethereum [67]). To correctly capture the ordering of transactions across  $L_x$  and  $L_y$ , we define a clock function  $\tau$  which maps a given round on any ledger to the time on a global, synchronized clock  $\tau : r \rightarrow t$ . We assume that the two chains are nevertheless synchronized and that there is no clock drift between them. We use this conversion implicitly in the rest of this paper. For conciseness, we will use the notation  $L^P[t]$  to mean the ledger state in the view of party  $P$  at the round  $r = \tau^{-1}(t)$  which corresponds to time  $t$ , namely  $L^P[\tau^{-1}(t)]$ .

**Persistence and Liveness.** Each participant  $P$  adopts and maintains a local ledger state  $L^P[t]$  at time  $t$ , i.e., her current view of the ledger. The views of two distinct participants  $P$  and  $Q$  on the same ledger  $L$  may differ at time  $t$  (e.g., due to network delay):  $L^P[t] \neq L^Q[t]$ . However, eventually, all honest parties in the ledger will have the same view. This is captured by the persistence and liveness properties of distributed ledgers [93]:

**Definition 1 (Persistence).** *Consider two honest parties  $P, Q$  of a ledger  $L$  and a persistence (or “depth”) parameter  $k \in \mathbb{N}$ . If a transaction  $\text{TX}$  appears in the ledger of party  $P$  at time  $t$ , then it will eventually appear in the ledger of party  $Q$  at a time  $t' > t$  (“stable” transaction). Concretely, for all honest parties  $P$  and  $Q$ , we have that  $\forall t \in \mathbb{N} : \forall t' \geq t + k : L^P[t] \preceq L^Q[t']$ , where  $L^P[t] \preceq L^Q[t']$  denotes that  $L^P$  at time  $t$  is a (not necessarily proper) prefix of  $L^Q[t']$  at time  $t'$ .*

As parties will eventually come to agreement about the blocks in their ledgers, we use the notation  $L[t]$  to refer to the ledger state at time  $t$  shared by all parties; similarly, we use the notation  $L[r]$  for the shared view of all parties at round  $r$ . This notation is valid when  $t$  is at least  $k$  time units in the past.

**Definition 2 (Liveness).** *Consider an honest party  $P$  of a ledger  $L$  and a liveness delay parameter  $u$ . If  $P$  attempts to write a transaction  $\text{TX}$  to its ledger at time  $t \in \mathbb{N}$ , then  $\text{TX}$  will appear in its ledger at time  $t'$ , i.e.,  $\exists t' \in \mathbb{N} : t' \geq t \wedge \text{TX} \in L^P[t']$ . The interval  $t' - t$  is upper bound by  $u$ .*

**Transaction Model.** A transaction  $\text{TX}$ , when included, alters the state of a ledger  $L$  by defining operations to be executed and agreed upon by consensus participants  $P_1, \dots, P_n$ . The expressiveness of operations is thereby left for the particular system to define, and can range from simple payments to execution of complex programs [181]. For generality, we do not differentiate between specifics transactions models (e.g. UTXO [152] or account-based models [181]).

### 2.3 Cross-Chain Communication System Model

Consider two independent distributed systems  $X$  and  $Y$  with underlying ledgers  $L_x$  and  $L_y$ , as defined in Section 2.2. We assume a *closed* system model as in [135] with a process  $P$  running on  $X$  and a process  $Q$  running on  $Y$ . A process can influence the state evolution of the underlying system by (i) writing a transaction  $\text{TX}$  to the underlying ledger  $L$  (commit), or (ii) by stopping to interact with the system (abort). We assume that  $P$  possesses transaction  $\text{TX}_P$ , which can be written to  $L_x$ , and  $Q$  possesses  $\text{TX}_Q$ , which can be written to  $L_y$ . A function

$desc$  maps a transaction to some “description” which can be compared to an expected description value, e.g., specifying the transaction value and recipient (the description differs from the transaction itself in that it may not, for example, contain any signature).  $P$  possesses a description  $d_Q$  which characterizes the transaction  $TX_Q$ , while  $Q$  possesses  $d_P$  which characterizes  $TX_P$ . Informally,  $P$  wants  $TX_Q$  to be written to  $L_y$  and  $Q$  wants  $TX_P$  to be written to  $L_x$ . Thereby,  $d_P = desc(TX_P)$  implies  $TX_P$  is valid in  $X$  (at time of CCC execution), as it cannot be written to  $L_x$  otherwise (analogous for  $d_Q$ ).

For the network, we assume no bounds on message delay or deviations between local clocks, unless the individual blockchain protocols require this. We treat failure to communicate as adversarial behavior. We note that, in the anonymous blockchain setting, more synchrony requirements are imposed than in the byzantine setting. Our construction does not impose any *additional* synchrony requirements than the individual ledger protocols. Hence, if  $P$  or  $Q$  become malicious, we indicate this using boolean “error variables” [95]  $m_P$  and  $m_Q$ . We assume  $P$  and  $Q$  know each other’s identity and no (trusted) third party is involved in the communication between the two processes.

## 2.4 Formalization of Correct Cross-Chain Communication

The goal of cross-chain communication can be described as the synchronization of processes  $P$  and  $Q$  such that  $Q$  writes  $TX_Q$  to  $L_y$  if and only if  $P$  has written  $TX_P$  to  $L_x$ . Thereby, it must hold that  $desc(TX_P) = d_Q \wedge desc(TX_Q) = d_P$ . The intuition is that  $TX_P$  and  $TX_Q$  are two transactions which must either both, or neither, be included in  $L_x$  and  $L_y$ , respectively. For example, they can constitute an exchange of assets which must be completed atomically.

To this end,  $P$  must convince  $Q$  that it created a transaction  $TX_P$  which was included in  $L_x$ . Specifically, process  $Q$  must verify that at given time  $t$  the ledger state  $L_x[t]$  contains  $TX_P$ . A cross-chain communication protocol which achieves this goal, i.e., is correct, must hence exhibit the following properties:

**Definition 3 (Effectiveness).** *If both  $P$  and  $Q$  behave correctly and  $TX_P$  and  $TX_Q$  match the expected descriptions (and are valid), then  $TX_P$  will be included in  $L_x$  and  $TX_Q$  will be included in  $L_y$ . If either of the transactions are not as expected, then both parties abort.*

$$(desc(TX_P) = d_Q \wedge desc(TX_Q) = d_P \wedge m_P = m_Q = \perp \implies TX_P \in L_x \wedge TX_Q \in L_y) \\ \wedge (desc(TX_P) \neq d_Q \vee desc(TX_Q) \neq d_P \implies TX_P \notin L_x \wedge TX_Q \notin L_y)$$

**Definition 4 (Atomicity).** *There are no outcomes in which  $P$  writes  $TX_P$  to  $L_x$  at time  $t$  but  $Q$  does not write  $TX_Q$  before  $t'$ , or  $Q$  writes  $TX_Q$  to  $L_y$  at  $t'$  but  $P$  did not write  $TX_P$  to  $L_x$  before  $t$ .*

$$\neg((TX_P \in L_x \wedge TX_Q \notin L_y) \vee (TX_P \notin L_x \wedge TX_Q \in L_y))$$

**Definition 5 (Timeliness).** *Eventually, a process  $P$  that behaves correctly will write a valid transaction  $TX_P$ , to its ledger  $L$ .*

From Persistence and Liveness of  $\mathbf{L}$ , it follows that eventually  $P$  writes  $\text{TX}_P$  to  $\mathbf{L}_x$  and  $Q$  becomes aware of and verifies  $\text{TX}_P$ .

**Definition 6 (Correct Cross-Chain Communication (CCC)).** *Consider two systems  $X$  and  $Y$  with ledgers  $\mathbf{L}_x$  and  $\mathbf{L}_y$ , each of which has Persistence and Liveness. Consider two processes,  $P$  on  $X$  and  $Q$  on  $Y$ , with to-be-synchronized transactions  $\text{TX}_P$  and  $\text{TX}_Q$ . Then a correct cross-chain communication protocol is a protocol which achieves  $\text{TX}_P \in \mathbf{L}_x \wedge \text{TX}_Q \in \mathbf{L}_y$  and has Effectiveness, Atomicity, and Timeliness.*

Summarizing, Effectiveness and Atomicity are safety properties. Effectiveness determines the outcome if transactions are not as expected or both transaction match descriptions and both processes are behaving correctly. Atomicity globally restricts the outcome to exclude behaviors which place a disadvantage on either process. Timeliness guarantees eventual termination of the protocol, i.e., is a liveness property.

## 2.5 The Generic CCC Protocol

We now describe the main phases of a Generic CCC Protocol, which can represent the transfer of good, assets or objects, between any two blockchain-based distributed systems  $X$  and  $Y$ . A visual representation is provided in Figure 1.

**1) Setup.** A CCC protocol is parameterized by the involved distributed systems  $X$  and  $Y$  and the corresponding ledgers  $\mathbf{L}_x$  and  $\mathbf{L}_y$ , the involved parties  $P$  and  $Q$ , the transactions  $\text{TX}_P$  and  $\text{TX}_Q$  as well as their descriptions  $d_P$  and  $d_Q$ . The latter ensure the validity of  $\text{TX}_P$  and  $\text{TX}_Q$  and determine the application-level specification of a CCC protocol. For example, in the case of an exchange of digital assets,  $d_P$  and  $d_Q$  define the asset types, transferred value, time constraints and any additional conditions agreed by parties  $P$  and  $Q$ . Typically, the setup occurs out-of-band between the involved parties and we hence omit this step hereby.

**2) (Pre-)Commit on  $X$ .** Upon successful setup, a publicly verifiable commitment to execute the CCC protocol is published on  $X$ :  $P$  writes<sup>2</sup> transaction  $\text{TX}_P$  to its local ledger  $\mathbf{L}_X^P$  at time  $t$  in round  $r$ . Due to Persistence and Liveness of  $\mathbf{L}_x$ , all honest parties of  $X$  will report  $\text{TX}_P$  as stable ( $\text{TX}_P \in \mathbf{L}_x$ ) in round  $r + u_x + k_x$ .

**3) Verify.** The correctness of the commitment on  $X$  by  $P$  is verified by  $Q$  checking (or receiving a proof from  $P$ ) that (i)  $d_P = \text{desc}(\text{TX}_P)$  and (ii)  $\text{TX}_P \in \mathbf{L}_x$  hold. From Persistence and Liveness of  $X$  we know the latter check will succeed at time  $t'$  which corresponds to round  $r + u_x + k_x$  on  $X$ , if  $P$  executed correctly.

**4a) Commit on  $Y$ .** Upon successful verification, a publicly verifiable commitment is published on  $Y$ :  $Q$  writes transaction  $\text{TX}_Q$  to its local ledger  $\mathbf{L}_Y^Q$  at time  $t'$  in round  $r'$  on  $Y$ . Due to Persistence and Liveness of  $\mathbf{L}_y$ , all honest parties of  $Y$  will report  $\text{TX}_Q$  as stable ( $\text{TX}_Q \in \mathbf{L}_y$ ) in round  $r' + u_y + k_y$ , where  $u_y$  is the liveness delay and  $k_y$  is the “depth” parameter of  $Y$ .

<sup>2</sup> In off-chain protocols [102], the commitment can be done by exchanging pre-signed transactions or channel states, which will be written to the ledger at a later point.

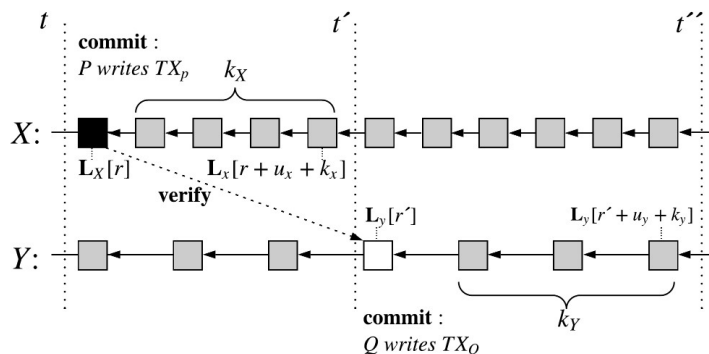


Fig. 1: CCC between  $X$  and  $Y$ . Process  $Q$  writes  $TX_Q$  only if  $P$  has written  $TX_P$ . We set exemplary persistence delays for  $X$  and  $Y$  as  $k_X = 4$  and  $k_Y = 3$ , and liveness delays as  $u_x = u_y = 0$ . We omit the optional the abort phase.

**4b) Abort.** If the verification fails and / or  $Q$  fails to execute the commitment on  $Y$ , a CCC protocol can exhibit an abort step on  $X$ , i.e., “reverting” the modifications  $TX_P$  made to the state of  $L_x$ . As blockchains are append-only data structures, reverting requires broadcasting an additional transaction  $TX_{P'}$ , which resets  $X$  to the state before the commitment of  $TX_P$ .

It is worth noting that some CCC protocols, specifically those facilitating *exchange* of assets, follow a two-phase commit design. In this case, steps 2 and 4a are executed in parallel, followed by the verification and (optional) abort steps on *both*  $X$  and  $Y$ . A further observation is that a CCC protocol necessarily requires a *conditional state transition* to occur on  $Y$ , given a state transition on  $X$ . As such, we do *not* consider (oracle) protocols which merely relay data across distributed ledgers [175, 62, 65, 5, 68], as CCC protocols by themselves.

### 3 Impossibility of CCC without a Trusted Third Party

In this section we show that, in the asynchronous setting, CCC is impossible without a trusted third party by reducing it to the Fair Exchange problem [38, 155].

**Fair Exchange.** On a high level, an exchange between two (or more) parties is considered fair if either both parties receive the item they expect, or neither do [40]. Fair exchange can be considered a sub-problem of fair secure computation [54], and is known to be impossible without a trusted third party [155, 182, 87, 86]. We recall the definition of Fair Exchange in Appendix A.

#### 3.1 What is a Trusted Third Party?

Numerous recent works use *a single* distributed ledger such as Bitcoin and Ethereum to construct (optimistic) fair exchange protocols [54, 37, 131, 124, 84, 128]. They leverage smart contracts (i.e., programs or scripts), the result of which

is agreed upon and enforced by consensus participants, to ensure the correctness of the exchange. These protocols thus use the consensus of the distributed ledgers as an abstraction for a trusted third party. If the majority of consensus participants are honest, correct behavior of processes/participants of the fair exchange is enforced – typically, the correct release of  $a_Q$  to  $P$  if  $Q$  received  $a_P$ .

A CCC protocol aims to achieve synchronization between *two* such distributed ledgers, both of which are inherently trusted to operate correctly. As we show below, a (possibly additional) TTP can be used to (i) confirm to the consensus participants of  $Y$  that  $\text{TX}_P$  was included in  $\mathbb{L}_x$ , who in turn enforce the inclusion of  $\text{TX}_Q$  in  $\mathbb{L}_y$ ; or (ii) directly enforce correct behavior of  $Q$ , such that  $\text{TX}_Q \in \mathbb{L}_y$ .

Similar to the abstraction of TTPs used in fair exchange protocols, in CCC it does not matter how exactly the TTP is implemented, as long as it enforces correct behavior of the participants. Strictly speaking, from the perspective of CCC there is little difference between a TTP consisting of a single individual and a committee where  $N$  out of  $M$  members must agree to take action (even though a committee is, without question, more resilient against failures) – contrary to the common assumptions made by the blockchain community.

### 3.2 Relating CCC to Fair Exchange.

We proceed to show that Correct Cross-Chain Communication is impossible in the asynchronous setting without a trusted third party (TTP), under the deterministic system model of distributed ledgers, by reducing CCC to Fair Exchange [40, 38, 155]. We recall, a fair exchange protocol must fulfill three properties: *Effectiveness*, *(Strong) Fairness* and *Timeliness* [155, 38] (cf. Appendix A).

**Lemma 1.** *Let  $M$  be a system model. Let  $C$  be a protocol which solves CCC in  $M$ . Then there exists a protocol  $S$  which solves Fair Exchange in  $M$ .*

*Proof (sketch).* Consider that the two processes  $P$  and  $Q$  are parties in a fair exchange. Specifically,  $P$  owns an item (or asset)  $a_P$  and wishes to exchange it against an item (or asset)  $a_Q$  owned by  $Q$ . Assume  $\text{TX}_P$  assigns ownership of  $a_P$  to  $Q$  and  $\text{TX}_Q$  transfers ownership of  $a_Q$  to  $P$  (specified in the “descriptions”  $d_P$  of  $\text{TX}_P$  and  $d_Q$  of  $\text{TX}_Q$ ). Then,  $\text{TX}_P$  must be included in  $\mathbb{L}_x$  and  $\text{TX}_Q$  must be included in  $\mathbb{L}_y$  to correctly execute the exchange. In other words, if  $\text{TX}_Q \in \mathbb{L}_y$  and  $\text{TX}_P \in \mathbb{L}_x$ , then  $P$  receives desired  $a_Q$  and  $Q$  receives desired  $a_P$ , i.e.,  $P$  and  $Q$  fairly exchange  $a_P$  and  $a_Q$ .

We observe the definition of Timeliness in CCC is equivalent to the definition of Timeliness in fair exchange protocols, as defined in [155]. Effectiveness in fair exchange states that if  $P$  and  $Q$  behave correctly and do not want to abandon the exchange (i.e.,  $m_P = m_Q = \perp$ ), and items  $a_P$  and  $a_Q$  are as expected by  $Q$  and  $P$ , then at the end of the protocol,  $P$  will own the desired  $a_Q$  and  $Q$  will own the desired  $a_P$  [155]. It is easy to see Effectiveness in CCC achieves exactly this property: if  $P$  and  $Q$  behave correctly and  $\text{desc}(\text{TX}_P) = d_P$  and  $\text{desc}(\text{TX}_Q) = d_Q$ , i.e.,  $\text{TX}_P$  transfers  $a_P$  to  $Q$  and  $\text{TX}_Q$  transfers  $a_Q$  to  $P$ , then  $P$  will write  $\text{TX}_P$  to  $\mathbb{L}_y$  at time  $t$  and  $Q$  will write  $\text{TX}_Q$  to  $\mathbb{L}_x$  before time  $t'$ . From Persistence and



Liveness of  $L_x$  and  $L_y$  we know both transactions will eventually be written to the local ledgers of  $P$  and  $Q$ , consequently all other honest participants of  $X$  will report  $TX_P \in L_X$  and honest participants of  $Y$  will report  $TX_Q \in L_Y$ . From our model we know that honest participants constitute majorities in both  $X$  and  $Y$ . Hence,  $P$  will receive  $a_Q$  and  $Q$  will receive  $a_P$ .

Strong Fairness in fair exchange states that there is no outcome of the protocol, where  $P$  receives  $a_Q$  but  $Q$  does not receive  $a_P$ , or, vice-versa,  $Q$  receives  $a_P$  but  $P$  does not receive  $a_Q$  [155]. In our setting, such an outcome is only possible if  $TX_P \in L_x \wedge TX_Q \notin L_y$  or  $TX_P \notin L_x \wedge TX_Q \in L_y$ , which contradicts the Atomicity property of CCC.

We construct a protocol for Fair Exchange using CCC in Appendix B. It is now left to show that CCC is defined under the same model as Fair Exchange. The distributed ledger model [93] used in CCC assumes the same asynchronous (explicitly) and deterministic (implicitly) system model (cf. Section 2.3) as [155, 90]. Since  $P$  and  $Q$  by definition can stop participating in the CCC protocol at any time, CCC exhibits the same crash failure model as Fair Exchange [39, 155] (and in turn Consensus [90]). Hence, we conclude:

**Theorem 1.** *There exists no asynchronous CCC protocol tolerant against misbehaving nodes without a trusted third party.*

*Proof.* Assume there exists an asynchronous protocol  $C$  which solves CCC. Then, due to Lemma 1 there exists a protocol which solves strong fair exchange. As this is a contradiction, there cannot exist such a protocol  $C$ .

Our result currently holds for the closed model, as in [155, 90]. In the open model,  $P$  and  $Q$  can be forced to make a decision by the system (or environment), i.e., transactions can be written on their behalf if they crash [129]. In the case of CCC, this means that distributed system  $Y$ , or more precisely, the consensus of  $Y$ , can write  $TX_Q$  to  $L_y$  on behalf of  $Q$  (if  $P$  wrote  $TX_P$  to  $L_x$ ). We observe that the consensus of  $Y$  becomes the TTP in this scenario: both  $P$  and  $Q$  must agree that the consensus of  $Y$  enforce correct execution of CCC. In practice, this can be achieved by leveraging smart contracts, similar to blockchain-based fair exchange protocols, e.g. [84]. As such, we can construct a smart contract, the execution of which is enforced by consensus of  $Y$ , that will write  $TX_Q$  to  $L_y$  if  $P$  includes  $TX_P$  in  $L_x$ , i.e.,  $Q$  is allowed to crash.

However, it remains the question how the consensus participants of  $Y$  become aware that  $TX_P \in L_x$ . In practice, a smart contract, can only perform actions based on some input. As such, before writing  $TX_Q$  the contract / consensus of  $Y$  must observe and verify that  $TX_P$  was included in  $L_x$ . A protocol achieving CCC must hence make one of the following assumptions. Either, there exists a TTP that will ensure correct execution of CCC; or the protocol assumes  $P$ , or  $Q$ , or some other honest, online party (this can again be consensus of  $Y$ ) will always deliver a proof for  $TX_P \in L_x$  to  $Y$  within a known, upper bounded delay, i.e., the protocol introduces some form of *synchrony* assumption. As argued in [155], we observe that *introducing a TTP and relying on a synchrony assumption are equivalent*:

*Remark 1.* When designing a CCC protocol, a choice must be made between introducing a trusted third party, or, equivalently, assuming some synchrony on the network.

The intuition behind this result is as follows. If we assume that process  $P$  does not crash and hence submits the necessary proof to the smart contract on  $Y$ , and that this message is delivered to the smart contract within a known upper bound, then we can be sure that CCC will occur correctly. Thereby,  $P$  intuitively represents its own trusted third party. However, if we cannot make assumptions on when the message will be delivered to the smart contract, as is the case in the asynchronous model, a trusted third party is necessary to determine the outcome of the CCC: the TTP observes  $\text{TX}_P \in \mathbb{L}_x$  and informs the smart contract or directly enforces the inclusion of  $\text{TX}_Q$  in  $\mathbb{L}_y$ . This illustrates how a TTP can be leveraged to enforce synchrony, i.e., timely delivery of messages, in CCC protocols. While the two models yield equivalent results, the choice between a TTP and network synchrony impacts the implementation details of a CCC protocol.

### 3.3 Incentives and Rational CCC

Several workarounds to the fair exchange problem, including gradual release mechanisms, optimistic models, and partially fair secure computation [40, 71, 132, 54], have been suggested in the literature. These workarounds suffer, among others, from a common drawback: they require some form of trusted party that does not collude with the adversary. Further, in case of an adversary-caused abort, honest parties must spend extra efforts to restore fairness, e.g., in the optimistic model the trusted server must be contacted each time fairness is breached.

First suggested in the context of rational exchange protocols [171], the economic dimension of blockchains enabled a shift in this paradigm: Rather than forcing an honest user to invest time and money to achieve fairness, the malicious user is economically punished when breaching fairness and the victim is reimbursed. This has paved the way to design *economically trustless* CCC protocols that follow a game theoretic model under the assumption that actors behave rationally [187]. We remark that malicious/altruistic actors can nevertheless breach CCC properties: even if there is no economic damage to parties  $P$  or  $Q$ , the correct execution of the communication protocol itself still fails.

## 4 The CCC Design Framework

With the impossibility result 3 and CCC model (Section 2.2) in mind, we now introduce a new framework for creating and evaluating CCC protocols. A generic CCC protocol consists of three main phases: commit (on  $X$ ), verify (and commit on  $Y$ ), and an optional abort. The main challenge of designing a CCC protocol is hence to determine the necessary trust model for each phase, from one of the following: (i) relying outright on a TTP, (ii) relying on an explicit synchrony

assumption, or (iii) a hybrid approach, where a TTP is only involved if synchrony is breached. The framework introduced below is structured as follows: for each CCC phase (subsection), we systematize the three possible trust models (TTP, synchrony, hybrid), outlining possible implementations and reasoning about practical consideration. This enables systematic evaluation of existing protocols and, at the same time, acts as a step-by-step guide for creating new CCC schemes.

#### 4.1 (Pre-)Commit Phase

The commit phase(s) of a CCC protocol typically involves the locking and unlocking of assets on chains  $X$  and  $Y$ , determined by the outcome of the protocol. **Model 1: Trusted Third Party (Coordinators)** A *coordinator* is a TTP that is tasked with ensuring correct execution of a CCC protocol. We classify coordinator implementations attending to two criteria: *custody of assets* and *involvement in blockchain consensus*. A coordinator (committee) can thereby be *static* (pre-defined) or *dynamic* (any user can join). And, finally, a CCC protocol can utilize *collateral* to incentivize correct behavior. We first introduce the classification criteria and then detail possible implementations of coordinators.

- *Custody of Assets*. Custody determines with whom the control over assets of (honest) participants resides. We differentiate between *custodians* and *escrows*. *Custodians* receive *unconditional* control over the participant’s funds and are thus *trusted* to release them as instructed by the protocol rules. *Escrows* receive control over the participant’s funds *conditional* to certain prearranged constraints being fulfilled. Contrary to custodians, escrows can fail to take action, e.g. freeze assets, but cannot commit theft.

- *Involvement in Consensus*. Coordinators can optionally also take part in the blockchain consensus protocol. *Consensus-level* coordinators refer to TTPs that are additionally consensus participants in the underlying chain. This is the case, for example, if the commit step is performed on chain  $X$  and enforced directly by the consensus participants of  $X$ , e.g. through a smart contract or directly a multi-/threshold signature. *External* coordinators, on the other hand, refer to TTPs which are not represented by the consensus participants of the underlying blockchain. This is the case if (i) the coordinators are external to the chain  $X$ , e.g. the consensus participants of chain  $Y$  or other parties, or (ii) less than the majority of consensus participants of chain  $X$  are involved.

- *Election*. An important distinction to make is between *static*, i.e., unchanged over time (usually permissioned), and *dynamic* coordinator sets. A dynamic coordinator can be chosen by CCC participants for each individual execution, or can be sampled by a pre-defined mechanism, as e.g. studied in [78, 126, 127, 156] for Proof-of-Work and in [148, 123, 77, 55] for Proof-of-Stake blockchains. We consider CCC protocols where any user can become a coordinator as *unrestricted*, while protocols that require coordinators to be approved by some third party are referred to as *restricted*.

- *Incentives and Collateralization*. Instead of following a *prohibitive* approach, i.e., technically preventing or limiting coordinators from deviating from

protocol rules, a CCC protocol can follow a *punishable* approach. That is, ensure misbehavior can be proven and penalized retrospectively. In the latter case, a coordinator will typically be required to lock collateral that can be *slashed* and allocated to (financially) damaged CCC participants.

*Coordinator Implementations.* We now detail the different coordinator types according to the aforementioned criteria and how they are implemented in practice.

- *External Custodians (Committees).* Instead of relying on the availability and honest behavior of a single external coordinator, trust assumptions can be distributed among a set of  $N$  committee members. Decisions require the acknowledgment (e.g. digital signature) of at least  $M \leq N$  members, whereby consensus can be achieved via Byzantine Fault Tolerant (BFT) agreement protocols such as PBFT [72, 126]. External custodians can be both static or dynamic, and collateralization can be added on involved blockchains to incentivize honest behavior.

- *Consensus-level Custodians (Consensus Committee)* are identical to external custodians, except that they are also responsible for agreeing on the state of the underlying ledger. This model is typically used in blockchain sharding [129, 34], where the blockchain  $X$  on which the commit step is executed runs a BFT consensus protocol, i.e., there already exists a *static* committee of consensus participants that much be trusted for correctness of CCC (Persistence and Liveness of  $X$ ). Collateralization of Consensus Custodians is best handled on another blockchain, i.e., where the coordinators have no influence on consensus.

- *External Escrows (Multisignature Contracts).* External Escrows are a special case of External Custodians, where the coordinator is transformed from Custodian to Escrow by means of a *multisignature contract*. Multisignature contracts require signatures of a subset (or majority) of committee members *and* the participant  $P$  (e.g., the asset owner), i.e.,  $P + M, M \leq N$ . The committee can thus only execute actions pre-authorized by the participant: it can at most freeze assets, but not commit theft.

**Model 2: Synchrony Assumptions (Lock Contracts)** An alternative to coordinators consists in relying on synchronous communication between participants and leveraging locking mechanisms which harvest security from cryptographic hardness assumptions. Such protocols are often referred to as *non-custodial*, as they avoid transferring custody over assets to a TTP – failures, in the worst case, result in a permanent lockup of funds without explicit (financial) benefits to a third party. We differentiate between *symmetric* contracts, where identical locks are created on both chains and released atomically, and *asymmetric* contracts where the main protocol logic is hosted on a single chain.

- *Hash Locks (symmetric).* A protocol based on hash locks relies on the *preimage resistance* property of hash functions: participants  $P$  and  $Q$  transfer assets to each other by means of transactions that must be complemented with the preimage of a hash  $h := H(r)$  for a value  $r$  chosen by  $P$  – the initiator of the protocol – typically uniformly at random [20, 19, 107, 140].

- *Signature-based Locks (symmetric).*  $P$  and  $Q$  can transfer assets to each other by means of transactions that require to solve the discrete logarithm problem of a value  $Y := g^y$  for a value  $y$ , chosen uniformly at random by  $P$  (i.e., the

initiator of the protocol). The functionality of embedding the discrete logarithm problem in the creation of a digital signature was put forward by the community under the term *adaptor signatures* [157] and formally defined in [41]. In practice, it has been shown that it is possible to implement adaptor signatures leveraging virtually any digital signature scheme [177], including ECDSA and Schnorr which are used for authorization in most blockchains today [60, 58, 141, 172, 85, 157, 151].

- *Timelock Puzzles and Verifiable Delay Functions (symmetric)*. An alternative approach is to construct (cryptographic) challenges, the solution of which will be made public at a predictable time in the future. Thus,  $P$  and  $Q$  can commit to the cross-chain transfer conditioned on solving one of the aforementioned challenges. Concrete constructions include timelock puzzles and verifiable delay functions. Timelock puzzles [160] build upon inherently sequential functions where the result is only revealed after a predefined number of operations are performed. Verifiable delay functions [58] improve upon timelock puzzles on that the correctness of the result for the challenge is publicly verifiable. This functionality can also be simulated by releasing parts of the preimage of a hash lock interactively bit by bit, until it can be brute forced [53].

- *Smart Contracts (asymmetric)* are programs stored in a ledger which are executed and their result agreed upon by consensus participants [67, 70]. As such, trusting in the correct behavior of a smart contract is essentially trusting in the secure operation of the underlying chain, making this a useful construction for (Consensus-level) Escrows. Contrary to Consensus-level Custodians, who must actively follow the CCC protocol and potentially run additional software, with smart contracts, consensus participants are not directly involved in the CCC protocol: an interaction with the CCC smart contract is, by default, treated like any other state transition and no additional software/action is required.

**Model 3: Hybrid (Watchtowers)** Instead of fully relying on coordinators being available or synchrony assumptions among participants holding, it is possible to employ so called *watchtowers*, i.e., service providers which act as a fallback if CCC participants experience crash failures. We observe strong similarities with optimistic fair exchange protocols [40, 39, 71]. Specifically, watchtowers take action to enforce the commitment, if one of the parties crashes or synchrony assumptions do not hold, i.e., after a pre-defined timeout [119, 45, 142, 43]. This construction was first introduced and applied to off-chain payment channels [102].

## 4.2 Verification Phase

The verification phase, during which the commitment on  $X$  is verified on  $Y$  (or vice-versa), can similarly be executed under different trust models, as detailed in the following. An important distinction concerns the type of verification performed: while most CCC protocols verify the inclusion of a transaction executing the commitment on  $X$ , full validation of correctness under  $X$ 's protocol rules is typically avoided due to the incurred computational overhead. A detailed analy-

sis and taxonomy of different verification techniques is provided in Appendix C

**Model 1: Trusted Third Party (Coordinators).** The simplest approach to cross-chain verification is to rely on a trusted third party (also referred to as *validators* [180]) to handle the verification of the state changes on interlinked chains during CCC execution.

- *External Validators.* A simple approach is to outsource the verification step to a (trusted) third party, external to the verifying ledger (in our case  $Y$ ), as in [176, 18]. The TTP can then be the same as in the commit/abort steps.

- *Consensus Committee / Smart Contracts.* Alternatively, the verification can be handled by the consensus participants of the verifying chain [129, 81, 137], leveraging the assumption that misbehavior of consensus participants indicates a failure of the chain itself.

**Model 2: Synchrony Assumption.** Instead of explicitly relying on a TTP, the verification phase can be implemented using:

- *Direct Observation.* Similar to the commit phase of CCC, one can require all participants of a CCC protocol to execute the verification phase individually: i.e., to run (fully validating) nodes in all involved chains. This is often the case in exchange protocols, such as atomic swaps using symmetric locks such as HTLCs [20, 107], but also in parent-child settings where one chain by design verifies or validates the other [47, 97, 136]. This relies on a synchrony assumption, i.e., requires CCC participants to observe commitments and act within a certain time, in order to complete the CCC.

- *Chain Relay Smart Contracts.* The verification process can be encoded in so called *chain relays* [5, 68, 187] – smart contracts deployed on  $Y$  capable of verifying the of state and hence the commitments executed on  $X$ . Chain relays resemble cryptocurrency light (or SPV) clients, i.e., store only the bare minimum data to verify the inclusion of transactions in the respective blockchain [152, 121, 139]. Accordingly, chain relays can only *verify* that a commitment was executed on  $X$  – yet not if it was valid under  $X$ ’s consensus rules. Instead, the “SPV assumption” is applied: if  $X$  has Persistence and Liveness, then a commitment (transaction) written to  $X$  must be valid [152, 139]. To fully *validate* the correctness of a commitment, one must either (i) download the entire state of chain  $X$  (infeasible for CCC), or (ii) encode the state of  $X$  in succinct proofs of knowledge [57, 49, 64] (c.f. Appendix C ).

**Model 3: Hybrid.** Verificaiton via TTPs and synchrony can be combined:

- *Watchtowers.* Just like in the commit phase, synchrony and TTP assumptions can be combined in the verification phase, such that a CCC protocol initially relies on a synchrony assumption, but can fall back to a TTP (*watchtowers*, c.f. Section 4.1) to ensure correct termination if messages are not delivered within a per-defined period.

- *Verification Games.* Inversely, *verification games* by default rely on TTPs for verification (mostly for performance improvements) and implement dispute resolution mechanisms as fall-back: users can provide (reactive) fraud proofs [35]

or accuse coordinators of misbehavior requiring them to prove correct operation [173, 104, 117].

### 4.3 Abort Phase

The abort of a CCC protocol is optional and is encountered typically in exchange protocols. Most other CCC protocols assume that once a commit is executed on  $X$ , no abort will be necessary.

**Model 1: Trusted Third Party (Coordinators)** Similarly to the commit phase, an abort can be handled by a trusted third party and the possible implementations are the same as in Section 4.1. If a TTP was introduced in the commit phase, the abort phase will be typically handled by the exact same TTP.

**Model 2: Synchrony Assumptions (Timelocks)** Alternatively, it is possible to enforce synchrony by introducing timelocks, after the expiry of which the protocol is aborted. Specifically, to ensure that assets are *not locked up indefinitely* in case of a crash failure of a participant or misbehavior of a TTP entrusted with the commit step, all commit techniques can be complimented with *timelocks*: after expiry of the timelock, assets are returned to their original owner. We differentiate between two types of timelocks. *Absolute timelocks* yield a transaction valid only after a certain point in time, defined in by a timestamp or a block (ledger at index  $i$ ,  $L[i]$ ) located in the future. *Relative timelocks*, on the other hand, condition  $TX_2$  on the existence of another transaction  $TX_1$ :  $TX_2$  only becomes valid and can be written to the underlying ledger if  $TX_1$  has already been included and a certain number of blocks (*confirmations* [7]) have passed.

**Model 3: Hybrid (Watchtowers)** As an additional measure of security, TTPs can be introduced as a fallback to timelocks in case CCC participants experience crash failures, e.g. in form of a watchtower [119, 45, 142, 43] that recovers otherwise potentially lost assets. This is specifically useful in the case of atomic swaps using Hased Timelock Contracts (HTLCs) [20, 107, 3, 158], when either party crashes after the hashlock’s secret has been revealed.

## 5 Classification of Existing CCC Protocols

We now apply the CCC Design Framework introduced in Section 4 to classify existing CCC protocols. All CCC protocols observed in practice follow the Generic CCC Protocol model (cf. Section 2.5). For each protocol, we hence study and reason about the trust model (TTP, synchrony, hybrid) selected for each phase of the CCC process, and summarize our classification in Table 1. Our analysis thereby focuses on well-documented or deployed protocols - which in turn have seen numerous implementations that are not referenced in this paper.

In addition to applying the CCC Design Framework, we split existing proposals into two protocol families, based on their design rationale and use case, which have direct implications on the design choices: (i) *exchange* protocols, which synchronize the exchange of assets across two ledgers (Section 5.1), and (ii) *migration* protocols, which allow to move an asset or object to a different ledger (Section 5.2).

Table 1: Classification of existing of Cross-Chain Communication protocols, in consideration of the selected TTP model (cf. Section 4) at each protocol step (commit, verify, abort). Notation for non-binary TTP values: ● uses a TTP, ○ fully relies on synchrony and availability of participants, ◐ hybrid. We also highlight if the TTP (committee) is static or changes dynamically, and whether collateral is utilized to incentivize correct behavior of TTPs. We use the following abbreviations: **EC** for External Custodian, **CC** for Consensus Custodian, **EE** for External Escrow, **SC** for Smart Contract, **EV** for External Validator, **CM** for Consensus Committee, and **DO** for Direct Observation.

| Protocol            | Trust Model at each CCC Protocol Phase          |             |      |                                |      |                  |    |                         |
|---------------------|---|-------------|------|--------------------------------|------|------------------|----|-------------------------|
|                     | Commit on chain X                               |             |      | Verify & Commit on chain Y     |      |                  |    |                         |
|                     | TTP Dynamic?                                    | Collateral? | Type | TTP                            | Type | TTP              |    |                         |
| Exchange Protocols  | Traditional Custodial Exchanges (e.g., [1, 53]) | ●           | ×    | EC (single, restricted)        | ●    | EV               | ●  | EC (single, restricted) |
|                     | A2L [172]                                       | ◐           | ✓    | EE (multisig + signature Lock) | ○    | DO               | ◐  | EE + Timelock           |
|                     | Arwen [106]                                     | ◐           | ×    | EE (multisig + Hash Lock)      | ○    | DO               | ◐  | EE + Timelock           |
|                     | Notarized HTLC Atomic Swaps [176]               | ○           | -    | Hash Lock                      | ●    | EV               | ◐  | EE + Timelock           |
|                     | HTLC Atomic Swaps [20, 107, 19, 176]            | ○           | -    | Hash Lock                      | ○    | DO               | ○  | Timelock                |
|                     | ECDSA/DLSAG Atomic Swaps [141, 151]             | ○           | -    | Signature Lock                 | ○    | DO               | ○  | Timelock                |
|                     | SPV Atomic Swaps [10, 125, 187, 108]            | ○           | -    | Standard payment               | ○    | SC(chain relay)  | ○  | Timelock                |
|                     | (Bidirectional) Chain Relays [125, 97]          | ○           | -    | SC                             | ○    | SC (chain relay) | -  | -                       |
|                     | XCLAIM [187], Dogethereum [174]                 | ●           | ✓    | EC (single, unrestricted)      | ○    | SC (chain relay) | -* | -                       |
|                     | tBTC [17]                                       | ●           | ✓    | EC (committee, restricted)     | ○    | SC (chain relay) | -* | -                       |
| Migration Protocols | Custodial Wrapped Assets (e.g., [18, 30, 29])   | ●           | ×    | EC (single, restricted)        | ●    | EV               | ●  | EC (single, restricted) |
|                     | Federated Sidechains/Pegs [47, 81, 97]          | ●           | ×    | EC (consensus of Y)            | ●    | CM               | -* | -                       |
|                     | BSK [137, 136]                                  | ●           | ×    | EC (consensus of Y)            | ●    | CM               | -* | -                       |
|                     | ATOMIX[129], SBAC[34], Fabric Channels[36]      | ●           | ×    | CC (shard X)                   | ●    | CM               | ●  | CC (shard X)            |
|                     | Rapidchain [185]                                | ●           | ×    | CC (shard X)                   | ●    | CM               | -* | -                       |
|                     | XCMP [66]                                       | ●           | ×    | EC (parent consensus)          | ●    | CM               | -* | -                       |
|                     | Proof-of-Burn (Federated) [169, 118]            | ○           | -    | SC / Burn address              | ●    | CM               | -  | -                       |
|                     | Proof-of-Burn (SPV) [118]                       | ○           | -    | SC / Burn address              | ○    | SC (chain relay) | -  | -                       |
|                     | Merged Mining/Staking [116, 97]                 | ●           | ×    | CC (consensus of X)            | ●    | CM               | -  | -                       |

\* While not explicitly considered by the protocol, the TTP used for the commitment on X can, at its discretion, abort the CCC protocol manually/out-of-band in case of failure on Y.



## 5.1 Exchange Protocols

Exchange protocols synchronize an atomic exchange of digital goods:  $x$  on chain  $X$  against  $y$  on  $Y$  (c.f. Fair Exchange in Appendix A). In practice, such protocols implement a *two-phase commit* mechanism, where parties first pre-commit to the exchange and can *explicitly abort* the protocol in case of disagreement or failure during the commit step.

**(Pre-)Commit.** Trivially, the commit phase can be handled by External Custodians: traditional, centralized exchanges require to deposit (commit) assets with a TTP before trading.

The longest-standing alternative to centralized solutions are atomic swaps via *symmetric locks* which rely on synchrony and cryptographic hardness assumptions. Counterparties  $P$  and  $Q$  lock (pre-commit) assets in on-chain contracts with identical release conditions on  $X$  and  $Y$ : spending from one lock releases the other, ensuring Atomicity of CCC. The first and most adopted implementation of symmetric locks are *hashed timelock contracts* (HTLCs) [19, 20, 107, 176], where the same secret (selected by  $P$ ) is used as pre-image to identical Hash Locks on  $X$  and  $Y$ . To improve cross-platform compatibility, Hash Locks, which require both chains to support (the same) hash functions, can be replaced with Signature Locks e.g., using ECDSA [141] or group/ring signature schemes [151].

On blockchains which support (near) Turing complete programming languages (e.g., Ethereum [67]) the commitment on  $X$  can exhibit more complex locking conditions via smart contracts. In SPV atomic swaps [10, 125, 187, 108], assets of a party  $P$  are locked in a smart contract on  $X$  which is capable of verifying the state of chain  $Y$  (chain relay, cf. Section 4.2) - and unlocked only if counterparty  $Q$  submits a correct proof for the expected payment (commitment) on  $Y$ . The smart contract can be further extended to support collateralization and penalties for misbehaving counterparties (e.g., to mitigate optionally and improve fairness [103, 187]).

Both symmetric and SPV atomic swaps suffer from usability challenges impeding adoption: they require users to be online and execute commitments in a timely manner to avoid financial damage (built-in abort mechanisms discussed later). Hybrid protocols seek to combine symmetric locks with TTP models to mitigate usability issues while avoiding full trust in a central provider. In Arwen [106], parties  $P$  and  $Q$  commit to-be-exchanged assets into on-chain multisignature contracts on  $X$  and  $Y$ , establishing shared custody with an External Escrow (EE). Trades are executed similar to HTLC swaps, yet utilize the escrow to ensure correct and timely execution. A2L [172] follows a similar multisignature setup but utilizes adaptor signatures [41] to ensure Atomicity of trades: the escrow only forwards  $P$ 's assets to  $Q$  if  $Q$  solves a cryptographic challenge, for which  $Q$  needs the help of  $P$ . Both Arwen and A2L require a complex on-chain setup process (similar to payment channels [158]) and rely on pre-paid fees (Arwen) or collateral (A2L) to protect the escrow from griefing attacks [75] - yielding them inefficient for one-time exchanges.

**Verify.** Contrary to traditional exchanges, where the custodial (operator) is also responsible for the verification phase, symmetric atomic swap protocols

(including Arwen and A2L) require users to *directly observe* all chains involved in the CCC to verify correct execution of the (pre-)commit phase. Notarized atomic swaps (e.g., as in InterLedger [176]) remove the online requirement for users by entrusting an External Validator (EV) e.g., a set of notaries, with the verification of (and timely reaction to) the commitment on  $X$ - at the risk of the EV colluding with the a counterparty to commit theft. A more robust approach, implemented in SPV atomic swaps, is the use of *chain relays*: the verification of the commit on  $X$  and the correct finalization of the CCC protocol (commit on  $Y$ ) is executed by a smart contract on  $Y$ , enforced by the consensus of  $Y$ .

**Abort.** Exchange CCC protocols typically add timelocks to the release conditions of the commitments of  $X$  and  $Y$  to ensure an automatic abort of the CCC protocol after a pre-defined delay. This is to prevent indefinite lock-up of assets, should a party crash or misbehave. However, CCC protocols implementing timelocks impose strict online requirements on participants and expose them to race conditions. The initiator  $P$  of e.g., a HTLC swap can defraud counterparty  $Q$  by recovering assets on  $X$  if they remain unclaimed upon expiry of the time-lock (e.g., if  $Q$  crashed). Some protocols, including A2L and Arwen, partially outsource this responsibility to TTPs [106, 172].

## 5.2 Migration Protocols

Migration protocols temporarily or permanently move digital goods from one blockchain to another. Typically, this is achieved by obtaining a “write lock” on an asset  $x$  on chain  $X$ , preventing any further updates to  $x$  on chain  $X$ , and consequently creating a *representation*  $y(x)$  on  $Y$ . The state of  $x$  can only be updated by modifying its “wrapped” version  $y(x)$  on  $Y$ - comparable to the concept of *mutual exclusion* in concurrency control [80]. The state changes of  $y(x)$  will typically be reflected back to chain  $X$  by locking or destroying (“burning”)  $y(x)$  and applying the updates to  $x$  when it is unlocked.

Migration protocols only require to execute CCC synchronization across  $X$  and  $Y$  twice: creating and destroying  $y(x)$ . The “wrapped” representation  $y(x)$  typically exhibits the same properties as “native” assets  $y$ , allowing seamless integration with applications on  $Y$ . For comparison, Exchange protocols require to setup and execute CCC for *each trade*. The main drawback of Migration protocols is the requirement of giving up custody over  $x$ , in the majority of cases to a TTP (cf. Table 1).

In practice, we identify four main use cases for Migration protocols: (i) *cryptocurrency-backed assets* used for transfers across heterogeneous blockchains (e.g., “wrapped” Bitcoin on Ethereum), (ii) communication across homogeneous chains (shards) in *sharded* blockchains, (iii) *sidechains* where a child chain is “pegged” to a parent for feature extensions, and (iv) *bootstrapping* of new blockchains using existing systems.

**(Pre-)Commit.** The simplest implementation of a Migration protocol (e.g., for cryptocurrency-backed assets) relies on a single, static TTP which receives unrestricted custody over the to-be-migrated assets during the commit phase

(External Custodian) – for example, implemented by wBTC [18], a custodial platform for migrating Bitcoin to Ethereum.

Instead of relying on a single TTP, most CCC rely on a TTP committee to improve robustness against failures. Protocols connecting heterogeneous blockchains via cryptocurrency-backed assets, notably tBTC [17], utilize a set of External Custodians (EC). In the tBTC protocol, currently deployed between Bitcoin and Ethereum, ECs construct a jointly controlled deposit public key on  $X$  via (ECDSA) threshold signatures [98], to which users send (commit) to-be-migrated assets. The ECs must thereby lock up collateral on  $Y$  which is used to reimburse users in case the EC committee commits theft or crashes. At the time of writing, the implemented threshold signature scheme does not support fault attribution, i.e., it is impossible to distinguish between honest and malicious committee members when slashing collateral, requiring the EC set to be static and restricted. RenVM [30] aims to replace threshold signatures with distributed key generation via secure multi-party computation [100] but implements a centralized approach at the time of writing.

*Sidechains* [47, 81, 97] establish a parent-child relationship between  $X$  and  $Y$ : the consensus committee of  $X$  (Consensus Custodian, CC) or  $Y$  (External Custodian, EC) is responsible for handling the correct deposit (commit) of  $x$  on  $X$ . In practice, implementations follow a similar approach to the heterogeneous setting: users deposit assets  $x$  to a public key with shared control among committee members, implemented e.g., via threshold / multisignature [111] schemes. Liquid [47, 81], which coined the “sidechain” terminology, maintains an 11-of-15 multisignature, controlled by its consensus participants, to migrate (lock/unlock) Bitcoin to and from the Liquid blockchain. RSK [137, 136], a merge-mined [116] Bitcoin sidechain, currently follows the same approach as Liquid but envisions a Bitcoin protocol upgrade enabling miners to vote on migrating assets to RSK.

Similarly, *sharded blockchains*, which consist of a set of homogeneous shard-chains with a homogeneous, shared security model, utilize the consensus committee(s) available within the system for securing cross-shard migrations. While often considered as a separate topic in research, sharded blockchains exhibit built-in CCC protocols [44]: Migrated assets  $x$  are locked with the consensus of  $X$  (Consensus Custodian, CC) during the commit phase. A novelty compared to heterogeneous systems is the explicit consideration of  $n$ -to- $m$  CCC protocols, such as ATOMIX [129], SBAC [34], and Fabric Channels [36], which require an explicit abort step as part of the two-phase commit design.

Recently, a new family of protocols following a permissionless design, was introduced. XCLAIM [187] and Dogethereum [174] allow anyone to become a TTP and accept deposits (commits) of  $x$  on  $X$ , establishing a dynamic and unrestricted set of coordinators (External Custodians, ECs). The only requirement for registering as an EC is to lock collateral  $y$  on  $Y$  – the amount of  $y$  locked thereby determines the amount of  $x$  deposits (and hence minted  $y(x)$ ) an EC can accept. While Dogethereum assumes a constant exchange rate between migrated  $x$  (equiv.  $y(x)$ ) and collateral asset  $y$ , XCLAIM utilizes a multi-stage over-collateralization scheme to re-balance the economic value of committed  $x$

and locked collateral  $y$ . To enable ECs to join and leave the system at any point in time, XCLAIM implements a replacement/auction mechanism via cross-chain SPV atomic swaps, where collateral  $y$  can be exchanged for committed  $x$  held in custody.

In cases where  $X$  and  $Y$  support smart contracts, specifically chain relays, *bidirectional* chain relays [125, 97] can be utilized, enabling non-custodial commitments on  $X$  and  $Y$ : locking of  $x$  and unlocking/minting of  $y(x)$  is handled exclusively by smart contracts under the assumption of synchrony.

Proof-of-Burn [169, 118] resembles follows a similar design, yet implements a unidirectional protocol: instead of being locked,  $x$  is provably destroyed (“burned”), and newly minted as  $y(x)$  on  $Y$ . As such, Proof-of-Burn is mostly used for bootstrapping of new blockchains. Merged mining [116] was the first CCC protocol deployed in practice (2011 in Namecoin) and is used explicitly for bootstrapping purposes. Miners (stakers) of  $X$  can reuse PoW solutions (stake) to progress consensus on  $Y$  by including a commitment to  $Y$ ’s state in the ledger of  $X$ .

**Verify.** Migration protocols - with the exception of centralized, custodial services - rely the on consensus of chain  $Y$  to correctly verify the commitment on  $X$ . We observe two main implementation techniques: (i) under synchrony assumptions by using *chain relay* smart contracts, which cryptographically verify the correctness of the commitment on  $X$ , or (ii) by requesting the consensus committee of  $Y$  to explicitly sign off on the CCC execution. XCMP [66], a cross-shard protocol, adds an additional verification step: cross-shard transfers are verified by and included in a hierarchically “superior” parent chain – which in turn is verified by the target shard  $Y$  before commitment.

**Abort.** We observe that Migration protocols generally do not implement an explicit abort phase. Instead, they assume that if the commitment on  $X$  is executed correctly it will eventually be verified by chain  $Y$ , which in turn will result in a correct commitment on  $Y$ . An exception hereof are  $n$ -to- $m$  transfers in sharded blockchains (e.g., ATOMIX [129] and SBAC [34]) which require an explicit abort phase. Such transfers follow a two-phase-commit protocol: assets on all source shards  $X_1, \dots, X_n$  are pre-committed and verified on all target shards  $Y_1, \dots, Y_n$ , which in turn execute a pre-commitment. If a single target shard fails to reply with a pre-commitment (within some period), the CCC protocol is aborted on all other source and target shards.

### 5.3 Insights and General Observations

An interesting, yet expected insight is that performance and usability outweigh security considerations from a user’s perspective. Decentralized and non-custodial CCC solution have been proposed as early as 2013 (symmetric swaps [19]) and 2015 (SPV swaps [10]), yet centralized providers remain the dominant cross-chain asset exchange facilitator. The recent rise of decentralized exchanges, which mostly operate within a single chain [24], has boosted the adoption of cryptocurrency-backed assets, although predominantly via custodial approaches: at the time of writing, 99% of “wrapped” Bitcoin on Ethereum has been issued through trusted, custodial services [26].

Decentralized CCC protocols still suffer from practical drawbacks hindering adoption. Symmetric atomic swaps impose strict online requirements on users. SPV atomic swaps, and similarly migration protocols such as XCLAIM and tBTC, make use of chain relays which are only feasible if  $Y$  supports smart contracts and the cryptographic primitives used in  $X$ . Orthogonal, collateralization, which allows to protect users from financial damage (cf. Section 3), incurs high capital requirements and opportunity cost – leading most users to resort to trusted, centralized solutions.

An interesting observation hereby is that sharded systems and sidechains do not necessarily benefit from decentralized CCC protocols. In fact, due to the homogeneous nature of the security models of  $X$  and  $Y$  in this setting, the use of the consensus committee(s) of  $X$  or  $Y$  as TTP for CCC does not introduce any additional (external) trust assumptions to the underlying systems.

## 6 CCC Challenges and Outlook

In this section, we provide an outlook on the (open) problems faced by CCC protocols and possible avenues for future work.

### 6.1 Heterogeneous Models and Parameters Across Chains

**Problems.** Different blockchains leverage different system models and parameterizations, which, if not handled correctly by CCC protocols, can lead to protocol failures. For instance, the absence of a global clock across chains requires CCC participants to either agree on a trusted third party as means of synchronization, or to rely on a chain-dependent time definition (e.g., block generation rates [93]) which are often non-deterministic and hence unsafe for strictly time-bound protocols [93, 187]. A practical example hereof are race-condition attacks on symmetric exchange protocols such as HTLC atomic swaps, discussed in Section 5.

Another consideration are the security models of interconnected chains: while  $X$  and  $Y$  may exhibit well defined security models, these are typically independent and not easily comparable (with the exception of sharding) – especially when combined within a CCC protocol. For instance,  $X$  may rely on PoW and thus assume that adversarial hash rate is bound by  $\alpha \leq 33\%$  [88, 99, 163]. On the other hand,  $Y$  may utilize PoS for consensus and similarly assume that the adversary’s stake in the system is bound by  $\beta \leq 33\%$ . While similar at first glance, the cost of accumulating stake [96, 89] may be lower than that of accumulating computational power, or vice-versa [61]. Since permissionless ledgers are not Sybil resistant [82], i.e., provide weak identities at best, quantifying adversary strength is challenging even within a single ledger [42]. This task becomes nearly impossible in the cross-chain setting: not only can consensus participants (i) “hop” between different chains [147, 134], destabilizing involved systems, but also (ii) be susceptible to bribing attacks executed cross-chain, against which there currently exist no countermeasures [144, 115].

Following from different security models, the lack of homogeneous finality guarantees [168] across blockchains poses another challenge for CCC. Consider the following:  $X$  accepts a transaction as valid when confirmed by  $k$  subsequent blocks e.g., as in PoW blockchains [93]; instead,  $Y$  deems transactions valid as soon as they are written to the ledger ( $k = 1$ , e.g. [31]). A CCC protocol triggers a state transition on  $Y$  conditioned on a transaction included in  $X$ , however, later an (accidental) fork occurs on  $X$ . While the state of  $X$  is reverted, this may not be possible on  $Y$  according to consensus rules – likely resulting in an inconsistent state on  $Y$  and financial damage to users.

**Outlook.** Considering the plethora of blockchain designs in practice, it is safe to assume a heterogeneous ecosystem for at least the near future. Protocol designers must hence carefully evaluate and consider the specifics of each interlinked chain when implementing CCC schemes: introduction of conservative lower bounds on transaction (commit) finality (hours / days rather than minutes), analysis of computation and communication capabilities of consensus participants, and accounting for peer-to-peer network delays when utilizing a trusted third party as global clock.

## 6.2 Heterogeneous Cryptographic Primitives Across Chains

**Problems.** Interconnected chains  $X$  and  $Y$  may rely on different cryptographic schemes, or different instances of the same scheme. CCC protocols, however, often require compatible cryptographic primitives: a CCC protocol between a system  $X$  using ECDSA [112] as its digital signature scheme and a system  $Y$  using Schnorr [164] is only seamlessly possible if both schemes are instantiated over the same elliptic curve [141]. This is, for example, the reason Ethereum uses the same `secp256k1` curve as Bitcoin [25].

Similarly, CCC protocols using Hash Locks, e.g. HTLC swaps, require that the domain of the hash function has the same size in both  $X$  and  $Y$  – otherwise the protocol is prone to *oversize preimage attacks* [114], i.e., an attack where a transaction cannot be accepted by a chain because the representation of the preimage requires more bits than those previously allocated to store it.

**Outlook.** A design challenge in CCC protocols is thus the interoperability of chains in terms of (cryptographic) primitives as required in CCC protocols. In cases where interlinked chains implement different elliptic curves, zero-knowledge proofs may provide a workaround, yet at the cost of increased protocol complexity, as well as computation and communication costs [154]. Our observations suggest that this is one of the main reasons for lack of interoperability across current blockchain networks.

## 6.3 Collateralization and Exchange Rates

**Problems.** In recent works [187, 174, 125, 17], we observe a trend towards collateralizing coordinators to prevent financial damage to users and incentivize correct behavior of TTPs. Thereby, it is crucial to ensure that the provided collateral has sufficient value to outweigh potential gains from misbehavior. However, in

the cross-chain setting, where insured asset and collateral are typically different, collateralized CCC protocols are forced to (i) implement measures against exchange rate fluctuations such as over-collateralization incurring capital inefficiencies for participants, and (ii) rely on (typically centralized) price oracles.

**Outlook.** Current CCC protocols, if at all, only provide minimal protection against exchange rate fluctuations, such as over-collateralization. An interesting avenue for future research is hence the design of dynamic collateralization e.g., based on the volatility of the locked/collateral assets. Decentralized price oracles already are an active field of research [159, 32, 175, 6, 188], yet as of this writing oracles remain single points of failure in collateralized CCC protocols. Cryptocurrency-backed assets traded on decentralized exchanges, where trading data is available on-chain, may thereby provide a valuable source of information for cross-verification with centralized providers [187].

#### 6.4 Lack of Formal Security Analysis

**Problems.** While numerous CCC protocols have been deployed and used in practice, handling value transfers worth millions, most lack formal and rigorous security analysis. This lack of formal security guarantees opens the door to possible security threats. For instance, *replay attacks* on state verification, i.e., where proofs are re-submitted multiple times or on multiple chains, can result in failures such as double spending [143] or counterfeited cryptocurrency-backed assets [187]. Another security issue arises with *data availability*. Protocols employing cross-chain verification via chain relays typically rely on timely arrival of proofs and metadata (block headers, transactions, ...). However, if an adversary can withhold this data from the verifying chain [35], such protocols not only become less efficient but potentially vulnerable to double spending and counterfeiting.

**Outlook.** This state of affairs calls for a rigorous and formal security analysis of existing CCC protocol – least those deployed in practice. In the meantime, ad-hoc solutions to the aforementioned security threats have been discussed in the community. For instance, protections against replay attacks involving the use of sequence numbers, or chains keeping track of previously processed proofs [143, 166, 69]. Similarly, first attempts to mitigate the data availability problem via erasure coding have been suggested in [35, 33, 184] – yet at the cost of protocol complexity and communication overhead.

#### 6.5 Lack of Formal Privacy Analysis.

**Problems.** Privacy is a crucial property of financial transactions and hence applies to CCC protocols. Ideally it should not be possible for an observer to determine which two events have been synchronized across chains (e.g., which assets have been exchanged and by whom). Unfortunately, CCC protocols deployed in practice lack formal privacy analysis and numerous privacy issues have already been detected. For instance, recent works [140, 101] leverage the fact that the same hash value is used on both chains involved in symmetric HTLC atomic

swaps to trivially link exchanged assets and accounts. Other de-anonymization techniques enabled by CCC protocols include miner address clustering via blocks merge-mined across different cryptocurrencies [116], cross-linking of miner and user accounts cross-chain by analyzing of blockchain forks [109, 170], and using public exchange datasets to trace cross-ledger trades [183].

**Outlook.** The academic community has developed formal frameworks that permit rigorous analysis of the privacy properties in the context of exchange protocols [101, 105, 172, 140, 141]. First techniques towards privacy-preserving CCC Exchange protocols via asymmetric and unlinkable locking techniques have been studied in [140, 141, 162, 79], yet, at the time of writing, we are not aware of privacy enhancements for the more-widely adopted Migration protocols – an interesting avenue for future research.

## 6.6 Upcoming Industrial and Research CCC Trends

**Interoperability Blockchains.** are specialized sharded distributed ledgers which aim to serve as communication layer between other blockchains [133, 180, 161, 167, 178, 110, 21] and exhibit implementations of existing CCC protocols. Individual shards, which are coordinated via a parent chain running a BFT agreement protocol, connect to and import assets from existing blockchains via Migration CCC protocols, most commonly cryptocurrency-backed assets [14]. A formal treatment of this design, also considering distributed computations, is presented in [138]. Cosmos [133] and Polkadot [180] also implement new standards for (internal) cross-shard communication (IBC [23] and XCMP [66] respectively). As of this writing, the aforementioned systems are under active development, making it difficult to argue about their security, feasibility, and long-term adoption - leaving room for future analysis.

**Efficient Light Clients.** Cross-chain state verification via chain relays is a fundamental part of robust CCC protocols. While current light/SPV clients suffice for e.g., mobile devices, they often remain infeasible for deployment on top of blockchains for CCC protocols, where storage and bandwidth are priced by the byte. Recent works on sub-linear light clients have achieved first significant theoretical [121, 139, 122] and practical performance improvements [76, 186, 179]. In parallel, recent developments in the field of zero-knowledge cryptography [57, 49, 64] pave the way towards (near)constant verification times and costs for chain relays. First schemes for blockchains with built-in support for such proof systems are put forth in [145, 59, 50].

**Off-Chain Protocols.** One of the most actively developed fields in blockchain research are off-chain (“L2”) communication networks [102], which aim to improve scalability (and privacy) of distributed ledgers: most transactions are executed off-chain and only channel opening and closure are written to the ledger. The influx of L2 solutions is thereby creating a new field for CCC research: (i) communication across off-chain channels [140, 141, 172, 105], and (ii) communication between off-chain and on-chain networks [22, 28]. While similar to conventional CCC protocols, the “off-chain” nature of L2 solutions requires more



complex techniques for the verification phase of CCC: intermediate states in off-chain protocols cannot be verified by existing chain relays, which only support verification of on-chain commitments, and must hence resort to cryptographic techniques such as adaptor signatures [41] or succinct proofs of knowledge [57, 49, 64].

## 7 Concluding Remarks

Our systematic analysis of cross-chain communication as a new problem in the era of distributed ledgers allows us to relate (mostly) community driven efforts to established academic research in database and distributed systems research. We formalize the cross-chain communication problem and show it cannot be solved without a trusted third party – contrary to the assumptions often made in the blockchain community. Following this result, we introduce a framework for evaluating existing and designing new cross-chain communication protocols, based on the inherent trust assumptions thereof. We then provide a classification and comparative evaluation, taking into account both academic research and the vast number of online resources, allowing us to better understand the similarities and differences between existing cross-chain communication approaches. Finally, by discussing implications and open challenges faced by cross-chain communication protocols, as well as the implications of interoperability on the security and privacy of blockchains, we offer a comprehensive guide for designing protocols, bridging multiple distributed ledgers.

## References

1. Binance exchange. Online. <https://www.binance.com/en>, Accessed: 2020-09-19.
2. Bitcoin Developer Guide: Simplified Payment Verification (SPV). <https://bitcoin.org/en/developer-guide#simplified-payment-verification-spv>. Accessed: 2018-05-16.
3. Bitcoin Wiki: Hashed Time-Lock Contracts. [https://en.bitcoin.it/wiki/Hashed\\_Timelock\\_Contracts](https://en.bitcoin.it/wiki/Hashed_Timelock_Contracts). Accessed: 2018-05-16.
4. Bitcoin wiki: Merged mining specification. [https://en.bitcoin.it/wiki/Merged\\_mining\\_specification](https://en.bitcoin.it/wiki/Merged_mining_specification). Accessed: 2018-05-03.
5. Btcrelay. <https://github.com/ethereum/btcrelay>. Accessed 2019-08-15.
6. Chainlink: A decentralized oracle network. Online. <https://link.smartcontract.com/whitepaper>, Accessed: 2020-09-19.
7. Confirmations. <https://en.bitcoin.it/wiki/Confirmation>. Accessed: 2018-11-28.
8. Dogerelay. <https://github.com/dogethereum/dogerelay>. Accessed 2019-08-15.
9. Eth-eos-relay. <https://github.com/EveripediaNetwork/eth-eos-relay>. Accessed 2019-08-15.
10. Ethereum contract allowing ether to be obtained with bitcoin. <https://github.com/ethers/EthereumBitcoinSwap>. Accessed: 2018-10-30.
11. Parity-Bridge. <https://github.com/paritytech/parity-bridge>. Accessed 2019-08-15.

12. The parity light protocol - wiki. [https://wiki.parity.io/The-Parity-Light-Protocol-\(PIP\)](https://wiki.parity.io/The-Parity-Light-Protocol-(PIP)). Accessed: 2018-10-30.
13. Peace relay. <https://github.com/loiluu/peacereley>. Accessed 2019-08-15.
14. Polkabtc: Trustless bitcoin on polkadot. Online. <https://github.com/interlay/BTC-Parachain> , Accessed: 2020-09-19.
15. Project alchemy. <https://github.com/ConsenSys/Project-Alchemy>. Accessed 2019-08-15.
16. Project waterloo. <https://blog.kyber.network/waterloo-a-decentralized-practical-bridge-between-eos-and-ethereum-1c230ac65524>. Accessed 2019-08-15.
17. tbtc: A decentralized redeemable btc-backed erc-20 token. <http://docs.keep.network/tbtc/index.pdf> . Accessed: 2019-11-15.
18. Wrapped bitcoin. <https://www.wbtc.network/assets/wrapped-tokens-whitepaper.pdf> . Accessed: 2018-05-03.
19. Alt chains and atomic transfers. [bitcointalk.org](http://bitcointalk.org), 2013.
20. Atomic swap. Bitcoin Wiki, 2013.
21. Wanchain whitepaper. <https://www.wanchain.org/files/Wanchain-Whitepaper-EN-version.pdf>, 2017.
22. Submarine swaps service. Online, 2018. <https://github.com/submarineswaps/swaps-service>.
23. Inter-blockchain communication protocol (ibc) specification. Online, 2019. <https://github.com/cosmos/ics/tree/master/ibc>.
24. Online, 2020. <https://coinmarketcap.com/rankings/exchanges/dex/>.
25. Online, 2020. <https://forum.etherbase.org/discussion/comment/53/Comment53>.
26. Bitcoin on ethereum. Online, 2020. <https://defipulse.com/btc>.
27. Bitcoin supply on ethereum tops \$1b. Coindesk, September 2020. <https://www.coindesk.com/bitcoin-supply-on-ethereum-tops-1b> .
28. Loop. Online, 2020. <https://lightning.engineering/loop/>.
29. Ptokens: How it works. Online, 2020. <https://ptokens.io/how-it-works> , Accessed: 2020-09-19.
30. Renvm. Online, 2020. <https://renproject.io/renvm> , Accessed: 2020-09-19.
31. I. Abraham, G. Gueta, and D. Malkhi. Hot-stuff the linear, optimal-resilience, one-message bft devil. *arXiv:1803.05069*, 2018.
32. J. Adler, R. Berryhill, A. Veneris, Z. Poulos, N. Veira, and A. Kastania. Astraea: A decentralized blockchain oracle. *arXiv preprint arXiv:1808.00528*, 2018.
33. M. Al-Bassam. Lazyledger: A distributed data availability ledger with client-side smart contracts, 2019.
34. M. Al-Bassam, A. Sonnino, S. Bano, D. Hryczyn, and G. Danezis. Chainspace: A sharded smart contracts platform. In *2018 Network and Distributed System Security Symposium (NDSS)*, 2018.
35. M. Al-Bassam, A. Sonnino, and V. Buterin. Fraud proofs: Maximising light client security and scaling blockchains with dishonest majorities. *CoRR*, abs/1809.09044, 2018.
36. E. Androulaki, C. Cachin, A. De Caro, and E. Kokoris-Kogias. Channels: Horizontal scaling and confidentiality on permissioned blockchains. In *European Symposium on Research in Computer Security*, pages 111–131. Springer, 2018.
37. M. Andrychowicz. Multiparty computation protocols based on cryptocurrencies, 2015. Accessed: 2017-02-15.
38. N. Asokan. Fairness in electronic commerce. 1998.
39. N. Asokan, V. Shoup, and M. Waidner. Asynchronous protocols for optimistic fair exchange. In *Proceedings. 1998 IEEE Symposium on Security and Privacy (Cat. No. 98CB36186)*, pages 86–99. IEEE, 1998.

40. N. Asokan, V. Shoup, and M. Waidner. Optimistic fair exchange of digital signatures. In *International Conference on the Theory and Applications of Cryptographic Techniques*, pages 591–606. Springer, 1998.
41. L. Aumayr, O. Ersoy, A. Erwig, S. Faust, K. Hostakova, M. Maffei, P. Moreno-Sanchez, and S. Riahi. Generalized bitcoin-compatible channels. *IACR Cryptol. ePrint Arch.*, 2020:476, 2020.
42. G. Avarikioti, L. Käppeli, Y. Wang, and R. Wattenhofer. Bitcoin security under temporary dishonest majority. In *23rd Financial Cryptography and Data Security (FC)*, 2019.
43. G. Avarikioti, E. K. Kogias, and R. Wattenhofer. Brick: Asynchronous state channels. *arXiv preprint arXiv:1905.11360*, 2019.
44. G. Avarikioti, E. Kokoris-Kogias, and R. Wattenhofer. Divide and scale: Formalization of distributed ledger sharding protocols. *arXiv preprint arXiv:1910.10434*, 2019.
45. G. Avarikioti, F. Laufenberg, J. Sliwinski, Y. Wang, and R. Wattenhofer. Towards secure and efficient payment channels. *arXiv preprint arXiv:1811.12740*, 2018.
46. O. Babaoglu and S. Toueg. Understanding non-blocking atomic commitment. *Distributed systems*, pages 147–168, 1993.
47. A. Back, M. Corallo, L. Dashjr, M. Friedenbach, G. Maxwell, A. Miller, A. Poelstra, J. Timón, and P. Wuille. Enabling blockchain innovations with pegged sidechains, 2014.
48. R. Belchior, A. Vasconcelos, S. Guerreiro, and M. Correia. A survey on blockchain interoperability: Past, present, and future trends. *arXiv preprint arXiv:2005.14282*, 2020.
49. E. Ben-Sasson, I. Bentov, Y. Horesh, and M. Riabzev. Scalable, transparent, and post-quantum secure computational integrity. *IACR Cryptology ePrint Archive*, 2018:46, 2018.
50. E. Ben-Sasson, A. Chiesa, and N. Spooner. Interactive oracle proofs. In *Theory of Cryptography Conference*, pages 31–60. Springer, 2016.
51. J. Benaloh and M. De Mare. One-way accumulators: A decentralized alternative to digital signatures. In *Workshop on the Theory and Application of Cryptographic Techniques*, pages 274–285. Springer, 1993.
52. P. Bennink, L. v. Gijtenbeek, O. v. Deventer, and M. Everts. An analysis of atomic swaps on and between ethereum blockchains using smart contracts. Tech. report, 2018. <https://work.delaat.net/rp/2017-2018/p42/report.pdf>.
53. I. Bentov, Y. Ji, F. Zhang, Y. Li, X. Zhao, L. Breidenbach, P. Daian, and A. Juels. Tesseract: Real-time cryptocurrency exchange using trusted hardware. *Cryptology ePrint Archive*, Report 2017/1153, 2017. Accessed:2017-12-04.
54. I. Bentov and R. Kumaresan. How to use bitcoin to design fair protocols. In *Advances in Cryptology—CRYPTO 2014*, pages 421–439. Springer, 2014.
55. I. Bentov, R. Pass, and E. Shi. Snow white: Provably secure proofs of stake, 2016. Accessed: 2016-11-08.
56. P. A. Bernstein, V. Hadzilacos, and N. Goodman. *Concurrency control and recovery in database systems*, volume 370. Addison-wesley New York, 1987.
57. N. Bitansky, R. Canetti, A. Chiesa, and E. Tromer. From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, pages 326–349. ACM, 2012.
58. D. Boneh, J. Bonneau, B. Bünz, and B. Fisch. Verifiable delay functions. In *CRYPTO*, 2018.

59. D. Boneh, B. Bünz, and B. Fisch. Batching techniques for accumulators with applications to iops and stateless blockchains. Cryptology ePrint Archive, Report 2018/1188, 2018. <https://eprint.iacr.org/2018/1188>.
60. D. Boneh and M. Naor. Timed commitments. In *Annual International Cryptology Conference*, pages 236–254. Springer, 2000.
61. J. Bonneau. Why buy when you can rent? bribery attacks on bitcoin consensus. In *BITCOIN '16: Proceedings of the 3rd Workshop on Bitcoin and Blockchain Research*, February 2016.
62. J. Bonneau, J. Clark, and S. Goldfeder. On bitcoin as a public randomness source, 2015. Accessed: 2015-10-25.
63. J. Bonneau, J. Clark, and S. Goldfeder. On bitcoin as a public randomness source. *IACR Cryptology ePrint Archive*, 2015:1015, 2015.
64. B. Bünz, J. Bootle, D. Boneh, A. Poelstra, P. Wuille, and G. Maxwell. Bulletproofs: Efficient range proofs for confidential transactions, 2017. Accessed:2017-11-10.
65. B. Bünz, S. Goldfeder, and J. Bonneau. Proofs-of-delay and randomness beacons in ethereum. 2017.
66. J. Burdges, A. Cevallos, P. Czaban, R. Habermeier, S. Hosseini, F. Lama, H. K. Alper, X. Luo, F. Shirazi, A. Stewart, et al. Overview of polkadot and its design considerations. *arXiv preprint arXiv:2005.13456*, 2020.
67. V. Buterin. Ethereum: A next-generation smart contract and decentralized application platform, 2014. Accessed: 2016-08-22.
68. V. Buterin. Chain interoperability. Tech. report, 2016. Accessed: 2017-03-25.
69. V. Buterin. Cross-shard contract yanking. <https://ethresear.ch/t/cross-shard-contract-yanking/1450>, 2018.
70. C. Cachin. Architecture of the hyperledger blockchain fabric, 2016. Accessed: 2016-08-10.
71. C. Cachin and J. Camenisch. Optimistic fair secure computation. In *Annual International Cryptology Conference*, pages 93–111. Springer, 2000.
72. M. Castro, B. Liskov, et al. Practical byzantine fault tolerance. In *OSDI*, volume 99, pages 173–186, 1999.
73. D. Catalano and D. Fiore. Vector commitments and their applications. In *International Workshop on Public Key Cryptography*, pages 55–72. Springer, 2013.
74. A. Chepurnoy, T. Duong, L. Fan, and H.-S. Zhou. Twinscoin: A cryptocurrency via proof-of-work and proof-of-stake, 2017. Accessed: 2017-03-22.
75. T. Chesney, I. Coyne, B. Logan, and N. Madden. Griefing in virtual worlds: causes, casualties and coping strategies. *Information Systems Journal*, 19(6):525–548, 2009.
76. S. Daveas, K. Karantias, A. Kiayias, and D. Zindros. A gas-efficient superlight bitcoin client in solidity. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, pages 132–144, 2020.
77. B. David, P. Gaži, A. Kiayias, and A. Russell. Ouroboros praos: An adaptively-secure, semi-synchronous proof-of-stake blockchain. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 66–98. Springer, 2018.
78. C. Decker and R. Wattenhofer. Bitcoin transaction malleability and mtgox. In *Computer Security-ESORICS 2014*, pages 313–326. Springer, 2014.
79. A. Deshpande and M. Herlihy. Privacy-preserving cross-chain atomic swaps. In *International Conference on Financial Cryptography and Data Security*, pages 540–549. Springer, 2020.
80. E. W. Dijkstra. Solution of a problem in concurrent programming control. In *Pioneers and Their Contributions to Software Engineering*, pages 289–294. Springer, 2001.

81. J. Dille, A. Poelstra, J. Wilkins, M. Piekarska, B. Gorlick, and M. Friedenbach. Strong federations: An interoperable blockchain solution to centralized third party risks. *arXiv preprint arXiv:1612.05491*, 2016.
82. J. R. Douceur. The sybil attack. In *International Workshop on Peer-to-Peer Systems*, pages 251–260. Springer, 2002.
83. T. Duong, L. Fan, and H.-S. Zhou. 2-hop blockchain: Combining proof-of-work and proof-of-stake securely. Cryptology ePrint Archive, Report 2016/716, 2016. Accessed: 2017-02-06.
84. S. Dziembowski, L. Eekey, and S. Faust. Fairswap: How to fairly exchange digital goods. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pages 967–984. ACM, 2018.
85. C. Egger, P. Moreno-Sanchez, and M. Maffei. Atomic multi-channel updates with constant collateral in bitcoin-compatible payment-channel networks. In *CCS*, 2019.
86. S. Even. A protocol for signing contracts. Technical report, Computer Science Department, Technion. Presented at CRYPTO’81, 1982.
87. S. Even and Y. Yacobi. Relations among public key signature systems. Technical report, Computer Science Department, Technion, 1980.
88. I. Eyal and E. G. Sirer. Majority is not enough: Bitcoin mining is vulnerable. In *Financial Cryptography and Data Security*, pages 436–454. Springer, 2014.
89. G. Fanti, L. Kogan, S. Oh, K. Ruan, P. Viswanath, and G. Wang. Compounding of wealth in proof-of-stake cryptocurrencies. *arXiv preprint arXiv:1809.07468*, 2018.
90. M. J. Fischer, N. A. Lynch, and M. S. Paterson. Impossibility of distributed consensus with one faulty process. volume 32, pages 374–382. ACM, 1985.
91. B. Ford, L. Gasser, E. K. Kogias, and P. Jovanovic. Cryptographically verifiable data structure having multi-hop forward and backwards links and associated systems and methods, Dec. 13 2018. US Patent App. 15/618,653.
92. A. Gabizon, K. Gurkan, P. Jovanovic, G. Konstantopoulos, A. Oines, M. Olszewski, M. Straka, and E. Tromer. Plumo: Towards scalable interoperable blockchains using ultra light validation systems. 2020.
93. J. A. Garay, A. Kiayias, and N. Leonardos. The bitcoin backbone protocol with chains of variable difficulty, 2016. Accessed: 2017-02-06.
94. A. Garoffolo, D. Kaidalov, and R. Oliynykov. Zendo: a zk-snark verifiable cross-chain transfer protocol enabling decoupled and decentralized sidechains. *arXiv preprint arXiv:2002.01847*, 2020.
95. F. C. Gärtner. Specifications for fault tolerance: A comedy of failures. 1998.
96. P. Gazi, A. Kiayias, and A. Russell. Stake-bleeding attacks on proof-of-stake blockchains. Cryptology ePrint Archive, Report 2018/248, 2018. Accessed:2018-03-12.
97. P. Gazi, A. Kiayias, and D. Zindros. Proof-of-stake sidechains. *IEEE Security and Privacy*. IEEE, 2019.
98. R. Gennaro, S. Goldfeder, and A. Narayanan. Threshold-optimal dsa/ecdsa signatures and an application to bitcoin wallet security. In *International Conference on Applied Cryptography and Network Security*, pages 156–174. Springer, 2016.
99. A. Gervais, G. O. Karame, K. Wüst, V. Glykantzis, H. Ritzdorf, and S. Capkun. On the security and performance of proof of work blockchains. In *Proceedings of the 2016 ACM SIGSAC*, pages 3–16. ACM, 2016.
100. O. Goldreich. Secure multi-party computation. *Manuscript. Preliminary version*, 78, 1998.
101. M. Green and I. Miers. Bolt: Anonymous payment channels for decentralized currencies. Cryptology ePrint Archive, Report 2016/701, 2016. Accessed: 2017-08-07.

102. L. Gudgeon, P. Moreno-Sanchez, S. Roos, P. McCorry, and A. Gervais. Sok: Off the chain transactions. Cryptology ePrint Archive, Report 2019/360, 2019. <https://eprint.iacr.org/2019/360>.
103. R. Han, H. Lin, and J. Yu. On the optionality and fairness of atomic swaps. Cryptology ePrint Archive, Report 2019/896, 2019. <https://eprint.iacr.org/2019/896>.
104. D. Harz and M. Boman. The scalability of trustless trust. arXiv:1801.09535, 2018. Accessed:2018-01-31.
105. E. Heilman, L. Alshenibr, F. Baldimtsi, A. Scafuro, and S. Goldberg. Tumblebit: An untrusted bitcoin-compatible anonymous payment hub, 2016. Accessed: 2017-09-29.
106. E. Heilman, S. Lipmann, and S. Goldberg. The arwen trading protocols. Whitepaper. <https://www.arwen.io/whitepaper.pdf>.
107. M. Herlihy. Atomic cross-chain swaps. arXiv:1801.09515, 2018. Accessed:2018-01-31.
108. M. Herlihy, B. Liskov, and L. Shrira. Cross-chain deals and adversarial commerce. *arXiv preprint arXiv:1905.09743*, 2019.
109. A. Hinteregger and B. Haslhofer. An empirical analysis of monero cross-chain traceability. *arXiv preprint arXiv:1812.02808*, 2018.
110. D. Hosp, T. Hoenisch, P. Kittiwongsunthorn, et al. Comit-cryptographically-secure off-chain multi-asset instant transaction network. *arXiv preprint arXiv:1810.02174*, 2018.
111. K. Itakura and K. Nakamura. A public-key cryptosystem suitable for digital multisignatures. *NEC Research & Development*, (71):1–8, 1983.
112. D. Johnson, A. Menezes, and S. Vanstone. The elliptic curve digital signature algorithm (ecdsa). *International journal of information security*, 1(1):36–63, 2001.
113. S. Johnson, P. Robinson, and J. Brainard. Sidechains and interoperability. *arXiv preprint arXiv:1903.04077*, 2019.
114. J. Jones and abitmore. Optional htlc preimage length and hash160 addition. BSIP 64, blog post. <https://github.com/bitshares/bsips/issues/163>.
115. A. Judmayer, N. Stifter, A. Zamyatin, I. Tsabary, I. Eyal, P. Gaži, S. Meiklejohn, and E. Weippl. Pay-to-win: Incentive attacks on proof-of-work cryptocurrencies. Cryptology ePrint Archive, Report 2019/775, 2019. <https://eprint.iacr.org/2019/775>.
116. A. Judmayer, A. Zamyatin, N. Stifter, A. G. Voyiatzis, and E. Weippl. Merged mining: Curse or cure? In *CBT'17: Proceedings of the International Workshop on Cryptocurrencies and Blockchain Technology*, Sep 2017.
117. H. Kalodner, S. Goldfeder, X. Chen, S. M. Weinberg, and E. W. Felten. Arbitrum: Scalable, private smart contracts. In *Proceedings of the 27th USENIX Conference on Security Symposium*, pages 1353–1370. USENIX Association, 2018.
118. K. Karantias, A. Kiayias, and D. Zindros. Proof-of-burn. In *International Conference on Financial Cryptography and Data Security*, pages 523–540. Springer, 2020.
119. M. Khabbazian, T. Nadahalli, and R. Wattenhofer. Outpost: A responsive lightweight watchtower. 2019.
120. A. Kiayias, N. Lamprou, and A.-P. Stouka. Proofs of proofs of work with sublinear complexity. In *International Conference on Financial Cryptography and Data Security*, pages 61–78. Springer, Springer, 2016.
121. A. Kiayias, A. Miller, and D. Zindros. Non-interactive proofs of proof-of-work. Cryptology ePrint Archive, Report 2017/963, 2017. Accessed:2017-10-03.
122. A. Kiayias, A. Polydouri, and D. Zindros. The Velvet Path to Superlight Blockchain Clients, 2020.
123. A. Kiayias, A. Russell, B. David, and R. Oliynykov. Ouroboros: A provably secure proof-of-stake blockchain protocol. In *Annual International Cryptology Conference*, pages 357–388. Springer, 2017.

124. A. Kiayias, H.-S. Zhou, and V. Zikas. Fair and robust multi-party computation using a global transaction ledger. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 705–734. Springer, 2016.
125. A. Kiayias and D. Zindros. Proof-of-work sidechains. In *International Conference on Financial Cryptography and Data Security*. Springer, 2018.
126. E. K. Kogias, P. Jovanovic, N. Gailly, I. Khoffi, L. Gasser, and B. Ford. Enhancing bitcoin security and performance with strong consistency via collective signing. In *25th USENIX Security Symposium (USENIX Security 16)*, Austin, TX, Aug. 2016. USENIX Association.
127. E. Kokoris-Kogias. Robust and scalable consensus for sharded distributed ledgers. Technical report, Cryptology ePrint Archive, Report 2019/676, 2019.
128. E. Kokoris-Kogias, E. C. Alp, S. D. Siby, N. Gailly, L. Gasser, P. Jovanovic, E. Syta, and B. Ford. Calypso: Auditable sharing of private data over blockchains. Technical report, Cryptology ePrint Archive, Report 2018/209, 2018.
129. E. Kokoris-Kogias, P. Jovanovic, L. Gasser, N. Gailly, E. Syta, and B. Ford. Omniledger: A secure, scale-out, decentralized ledger via sharding. In *2018 IEEE Symposium on Security and Privacy (SP)*, pages 583–598. IEEE, 2018.
130. L. Kokoris-Kogias, L. Gasser, I. Khoffi, P. Jovanovic, N. Gailly, and B. Ford. Managing identities using blockchains and CoSi. In *9th Workshop on Hot Topics in Privacy Enhancing Technologies (HotPETs 2016)*, 2016.
131. R. Kumaresan and I. Bentov. Amortizing secure computation with penalties. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 418–429. ACM, 2016.
132. A. K p c  and A. Lysyanskaya. Usable optimistic fair exchange. *Computer Networks*, 56(1):50–63, 2012.
133. J. Kwon and E. Buchman. Cosmos: A network of distributed ledgers. <https://github.com/cosmos/cosmos/blob/master/WHITEPAPER.md>, 2015.
134. Y. Kwon, H. Kim, J. Shin, and Y. Kim. Bitcoin vs. bitcoin cash: Coexistence or downfall of bitcoin cash? *arXiv:1902.11064*, 2019.
135. L. Lamport. A simple approach to specifying concurrent systems. *Communications of the ACM*, 32(1):32–45, 1989.
136. S. Lerner. Drivechains, sidechains and hybrid 2-way peg designs. Technical report, Tech. Rep. [Online], 2018.
137. S. D. Lerner. Rootstock: Bitcoin powered smart contracts. [https://docs.rsk.co/RSK\\_White\\_Paper-Overview.pdf](https://docs.rsk.co/RSK_White_Paper-Overview.pdf), 2015.
138. Z. Liu, Y. Xiang, J. Shi, P. Gao, H. Wang, X. Xiao, and Y.-C. Hu. Hyperservice: Interoperability and programmability across heterogeneous blockchains. *arXiv preprint arXiv:1908.09343*, 2019.
139. L. Luu, B. Buenz, and M. Zamani. Flyclient super light client for cryptocurrencies. Accessed 2018-04-17.
140. G. Malavolta, P. Moreno-Sanchez, A. Kate, M. Maffei, and S. Ravi. Concurrency and privacy with payment-channel networks. In *CCS*, pages 455–471, 2017.
141. G. Malavolta, P. Moreno-Sanchez, C. Schneidewind, A. Kate, and M. Maffei. Anonymous multi-hop locks for blockchain scalability and interoperability. In *NDSS*, 2019.
142. P. McCorry, S. Bakshi, I. Bentov, A. Miller, and S. Meiklejohn. Pisa: Arbitration outsourcing for state channels. *IACR Cryptology ePrint Archive*, 2018:582, 2018.
143. P. McCorry, E. Heilman, and A. Miller. Atomically trading with roger: Gambling on the success of a hardfork. In *CBT'17: Proceedings of the International Workshop on Cryptocurrencies and Blockchain Technology*, Sep 2017.

144. P. McCorry, A. Hicks, and S. Meiklejohn. Smart contracts for bribing miners. In *5th Workshop on Bitcoin and Blockchain Research, Financial Cryptography and Data Security 18 (FC)*. Springer, 2018.
145. I. Meckler and E. Shapiro. Coda: Decentralized cryptocurrency at scale. <https://cdn.codaprotocol.com/v2/static/coda-whitepaper-05-10-2018-0.pdf>, 2018.
146. R. C. Merkle. A digital signature based on a conventional encryption function. In *Conference on the Theory and Application of Cryptographic Techniques*, pages 369–378. Springer, 1987.
147. D. Meshkov, A. Chepurnoy, and M. Jansen. Revisiting difficulty control for blockchain systems. Cryptology ePrint Archive, Report 2017/731, 2017. Accessed: 2017-08-03.
148. S. Micali. Algorand: The efficient and democratic ledger, 2016. Accessed: 2017-02-09.
149. A. Miller. The high-value-hash highway, bitcoin forum post, 2012.
150. M. Miraz and D. C. Donald. Atomic cross-chain swaps: Development, trajectory and potential of non-monetary digital token swap facilities. *Annals of Emerging Technologies in Computing (AETiC) Vol, 3*, 2019.
151. P. Moreno-Sanchez, Randomrun, D. V. Le, S. Noether, B. Goodell, and A. Kate. Dlsag: Non-interactive refund transactions for interoperable payment channels in monero. Cryptology ePrint Archive, Report 2019/595, 2019. <https://eprint.iacr.org/2019/595>.
152. S. Nakamoto. Bitcoin: A peer-to-peer electronic cash system, Dec 2008. Accessed: 2015-07-01.
153. K. Nikitin, E. Kokoris-Kogias, P. Jovanovic, N. Gailly, L. Gasser, I. Khoffi, J. Cappos, and B. Ford. CHAINIAC: Proactive software-update transparency via collectively signed skipchains and verified builds.
154. S. Noether. Discrete logarithm equality across groups. Online, 2020. <https://www.getmonero.org/resources/research-lab/pubs/MRL-0010.pdf>.
155. H. Pagnia and F. C. Gärtner. On the impossibility of fair exchange without a trusted third party. Technical report, Technical Report TUD-BS-1999-02, Darmstadt University of Technology . . . , 1999.
156. R. Pass and E. Shi. Hybrid consensus: Scalable permissionless consensus, Sep 2016. Accessed: 2016-10-17.
157. A. Poelstra. Scriptless scripts. Presentation slides. <https://download.wpsoftware.net/bitcoin/wizardry/mw-slides/2017-03-mit-bitcoin-expo/slides.pdf>.
158. J. Poon and T. Dryja. The bitcoin lightning network, 2016. Accessed: 2016-07-07.
159. H. Ritzdorf, K. Wüst, A. Gervais, G. Felley, et al. Tls-n: Non-repudiation over tls enabling ubiquitous content signing. In *Network and Distributed System Security Symposium (NDSS)*, 2018.
160. R. L. Rivest, A. Shamir, and D. A. Wagner. Time-lock puzzles and timed-release crypto. 1996.
161. T. Rocket. Snowflake to avalanche: A novel metastable consensus protocol family for cryptocurrencies. *Available [online].[Accessed: 4-12-2018]*, 2018.
162. J. Rubin, M. Naik, and N. Subramanian. Merkelized abstract syntax trees. <http://www.mit.edu/~jlrubin/public/pdfs/858report.pdf>, 2014.
163. A. Sapirshstein, Y. Sompolinsky, and A. Zohar. Optimal selfish mining strategies in bitcoin, 2015. Accessed: 2016-08-22.
164. C.-P. Schnorr. Efficient signature generation by smart cards. *Journal of cryptology*, 4(3):161–174, 1991.
165. V. A. Siris, D. Dimopoulos, N. Fotiou, S. Voulgaris, and G. C. Polyzos. Interledger smart contracts for decentralized authorization to constrained things, 2019.



166. A. Sonnino, S. Bano, M. Al-Bassam, and G. Danezis. Replay attacks and defenses against cross-shard consensus in sharded distributed ledgers. *arXiv preprint arXiv:1901.11218*, 2019.
167. M. Spoke and Nuco Engineering Team. Aion: The third-generation blockchain network. [https://aion.network/media/2018/03/aion.network.technical-introduction\\_en.pdf](https://aion.network/media/2018/03/aion.network.technical-introduction_en.pdf). Accessed 2018-04-17.
168. A. Stewart and E. Kokoris-Kogia. Grandpa: a byzantine finality gadget. *arXiv preprint arXiv:2007.01560*, 2020.
169. I. Stewart. Proof of burn, 2012. Accessed: 2017-05-10.
170. N. Stifter, P. Schindler, A. Judmayer, A. Zamyatin, A. Kern, and E. Weippl. Echoes of the past: Recovering blockchain metrics from merged mining. In *Proceedings of the 23rd International Conference on Financial Cryptography and Data Security (FC)*. Springer, 2019.
171. P. Syverson. Weakly secret bit commitment: Applications to lotteries and fair exchange. In *Proceedings. 11th IEEE Computer Security Foundations Workshop (Cat. No. 98TB100238)*, pages 2–13. IEEE, 1998.
172. E. Tairi, P. Moreno-Sanchez, and M. Maffei. A<sup>2</sup>l: Anonymous atomic locks for scalability and interoperability in payment channel hubs. Cryptology ePrint Archive, Report 2019/589, 2019. <https://eprint.iacr.org/2019/589>.
173. J. Teutsch and C. Reitwießner. A scalable verification solution for blockchains, March 2017. Accessed:2017-10-06.
174. J. Teutsch, M. Straka, and D. Boneh. Retrofitting a two-way peg between blockchains. Technical report, 2018.
175. J. Teutsch and TrueBit Establishment. On decentralized oracles for data availability. 2017.
176. S. Thomas and E. Schwartz. A protocol for interledger payments. *URL https://interledger.org/interledger.pdf*, 2015.
177. S. A. K. Thyagarajan and G. Malavolta. Lockable signatures for blockchains: Scriptless scripts for all signatures. Cryptology ePrint Archive, Report 2020/1613, 2020. <https://eprint.iacr.org/2020/1613>.
178. G. Verdian, P. Tasca, C. Paterson, and G. Mondelli. Quant overledger whitepaper. <https://www.quant.network/>, 2018.
179. M. Westerkamp and J. Eberhardt. zkrelay: Facilitating sidechains using zksnark-based chain-relays. *Contract*, 1(2):3, 2020.
180. G. Wood. Polkadot: Vision for a heterogeneous multi-chain framework. *White Paper*, 2015.
181. G. Wood. Ethereum: A secure decentralised generalised transaction ledger eip-150 revision (759dced - 2017-08-07), 2017. Accessed: 2018-01-03.
182. A. C.-C. Yao. How to generate and exchange secrets. In *27th Annual Symposium on Foundations of Computer Science (sfcs 1986)*, pages 162–167. IEEE, 1986.
183. H. Yousaf, G. Kappos, and S. Meiklejohn. Tracing transactions across cryptocurrency ledgers. In *28th {USENIX} Security Symposium ({USENIX} Security 19)*, pages 837–850, 2019.
184. M. Yu, S. Sahraei, S. Li, S. Avestimehr, S. Kannan, and P. Viswanath. Coded merkle tree: Solving data availability attacks in blockchains. *arXiv preprint arXiv:1910.01247*, 2019.
185. M. Zamani, M. Movahedi, and M. Raykova. Rapidchain: A fast blockchain protocol via full sharding. Cryptology ePrint Archive, Report 2018/460, 2018.
186. A. Zamyatin, Z. Avarikioti, D. Perez, and W. J. Knottenbelt. Txchain: Efficient cryptocurrency light clients via contingent transaction aggregation. Sep 2020.

187. A. Zamyatin, D. Harz, J. Lind, P. Panayiotou, A. Gervais, and W. Knottenbelt. Xclaim: Trustless, interoperable, cryptocurrency-backed assets. *IEEE Security and Privacy. IEEE*, 2019.
188. F. Zhang, S. K. D. Maram, H. Malvai, S. Goldfeder, and A. Juels. Deco: Liberating web data using decentralized oracles for tls. *arXiv preprint arXiv:1909.00938*, 2019.

## A Strong Fair Exchange Definition

This section provides the definition of the strong Fair Exchange problem, as presented in [38–40, 155].

Fair Exchange considers two processes (or parties)  $P$  and  $Q$  that wish to exchange two items (or asset):  $a_P$  owned by  $P$  against  $a_Q$  owned by  $Q$ . There exists a function  $desc$  that maps any exchangeable item (or asset) to a string describing it in "sufficient" detail (e.g. the value and recipient of a payment). The inputs of  $P$  to a Fair Exchange protocol are an item  $a_P$  and a description  $d_Q$  of the desired item. Analogous, the inputs for  $Q$  are  $a_Q$  and  $d_P$ . To indicate that  $P$  is dishonest, an (boolean) error variable  $m_P$  is introduced (analogous,  $m_Q$  for  $Q$ ) [95]. A successful Fair Exchange is shown in Table 2 below.

Table 2: A successful Fair Exchange, as defined in [38–40, 155].

|  |               |  |
|--|---------------|--|
| <p><b>P</b></p> <p><b>Input</b> : <math>a_P, d_Q, Q</math></p> |               | <p><b>Q</b></p> <p><b>Input</b> : <math>a_Q, d_P, P</math></p> |
|  | fair exchange |  |
|  | ← →           |  |
| <b>Output</b> : $a_Q(desc(a_Q) = d_Q)$                         |               | <b>Output</b> : $a_P(desc(a_P) = d_P)$                         |
|  | or            |  |
| <i>aborted</i>   |               | <i>aborted</i>   |

A successful Fair Exchange protocol must thereby fulfill the following properties:

**Definition 7 (Effectiveness).** *If both  $P$  and  $Q$  behave correctly, i.e.,  $m_P = m_Q = \text{false}$ , and the items  $a_P$  and  $a_Q$  match the expected descriptions, i.e.,  $desc(a_Q) = d_Q \wedge desc(a_P) = d_P$ , then  $P$  will receive  $a_Q$  and  $Q$  will receive  $a_P$ . If the items are not as expected, i.e.,  $desc(a_Q) \neq d_Q \vee desc(a_P) \neq d_P$ , then both parties will abort the exchange.*

**Definition 8 (Timeliness).** *Eventually  $P$  will transfer  $a_P$  to  $Q$  or abort, and  $Q$  will transfer  $a_Q$  to  $P$  or abort.*

**Definition 9 (Strong Fairness).** *There are no outcomes in which  $Q$  receives  $a_P$  but  $P$  does not receive  $a_Q$  ( $Q$  aborts), or  $P$  receives  $a_Q$  but  $Q$  does not receive  $a_P$  ( $P$  aborts).*

Effectiveness determines the outcome of the exchange if  $P$  and  $Q$  are willing to perform the exchange and the item's match the expected descriptions, or the items do not match the expected descriptions. (Strong) Fairness restricts the outcomes of the exchange such that neither party is left at a disadvantage. Timeliness ensures eventual termination of the exchange protocol. Note: we do not provide a definition for "Non-repudiability" as this property is not a critical requirement for Fair Exchange protocols, but only becomes relevant in disputes after an exchange [39, 155].

## B Fair Exchange using CCC

We provide the intuition of how to construct a Fair Exchange protocol using a generic CCC protocol in Algorithm 1. Specifically,  $P$  and  $Q$  exchange assets  $a_P$  and  $a_Q$ , if transaction  $\text{TX}_P$  is written to  $L_x$  and transaction  $\text{TX}_Q$  is written to  $L_y$  (cf. Section 3.2).

---

### Algorithm 1: Fair Exchange using a generic CCC protocol

---

**Result:**  $\text{TX}_P \in L_x \wedge \text{TX}_Q \in L_y$  (i.e.,  $P$  has  $a_P$ ,  $Q$  has  $a_Q$ ) **or**  
 $\text{TX}_P \notin L_x \wedge \text{TX}_Q \notin L_y$  (i.e., no exchange)  
 setup( $L_x, L_y, \text{TX}_P, \text{TX}_Q, d_P, d_Q$ );  
**if**  $m_P = \text{false}$  **then**  
 | *commit*( $\text{TX}_P, L_x$ ); //  $P$  transfers  $a_P$  to  $Q$   
**end**  
**if** (*verify*( $\text{TX}_P, L_x, d_P$ ) = **true**)  $\wedge$   $m_Q = \text{false}$  **then**  
 | *commit*( $\text{TX}_Q, L_y$ ); //  $Q$  transfers  $a_Q$  to  $P$   
**else**  
 | *abort*( $\text{TX}_Q, L_y$ ); //  $Q$  does not transfer  $a_Q$  to  $P$   
**end**  
**if** *verify*( $\text{TX}_Q, L_y, d_Q$ ) = **false** **then**  
 | *abort*( $\text{TX}_P, L_x$ ); //  $P$  recovers  $a_P$   
**end**

---



---

### Algorithm 2: Commit( $\text{TX}, L$ )

---

**if** *valid*( $\text{TX}, L$ ) **then**  
 | Write  $\text{TX}$  to  $L$ ;  
**end**

---



---

### Algorithm 3: Verify( $\text{TX}, L, d$ )

---

**if**  $\text{TX} \in L \wedge \text{desc}(\text{TX}) = d$  **then**  
 | return **true**;  
**end**  
 return **false**;

---

---

**Algorithm 4:** Abort( $\text{TX}, \text{L}$ )

---

```

if  $\text{TX} \in \text{L}$  then
  | Revert  $\text{TX}$ ; // e.g. using a new transaction
else
  | //do nothing
end

```

---

## C Classification of Cross-Chain State Verification

As described in Section 2.5, a critical component of cross-chain communication is the verification of the state “evolution” of a chain  $X$  from within another chain  $Y$ , i.e., that  $X$  is in a certain state after the commit step. In this section we discuss the different elements of the chain that can be verified during the process, to complement the process of verifying state evolution. We show that there is a classification for what is verified (Section C.1), overview existing techniques for each class, and discuss the relation between the verification classes (Section C.2).

### C.1 Verification Classes

If a party  $P$  on  $X$  is misbehaving, it may withhold information from a party  $Q$  on  $Y$  (i.e., not submit a proof), but it should not be able to trick  $Q$  into accepting an incorrect state of  $\text{L}_x$  (e.g., convince  $Q$  that  $\text{TX}_1 \in \text{L}_x$  although  $\text{TX}_1$  was never written).

**Verification of State.** The simplest form of cross-chain verification is to check whether a specific state *exists*, i.e., is reachable but has not necessarily been agreed upon by the consensus participants. A representative example is the verification of Proof-of-Work in merged mining[4, 116]: the child chain  $Y$  only parses a given  $X$  block and verifies that the hash of the  $Y$  candidate block was included, and checks that the PoW hash exceeds the difficulty target of  $Y$ . Note that  $Y$  does not care whether the block is actually part of  $\text{L}_x$ . Another example is the use of blockchains as a public source of randomness [63, 65, 83, 74].

**Verification of State Agreement.** In addition to the existence of a state, a proof can provide sufficient data to verify that consensus has been achieved on that state. Typically, the functionality of this verification is identical to that of blockchain light clients [152, 2, 12]: instead of verifying the entire blockchain of  $X$ , all block headers and only transactions relevant to the CCC protocol are verified (and stored) on  $Y$ . The assumption thereby is that an invalid block will not be included in the verified blockchain under correct operation [152, 139]. Block headers can be understood as the meta-data for the block, including a commitment to all the transactions in the block, which are typically referenced using a vector commitment [73] (or some other form of cryptographic accumulator [51]), e.g. Merkle trees[146]. We discuss how proofs of state agreement differ depending on the underlying consensus mechanism below (non-exhaustive):

- **Proof-of-Work.** To verify agreement in PoW blockchains, a primitive called (*Non-interactive*) *Proofs of Proof-of-Work* [120, 121], also referred to as SPV (simplified payment verification) [152] is used. Thereby, the verifier of a proof must at least check for each block that (i) the PoW meets the specified difficulty target, (ii) the specified target is in accordance with the difficulty adjustment and (iii) the block contains a reference to the previous block in the chain [2, 187]. The first known implementation of cross-chain state agreement verification (for PoW blockchains) is BTCRelay [5]: a smart contract which allows to verify the state of Bitcoin on Ethereum<sup>3</sup>.
- **Proof-of-Stake.** If the verified chain uses Proof-of-Stake in its consensus, the proofs represent a dynamic collection of signatures, capturing the underlying stake present in the chain. These are referred to as *Proofs of Proof-of-Stake* (PoPoS) and a scheme in this direction was put forth in [97].
- **BFT.** In case the blockchain is maintained by a BFT committee, the cross-chain proofs are simplified and take the form of a sequence of signatures by  $2f + 1$  members of the committee, where  $f$  is the number of faulty nodes that can be tolerated [72]. If the committee membership is dynamically changing, the verification process needs to capture the rotating configuration of the committee [153].

*Sub-linear State Agreement Proofs.* Verifying all block headers results in proof complexity linear in the size of the blockchain. However, there exist techniques for achieving *sub-linear* (logarithmic in the size of the chain) complexity, which rely on probabilistic verification. For PoW blockchains, we are aware of two approaches: Superblock (Ni)PoPoWs [149, 47, 120, 121] and FlyClient [139]. Both techniques rely on probabilistic sampling but differ in the selection heuristic. Superblock (Ni)PoPoWs sample blocks which exceed the required PoW difficulty<sup>4</sup>, i.e., randomness is sourced from the unpredictability of the mining process, whereas FlyClient suggests to sample blocks using an optimized heuristic after the chain has been generated (using randomness from the PoW hashes [63]). For blockchains maintained by a static BFT committee, the verified signatures can be combined into aggregate signatures [126, 127] for optimization purposes. These signature techniques are well known and have been invented prior to blockchains, and we hence do not elaborate further on these schemes. In the dynamic setting, skipchains [91, 153, 130], i.e., double-linked skiplists which enable sub-linear crawling of identity chains, can reduce costs from linear to logarithmic (to the number of configurations). Recently, a number for light client protocols that leverage the compression properties of zero-knowledge proof systems have been proposed [94, 92, 179].

**Verification of State Evolution.** Once verified by some chain  $Y$  that chain  $X$  has reached agreement on a ledger state  $L_x[r]$ , it is then possible for (users on)  $Y$  to verify that certain transactions have been included in  $L_x$ . As mentioned, block headers typically reference included transactions via vector commitments.

<sup>3</sup> Similar contracts have been proposed for other chains [15, 13, 8, 16, 11, 9].

<sup>4</sup> It is a property of the PoW mining process that a certain percentage of blocks exceeds or fall short of the required difficulty.

As such, to verify that  $\text{TX} \in L_x[r]$  the vector commitment on  $L_x[r]$  needs to be opened at the index of that transaction, e.g. by providing a Merkle tree path to the leaf containing TX (e.g. as in Bitcoin). Thereby, multiple transactions can be aggregated in a single proof [186].

**Verification of State Validity.** Even though a block is believed to have consensus, it may not be a valid block if it contains invalid transactions or state transitions (e.g. a PoW block meeting the difficulty requirements, but containing conflicting transactions). Fully validating nodes will reject these blocks as they check all included transactions. However, in the case of cross-chain communication, where chains typically only verify state agreement but not validity, detection is not directly possible. We classify two categories of techniques to enable such chains, and non-full nodes (i.e., light clients), to reject invalid blocks:

- In **proactive** state validation, nodes ensure that blocks are valid before accepting them. Apart from requiring participants to run fully validating nodes, this can be achieved by leveraging “validity proofs” through succinct proofs of knowledge, using systems such as SNARKs [57], STARKs [49] or Bulletproofs [64]. First schemes for blockchains offering such proofs for each state transition are put forth in [145, 59, 50]. Informally speaking, this is a “guilty until proven innocent model”: nodes assume blocks that have consensus are invalid until proven otherwise.
- In **reactive** state validation, nodes follow an “innocent until proven guilty” model. It is assumed that blocks that have consensus only contain valid state transition, unless a state transition “fraud proof” [35] is created. Fraud proofs typically are proofs of state evolution, i.e., opening of the transaction vector commitment in the invalid block at the index of the invalid transaction, e.g. via Merkle tree paths. Depending on the observed failure, more data may be necessary to determine inconsistencies (e.g. Merkle paths for conflicting transactions in a double spend).

**Verification of Data Availability.** Consensus participants may produce a block header, but not release the included transactions, preventing other participants from validating the correctness of the state transition. To this end, verification of state validity can be complimented by verification of data availability. A scheme for such proofs was put forth in [35, 184], which allows to verify with high probability that all the data in a block is available, by sampling chunks in an erasure-coded version of a block.

## C.2 Relation between Verification Classes

Verification of State Agreement requires to first verify a specific state exists or has been proposed (Verification of State). To verify a transaction was included at  $L[r]$  (State Evolution), it is first necessary to verify that the ledger state at  $L[r]$  is indeed agreed upon (State Agreement). Finally, to verify that a state (transition) is indeed valid (State Validity), one must first verify that all associated transactions were indeed included in the ledger (State Evolution). Verification of

Data Availability serves as complimentary security measure, and can be added to any of the classes to protect against data withholding attacks. We illustrate this relationship in Figure 2.




Fig. 2: Venn diagram of cross-chain state verification classes. The red, dotted line highlights the minimum requirement for correctly operating light clients, i.e., verifying SPV / NiPoPoWs in the case of PoW blockchains.

## D Notation

Below is a table of notation used globally throughout the paper.

| Notation     | Meaning  |
|--------------|--|
| $X$          | System X   |
| $Y$          | System Y   |
| $L_x$        | Ledger X   |
| $L_y$        | Ledger Y   |
| $TX_L^{\in}$ | Transaction in $L$   |
| $TX$         | Transaction  |
| $P$ and $Q$  | Processes (parties)  |
| $a_P$        | Asset owned by P   |
| $a_Q$        | Asset owned by Q   |
| $L[r]$       | State of ledger $L$ at round $r$                                       |
| $L^P[r]$     | State of ledger $L$ at round $r$ in the view of P                      |
| $L^Q[r]$     | State of ledger $L$ at round $r$ in the view of Q                      |
| $desc(TX)$   | Description of TX ( <i>e.g.</i> , the transaction value and recipient) |
| $r$          | Ledger round   |
| $u$          | Ledger liveness delay parameter  |
| $k$          | Ledger persistence delay / “depth” parameter                           |



## E Acknowledgements

We would like express our gratitude to Georgia Avarikioti, Daniel Perez and Dominik Harz for helpful comments and feedback on earlier versions of this manuscript. We also thank Nicholas Stifter, Aljosha Judmayer, Philipp Schindler, Edgar Weippl, and Alistair Stewart for insightful discussions during the early stages of this research. We also wish to thank the anonymous reviewers for their valuable comments that helped improve the presentation of our results.

This research was funded by Bridge 1 858561 SESC; Bridge 1 864738 PR4DLT (all FFG); the Christian Doppler Laboratory for Security and Quality Improvement in the Production System Lifecycle (CDL-SQI); the competence center SBA-K1 funded by COMET; Chaincode Labs through the project SLN: Scalability for the Lightning Network; and by the Austrian Science Fund (FWF) through the Meitner program (project M-2608).

Mustafa Al-Bassam is funded by a scholarship from the Alan Turing Institute. Alexei Zamyatin conducted the early stages of this work during his time at SBA Research, and was supported by a Binance Research Fellowship.