

Contingent Payments

Mathias Hall-Andersen¹[0000-0002-0195-6659]

Aarhus University, mathias@hall-andersen.dk

This document serves as an introduction and extended motivation, the goal is to make intelligible and motivate the remaining article for people with an MSc in computer science. Therefore this document serves as a brief introduction to the problem statement, the notion of cryptographic assumptions, a crash course on the universal composability framework and the setting of FastSwap.

1 Introduction

1.1 Prisoners Dilemma

Consider the experience of buying/selling license keys over the internet:

Seller: Imagine you are running a shop which sells license keys to a wide selection of commercial software. A customer buys a license key for a piece of office software, he pays using Visa, you send him the license key. However a few days later he disputes the payment, saying that the key did not work. You know this to be false.

Buyers: Now suppose you wish to buy a license to a piece of commercial software. You go online, find a seeming legitimate looking website and buy a license key using Bitcoin. However when you receive the key, it does not work. You try to contact the seller, but no response.

What is described above is a modern variant of the classical prisoners dilemma: whichever party moves first is at risk of being cheated: the buyer can finalize the payment (using Bitcoin) and get scammed by the seller, or the seller can finalize the sale (sending the key) and get scammed by the buyer. What largely saves people in modern society from the constant experience of being defrauded by strangers is the existence of strong judicial systems which can enforce harsh punishments that far outweighs the limited gain from cheating in the interaction: essentially by arbitrarily imposing a cost to defecting in the game. However, this approach is ill suited if:

1. The participants cannot agree on a judicial system to trust, this issue compounds if the stakes are very high, which makes corrupting the judicial systems a more viable strategy.
2. The participants are anonymous, such that no penalty can be enforced.

The latter is especially devastating since it even curtails the possibility of locally keeping a score of ‘good’ and ‘bad’ merchants/customers.

1.2 Contingent Payments

In this thesis we, in particular, suggest a practical scheme for solving these kinds of ‘Contingent Payments’ where payment occurs if and only if the input (e.g. serial key) satisfies a predicate (e.g. serial key check algorithm). Further examples of such problems might be:

- Buying proofs. The predicate consists of a formal proof checker which checks every application of derivation rules in a proof provided by the seller.
- Out sourcing computation. Where the buyer encodes the properties of a solution to some large problem that he wishes to have solved, e.g. solving some resource allocation problem.
- Cross-chain swaps / interaction. Where the seller exchanges a signed transaction on one blockchain for monetary compensation on another. This can potentially be generalized to interactions with complex smart contracts across chains by including the smart contract code in the predicate.
- Trustless bug-bounties. Where the buyer compiles a vulnerable program e.g. a file-parser, into a contingent payment predicate and wishes to buy a file such that the seller clearly illustrates control of the program counter (e.g. by setting it to a magical value).
- File transfer. Where the buyer wishes to obtain a file with a particular hash. This is the particular use case considered in the FairSwap paper.

In general any kind of digital good that can be defined by a computable program. The goal of thesis has been to devise a scheme that would be practical for all the aforementioned applications and many more.

1.3 Cryptographic Assumptions

We will make wide application of cryptographic primitives and ultimately reduce security of the FastSwap protocol to cryptographic assumptions. As an example of both cryptographic primitives and cryptographic assumptions, consider a simple and widely applied primitive: authenticated encryption (Definition 1).

Definition 1 (Authenticated Encryption). *An authenticated encryption scheme consists of a family indexed by a security parameter κ of pairs of two PPT (polynomial in κ) algorithms:*

- *Seal* : $\mathcal{K}_\kappa \times \mathcal{M} \rightarrow \mathcal{C}$, *encrypts a plaintext yielding a ciphertext.*
- *Open* : $\mathcal{K}_\kappa \times \mathcal{C} \rightarrow \mathcal{M} \cup \{\perp\}$, *decrypts and authenticates the ciphertext, yielding either \perp or a decrypted message.*

Satisfying perfect completeness for every message.

$$\forall m \in \mathcal{M} : 1 = \mathbb{P}_{k \leftarrow \mathcal{K}_\kappa} [\text{Open}(k, \text{Seal}(k, m)) = m] \quad (1)$$

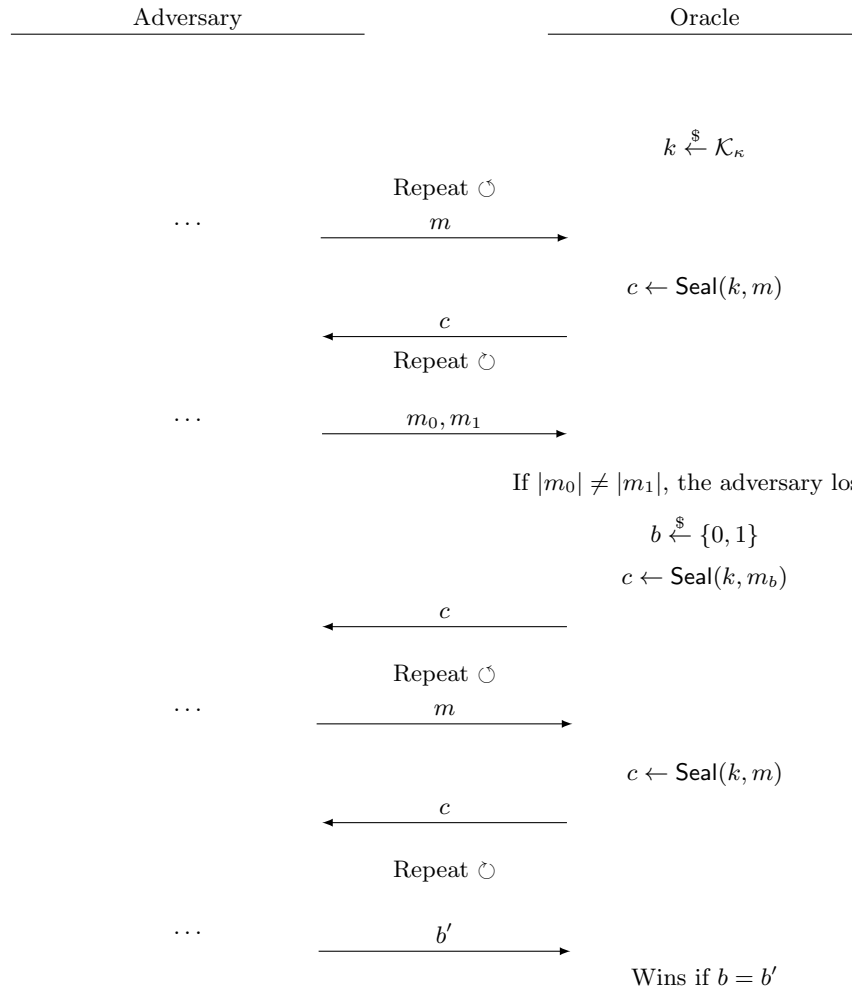
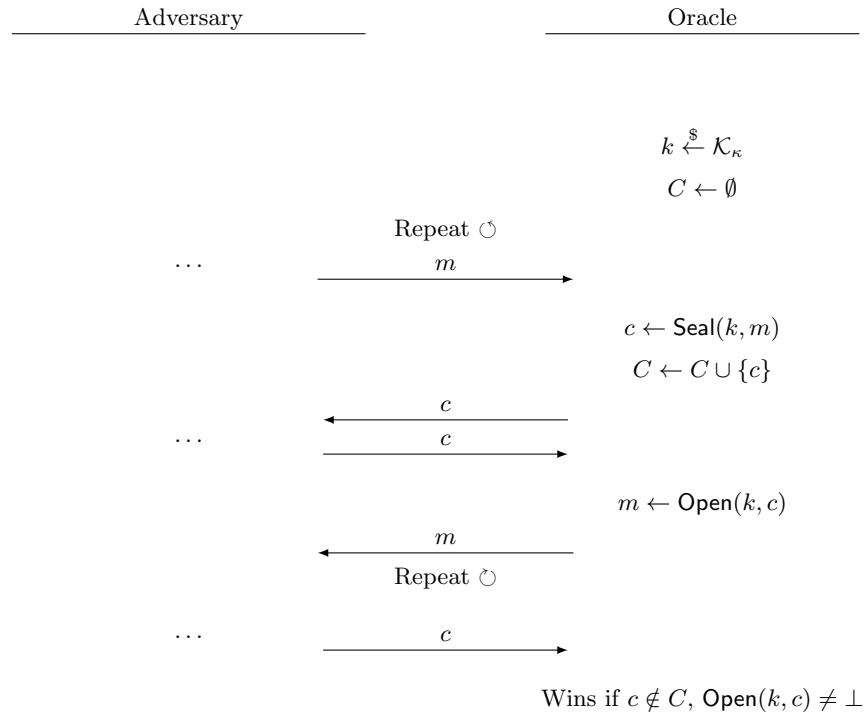


Fig. 1. IND-CPA Game

**Fig. 2.** SUF-CMA Game

The previous definition (Definition 1) is not concerned with security, in particular it does not rule out ‘clearly insecure’, but perfectly complete instantiations like $\text{Seal}(k, m) = m$. In order to define what ‘secure’ means for primitives, we formulate security games (see Figure 1, Figure 2) played by an ‘adversary’ against an ‘oracle’, ‘security’ then means that the winning probability for any polynomial adversary is bounded by some expression: usually reflecting that the adversary only does negligibly better than the naive strategy. For IND-CPA¹ security, we require that ciphertexts are indistinguishable to the adversary, even when he is allowed to encrypt arbitrary message:

Definition 2 (IND-CPA Security). *An authenticated encryption scheme is said to be IND-CPA if for every PPT adversary \mathcal{A} playing the IND-CPA game (Figure 1), there exists a negligible function $\text{negl}(\kappa)$ such that the probability of \mathcal{A} winning the IND-CPA game is bounded by $1/2 + \text{negl}(\kappa)$.*

Note in particular that it implies that encryption cannot be deterministic, observe also that in Figure 1 the game can only be won if $|m_0| = |m_1|$, reflecting that IND-CPA encryption does not hide the length of messages. For SUF-CMA² security we intuitively require that the adversary can only create valid ciphertexts by querying the oracle:

Definition 3 (SUF-CMA Security). *An authenticated encryption scheme is said to be SUF-CMA if for every PPT adversary \mathcal{A} playing the SUF-CMA game (Figure 2), there exists a negligible function $\text{negl}(\kappa)$ such that the probability of \mathcal{A} winning the SUF-CMA game is bounded by $\text{negl}(\kappa)$.*

A cryptographic assumption is a conjecture that there exists a primitive satisfying the security definition:

Conjecture 1. There exists an IND-CPA and SUF-CMA secure authenticated encryption scheme.

This is clearly an open conjecture³, since in particular it would imply $\mathbf{P} \neq \mathbf{NP}$: For both games the oracle \mathcal{O} is **PPT** if $\mathbf{P} = \mathbf{NP}$ then \mathcal{A} can non-deterministically guess the random tape R of \mathcal{O} , then simulate \mathcal{O}_R on his queries and compare the responses of the simulated oracle and the real transcript to distinguish between correct guesses. Concretely for the IND-CPA game \mathcal{A} can simply guess the key in non-deterministic polynomial time and accept iff. $\text{Open}(k, c) = m_0$, thus winning the IND-CPA game except with negligible probability.

1.4 Universal Composability & The Real/Ideal Paradigm

What does it mean for protocols to be ‘secure’? Merely claiming that a protocol is ‘secure’ is a meaningless statement, for two prime reasons: what is the exact

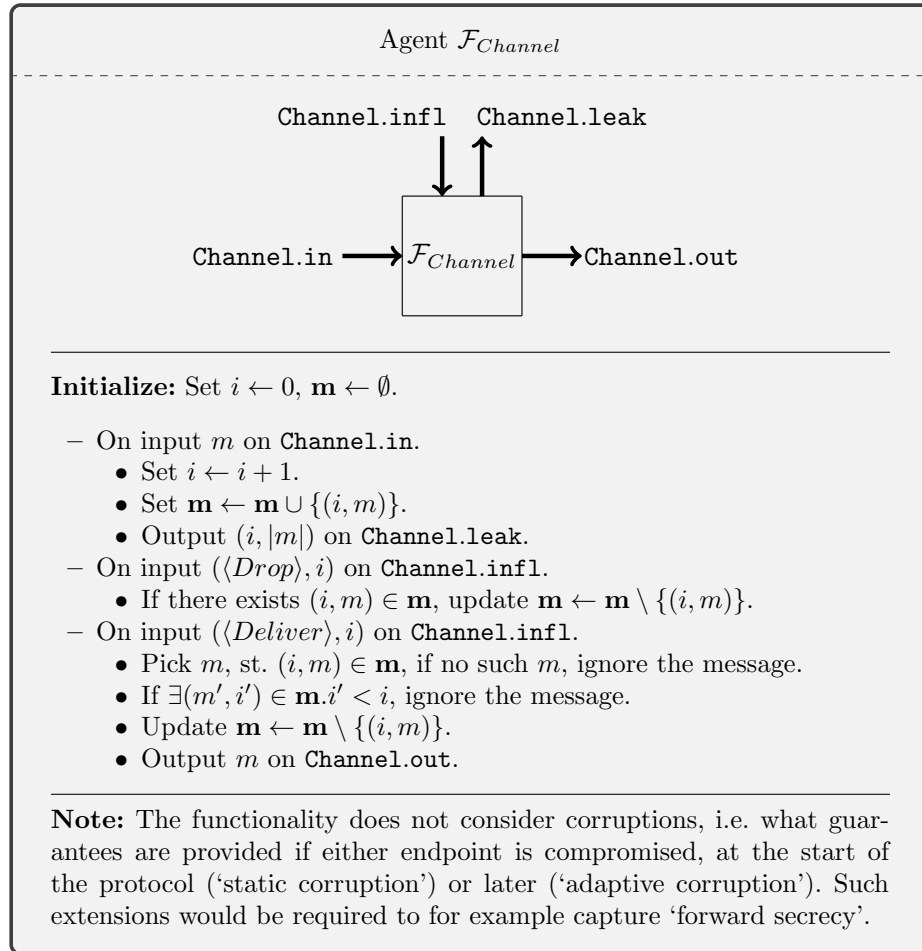
¹ Indistinguishability under Chosen Plaintexts

² Strong Unforgability under Chosen Message Attack

³ Like all cryptographic assumptions.

claimed function of the protocol? Under what conditions is the functionality ensured⁴?

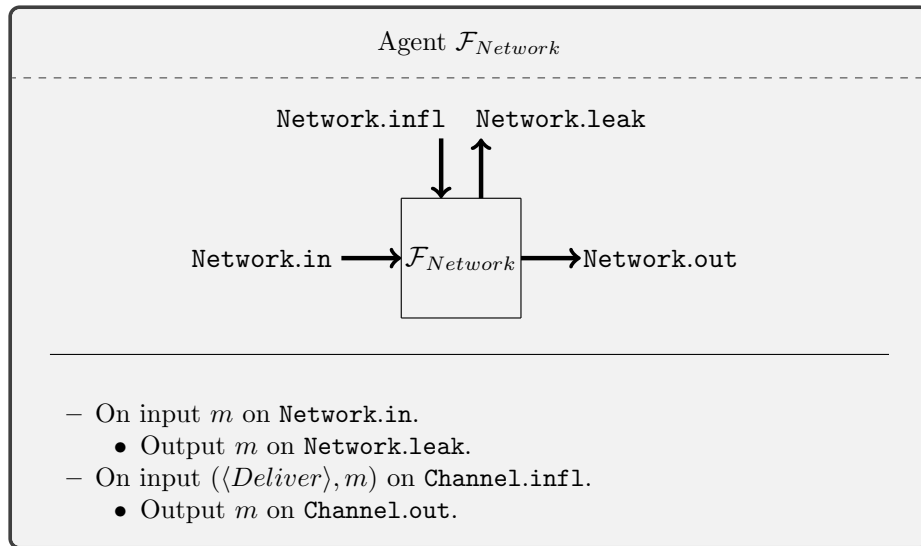
The first step in formalizing security consists in formulating exact answers to these two questions. In order to formalize exactly what the function of our protocol is we define so-called ‘ideal functionalities’, an ideal functionality in some ways corresponds to a trusted third-party which all parties can interact directly with. The desired behavior of the protocol is defined by the ideal functionality: the ideal functionality is the definition of ‘secure’ and in particular does not depend on cryptographic assumptions. This is best illustrated with an example: consider a formulation of a simple ideal functionality for an in-order one-directional encrypted authenticated channel:



⁴ e.g. what kind of adversaries is it secure against?

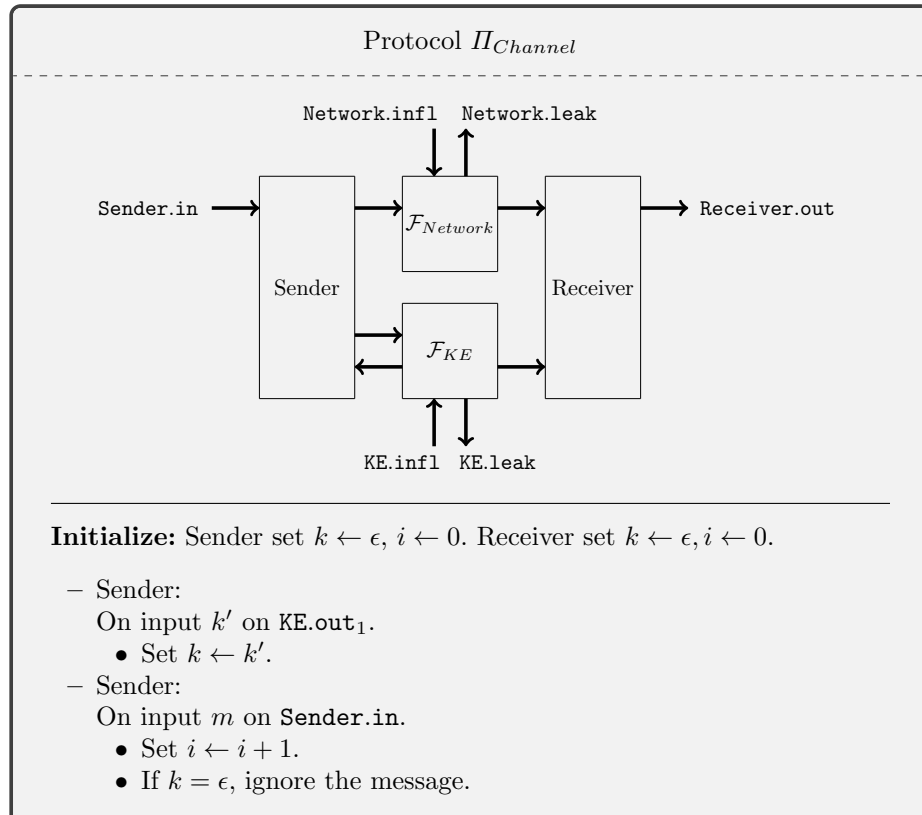
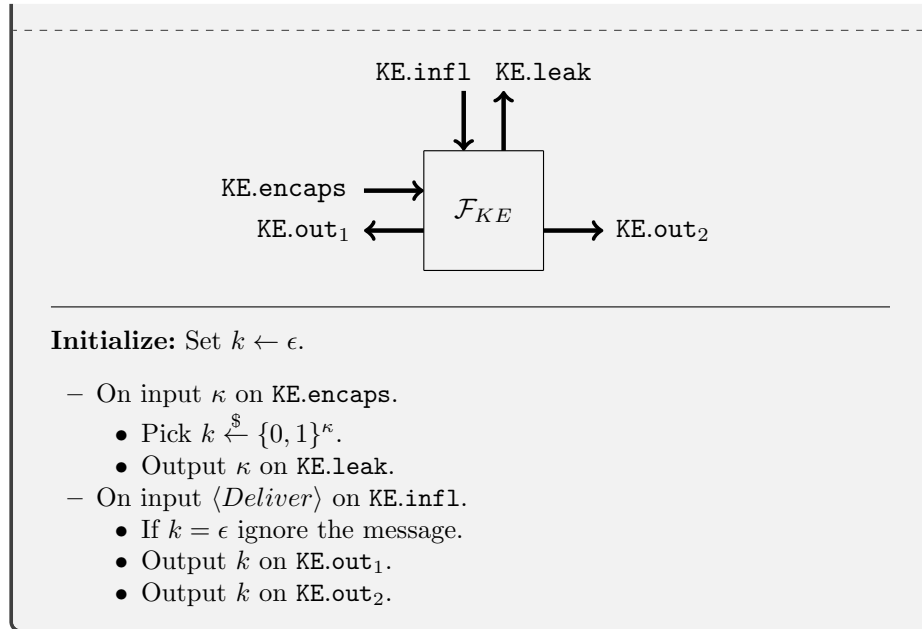
As the names indicate `Channel.infl` and `Channel.leak` are dedicated to the ‘adversary’. They reflect what the adversary can ‘influence’ and ‘learn’ respectively. Additionally the adversary can drop and deliver messages at his discretion using the influence port, this behavior corresponds roughly to the scenario where the authenticated channel allows for packets to be dropped (e.g. IPsec) and the adversary controls the network. Note that the adversary cannot inject messages, replay messages or deliver them out-of-order.

We also require ideal functionalities to model the underlying network, as well as some means of key-establishment. The $\mathcal{F}_{Network}$ functionality reflects the adversaries total control of the network:



The \mathcal{F}_{KE} allows the establishment of a shared secret with a length chosen by one of the parties, the ‘adversary’ learns only the length of the shared secret. In reality \mathcal{F}_{KE} is quite hard to instantiate in practice using e.g. $\mathcal{F}_{Network}$, due to Man-in-the-middle type attacks, but for illustration purposes we will simply assume that this is given.





- * Send κ on `KE.encaps`.
- If $k \neq \epsilon$:
 - * Compute $c \leftarrow \text{Seal}(k, (i, m))$.
 - * Output c on `Network.in`.
- Receiver:
 - On input k' on `KE.out2`:
 - Set $k \leftarrow k'$.
 - Receiver:
 - On input x on `Network.out`:
 - $p \leftarrow \text{Open}(k, x)$
 - If $p = \perp$ ignore the message.
 - Parse p as (i', p) , otherwise ignore the message.
 - If $i' \leq i$ ignore the message.
 - Set $i \leftarrow i'$.
 - Output m on `Receiver.out`.

Note: for simplicity reasons Π_{Channel} drops the messages if no key has been established. However the functionality also allows dropping of messages at will and we shall see that Π_{Channel} still instantiates $\mathcal{F}_{\text{Channel}}$.

Security in the universal composability framework is defined by indistinguishably between the ideal world (the ideal functionality) and the real world (the protocol) to an ‘environment’ which is given access to all the external ports i.e. the ports not used by the protocol. However, the difference between the ideal and the real world is allowed to be ‘trivial’, in particular the Π_{Channel} protocol and the $\mathcal{F}_{\text{Channel}}$ ideal functionality does not have the same set of ports, which would allow trivial distinguishing. To make precise what ‘trivial’ means, we introduce a simulator: an algorithm no more powerful than the environment which is given access to the ports of the functionality and can then emulate the exposed ports of the protocol. The justification is intuitively that any piece of information that the adversary (‘environment’) learns during interaction in the protocol, he could also obtain by interacting with the simulator and hence directly from the ideal functionality with no additional ‘powers’.

There are two ‘parameters’ in the definition above which can be tweaked: the ‘power’ of the environment and the notion of indistinguishability. In this article we shall concern ourselves only with PPT environments and negligible statistical distance (computational indistinguishability): the environment is a probabilistic polynomial time Turing machine and allowed to succeed with negligible advantage, which essentially allows the use of cryptographic assumptions in the protocol. One could consider stronger classes of adversaries e.g. computationally unbounded ones, or stricter notions of indistinguishability e.g. equality between the two distributions.

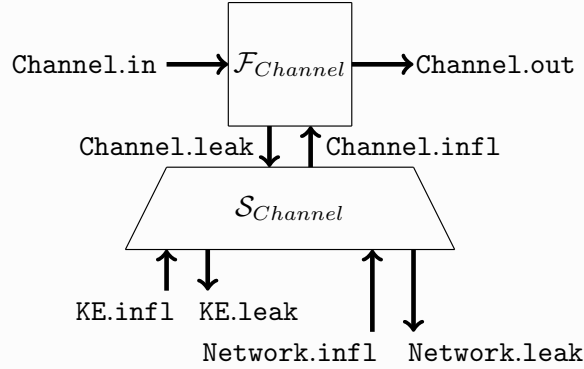
Intuitively, we can think of the $\mathcal{F}_{Network}$ and \mathcal{F}_{KE} functionalities as the premises for providing $\mathcal{F}_{Channel}$ by construction using our protocol and a corresponding simulator. Note for instance: the ‘weaker’⁵ $\mathcal{F}_{Network}$ and \mathcal{F}_{KE} are, and the ‘strong’⁶ the $\mathcal{F}_{Channel}$ functionality is, the harder it is to provide an instantiation of $\mathcal{F}_{Channel}$ from $\mathcal{F}_{Network}, \mathcal{F}_{KE}$ against a given set of adversaries.

We denote composition of agents by \diamond , e.g. $\Pi_{Channel} \diamond \mathcal{F}_{Network} \diamond \mathcal{F}_{KE}$ is the $\Pi_{Channel}$ protocol where the ports of $\Pi_{Channel}$ has been matched to the corresponding ports of the functionalities. We write $\Pi_{Channel} \diamond \mathcal{F}_{Network} \diamond \mathcal{F}_{KE} \geq_{comp} \mathcal{F}_{Channel}$ to indicate that $\Pi_{Channel} \diamond \mathcal{F}_{Network} \diamond \mathcal{F}_{KE}$ is ‘at least as secure’ as $\mathcal{F}_{Channel}$ with respect to computationally bounded environments. More formally, to show $\Pi_{Channel} \diamond \mathcal{F}_{Network} \diamond \mathcal{F}_{KE} \geq_{comp} \mathcal{F}_{Channel}$ our task is:

1. Define a simulator $\mathcal{S}_{Channel}$.
2. Show for any PPT environment \mathcal{Z} , the distributions $\mathcal{S}_{Channel} \diamond \mathcal{F}_{Channel} \diamond \mathcal{Z}$ and $\Pi_{Channel} \diamond \mathcal{F}_{Network} \diamond \mathcal{F}_{KE} \diamond \mathcal{Z}$ are statistically indistinguishable.

The latter is done by creating a reduction from the environment \mathcal{Z} which successfully distinguishes to an adversary \mathcal{A} breaking a cryptographic assumption; in a style superficially similar to Karp reductions known from complexity theory.

Simulator $\mathcal{S}_{Channel}$: simulate $\Pi_{Channel}$ using $\mathcal{F}_{Channel}$



Initialize: Let $C \leftarrow []$ (a map).

- On $\langle Deliver \rangle$ on $KE.infl$:
 - Set $k \xleftarrow{\$} \mathcal{K}_\kappa$.
- On $(i, |m|)$ on $Channel.leak$:
 - If $k = \epsilon$:
 - * Send $(\langle Drop \rangle, i)$ on $Channel.infl$.
 - Otherwise:
 - * Compute $c \leftarrow Seal(k, (i, 0^{|m|}))$.

⁵ Allowing more influence, leaking more state.

⁶ Allowing less influence, leaking less state.

- * Set $C[c] \leftarrow i$.
- * Output c on `Network.leak`.
- On $(\langle \text{Deliver} \rangle, x)$ on `Network.infl`:
 - If $x \notin C$, ignore the message.
 - Let $i \leftarrow C[x]$.
 - For $j \in [1, i - 1]$, output $(\langle \text{Drop} \rangle, j)$ on `Channel.infl`.
 - Output $(\langle \text{Deliver} \rangle, i)$ on `Channel.infl`.

Theorem 1 ($\Pi_{\text{Channel}} \diamond \mathcal{F}_{\text{Network}} \diamond \mathcal{F}_{\text{KE}} \geq_{\text{comp}} \mathcal{F}_{\text{Channel}}$).

Proof. Using the simulator $\mathcal{S}_{\text{Channel}}$, show:

$$\mathcal{S}_{\text{Channel}} \diamond \mathcal{F}_{\text{Channel}} \diamond \mathcal{Z} \equiv_{\text{stat}} \Pi_{\text{Channel}} \diamond \mathcal{F}_{\text{Network}} \diamond \mathcal{F}_{\text{KE}} \diamond \mathcal{Z} \quad (2)$$

We let $\mathcal{H}_0 = \mathcal{S}_{\text{Channel}} \diamond \mathcal{F}_{\text{Channel}} \diamond \mathcal{Z}$ be the distribution of the original simulation and consider indistinguishable hybrids \mathcal{H}_ℓ for all $\ell > 0$ defined as follows:

Consider the sequence of hybrids $\ell > 0$, where \mathcal{H}_ℓ is derived from \mathcal{H}_0 st. the simulator ‘extracts’ m_1, \dots, m_ℓ from $\mathcal{F}_{\text{Channel}}$ and leaks the encryptions of m_i on `Network.leak`, rather than letting $m_i = 0^{|m_i|}$ for $i \in [1, \ell]$ as done in $\mathcal{S}_{\text{Channel}}$. Note that $\mathcal{H}_{q(\kappa)}$ for some polynomial q is the distribution where every query made by \mathcal{Z} is answered with $\text{Seal}(k, m_i)$, i.e. $\forall \delta > 0 : \mathcal{H}_{q(\kappa)} = \mathcal{H}_{q(\kappa)+\delta}$. Suppose there exists an adversary that succeeds in distinguishing between \mathcal{H}_ℓ and $\mathcal{H}_{\ell+1}$, then there exists a PPT adversary \mathcal{A} winning the IND-CPA with non-negligible probability. Construct \mathcal{A} as follows: run $\mathcal{S}_{\text{Channel}} \diamond \mathcal{F}_{\text{Channel}} \diamond \mathcal{Z}$, whenever m_i is sent:

- If $i \leq \ell$ send m_i to the encryption oracle.
- If $i = \ell + 1$ send $(0^{|m_i|}, m_i)$ as the challenge to the oracle.
- If $i > \ell + 1$ send $0^{|m_i|}$ to the encryption oracle.

Clearly if $b = 0$, the distribution is \mathcal{H}_ℓ otherwise $\mathcal{H}_{\ell+1}$. Hence \mathcal{A} wins the IND-CPA game with non-negligible advantage.

It follows that \mathcal{H}_∞ derived from $\mathcal{S}_{\text{Channel}}$ where $c_i \leftarrow \text{Seal}(k, m_i)$ always, is indistinguishable from \mathcal{H}_0 . Lastly it remains to argue that \mathcal{H}_∞ is indistinguishable from $\Pi_{\text{Channel}} \diamond \mathcal{F}_{\text{Network}} \diamond \mathcal{F}_{\text{KE}} \diamond \mathcal{Z}$. Clearly the distributions of the ciphertexts c is equal (k is sampled with the same distribution), however the simulator replaces `Open` with the map C (crucial for the reduction to IND-CPA in the previous sequence of hybrids):

Suppose $\mathcal{H}_\infty \not\equiv_{\text{stat}} \Pi_{\text{Channel}} \diamond \mathcal{F}_{\text{Network}} \diamond \mathcal{F}_{\text{KE}} \diamond \mathcal{Z}$, then there exists \mathcal{A} winning the SUF-CMA game with non-negligible probability $p(\kappa)$. Claim: if $\mathcal{H}_\infty \not\equiv_{\text{stat}} \Pi_{\text{Channel}} \diamond \mathcal{F}_{\text{Network}} \diamond \mathcal{F}_{\text{KE}} \diamond \mathcal{Z}$, then \mathcal{Z} sends $(\langle \text{Deliver} \rangle, x)$ st. $\text{Open}(k, x) \neq \perp$ and $x \notin C$ with probability $p(\kappa)$. Produce \mathcal{A} as follows: let $q(\kappa)$ be a polynomial bound on the number of $(\langle \text{Deliver} \rangle, x)$ queries, pick $j \stackrel{\$}{\leftarrow} [1, q(\kappa)]$, let $i \leftarrow 0$. Run $\Pi_{\text{Channel}} \diamond \mathcal{F}_{\text{Network}} \diamond \mathcal{F}_{\text{KE}} \diamond \mathcal{Z}$, whenever `Seal` must be evaluated, call the tagging oracle in the SUF-CMA game, when \mathcal{Z} sends $(\langle \text{Deliver} \rangle, x)$ on `Network.infl`:

- If $i = j$, send x as the forgery to the oracle in the *SUF-CMA* game.
- If $i \neq j$ send x to the check oracle in the *SUF-CMA* game and obtain p .
- Set $i \leftarrow i + 1$.

With probability $1/q(\kappa)$ the guess for the message index is correct, hence with non-negligible probability $p(\kappa)/q(\kappa)$ \mathcal{A} wins the *SUF-CMA* game.

It is also possible to prove security solely in the game-based paradigm, however the universal composability framework provides one central advantage, the UC theorem:

Theorem 2 (Universal Composability). *Let Π_F, Π_G be protocols implementing $\mathcal{F}_F, \mathcal{F}_G$ respectively. If $\Pi_F \diamond \mathcal{F}_G \geq_{comp} \mathcal{F}_F$ and $\Pi_G \diamond \mathcal{F}_H \geq_{comp} \mathcal{F}_G$ then $(\Pi_F \diamond \Pi_G) \diamond \mathcal{F}_H \geq_{comp} \mathcal{F}_F$.*

Intuitively the UC theorem states that if a protocol (e.g. Π_F) uses an ideal functionality (e.g. \mathcal{F}_G), this functionality can be safely swapped for any sub-protocol implementing the functionality (e.g. Π_G). Besides providing modularity, this enables us to eliminate potential issues where the interplay between protocols leads to a system that is ‘insecure’ overall.

1.5 Blockchains

For this thesis we will concern ourselves less with the exact details of the underlying blockchain and merely use it as an instantiation of an ideal functionality. Essentially blockchains offer the promise of an ‘honest-but-curious’ party for hire: a smart contract can be created by any other participant in the protocol, its code inspected and a guarantee is provided that this simulated party will act according to the protocol specification.

Most current blockchains validate smart contract execution in time that is linear in the length of the computation triggered by the transactions. This is due to the naive method of ensuring the integrity of the state transition: every transaction is simply re-executed by every node in the network. This places a natural constraint on the amount of computation that can be ensured by the blockchain: no more than the slowest node in the network can execute. Succinct proofs of computational integrity with sub-linear verification time has only seen limited application.

Besides introducing a natural upper limit this also leads to a high price of executing code ‘on-chain’, since the resource is finite and users essentially engage in an auction for running time. This motivates our desire to reduce the amount of computation on-chain, both to enable scaling and to reduce costs. Additionally, smart contracts are ‘honest-but-curious’ which is undesirable when dealing with private data.

2 Contributions

I claim to have designed the first truly practical contingent payment protocol, FastSwap, described in the accompanying article. In particular the distinguishing features of FastSwap are:

- Concrete efficiency. FastSwap does not depend on zero-knowledge proofs which often entails significant computational overhead for the prover.
- Honest communication is optimal (the size of the program description and the witness). In particular it is independent of the length of the computation.
- Dispute communication is logarithmic in the length of the computation. The dispute resolution contract is independent of the predicate and need only be deployed once globally.
- The computational model is flexible, in particular branching RAM programs can be efficiently executed, which enables easy complication of existing languages to FastSwap predicates.
- Concrete optimizations can enable the executions of complex high-level languages, even when the language during dispute is very simple (leading to a small dispute resolution contract). This is likely to yield significant performance increases in practice for long-running predicates.

FastSwap

Mathias Hall-Andersen¹[0000-0002-0195-6659]

Aarhus University, mathias@hall-andersen.dk

Abstract. FastSwap enables a simple and concretely efficient contingent payments for complex predicates. FastSwap only relies on symmetric primitives (semantically secure encryption and cryptographic hash functions) and avoids ‘heavy-weight’ primitives such as general ZKP systems. FastSwap is particularly well-suited for applications where the witness or predicate is large (on the order of MBs / GBs) or expensive to calculate (e.g. 2^{30} computation steps or memory). Additionally FastSwap allows predicates to be implemented using virtually any computational model (including branching execution), which e.g. enables practitioners to efficiently express the predicate in imperative languages already familiar to them, without an expensive transformation to e.g. satisfiability of arithmetic circuits. The cost of this efficiency during honest execution is a logarithmic number of rounds during a dispute resolution in the presence of a corrupted party. Let the witness be of size $|w|$ and the predicate of size $|P|$, where computing $P(w)$ takes n steps. In the honest case the off-chain communication complexity is $|w| + |P| + c$ for a small constant c , the on-chain communication complexity is c' for a small constant c' . In the malicious case the on-chain communication complexity is $O(\log n)$ with small constants. Concretely with suitable optimizations the number of rounds (on-chain transactions) for a computation of 2^{30} steps can be brought to 2 in the honest case with an estimated cost of ≈ 2 USD on the Ethereum blockchain¹ and to 14 rounds with an estimated cost of ≈ 4 USD in case of a dispute. It is noted that the corrupted party can be made a penalty in case of dispute.

Keywords: Contingent payments, Fair exchange, Concrete efficiency, Smart contracts, Provable security, Universal composability, Authenticated computation structures, Authenticated data structures.

1 Introduction

FastSwap chiefly enables ‘Contingent Payments’, wherein a buyer holds a predicate P and funds, whereas a seller holds a purported witness w for which she claims $P(w) = 1$. The goal of a contingent payment protocol is to enable the buyer to learn w and ensure payment iff. w satisfies $P(w) = 1$. Examples of potential applications of contingent payments include:

¹ At the time of writing, using a gas price of 10 Gwei (1 ETH = 10^9 Gwei) and with price of Ethereum at 160 USD/ETH. Assuming a one-time library contract has already been published.

- **Buying signatures on documents.** For example paying for a certificate iff. the signature on the certificate is correct and particular fields are set.
- **Buying transactions.** Where the buyer wishes to purchase a signed transaction for some blockchain, which causes a particular smart contract (included in the predicate) to reach a given state.
- **Purchasing solutions to hard problems.** e.g. paying for the factorization of composite numbers or solutions to resource allocations problems.
- **Purchasing a file with a particular hash.** In which the buyer wishes to pay for a large file transfer iff. he receives the full file and its cryptographic hash matches a particular value.
- **Trustless bug-bounties.** Where the input (e.g. a file) must violate some safety constrains (e.g. demonstrate arbitrary control of the control flow) of a potentially vulnerable program (e.g. an image parser) to be considered a valid witness.
- **Buying formal proofs.** The predicate can contain a formal proof checker and the witness is required to be a formal proof which proves a desired statement.

We facility this by introduction a very limited ‘honest-but-curious’ judge, which possesses only *poly*(log) computation & communication in $|w|, |P|$ and n (the length of the computation $P(w)$). Furthermore we want to ensure that when both parties are honest w is not leaked to the environment. These constraints each preclude the naive protocol wherein the buyer sends the funds to the judge, the seller sends w to the judge and the judge pays the seller iff. $P(w) = 1$. The motivation for this setting is enabling the instantiation of the judge efficiently via smart contract.

Like prior works [1] [4], FastSwap enables the judge to learn the whether the exchange was successful, therefore the action of the judge upon learning the output of the predicate can be made essentially arbitrary, additionally FastSwap can easily be extended to multiple ‘buyers’. Due to these features we believe it might be useful to consider the more general application of enabling ‘off-chain’ computation of predicate by a (potentially malicious) ‘prover’ under the scrutiny of one/more (potentially malicious) ‘auditors’, in the presence of a judge incapable of computing the predicate alone. We note that constructions like FastSwap has the ability to be significantly more efficient in practice than the application of current Zero-Knowledge proof systems to elevate the problem. Hence we frame FastSwap in this more general setting.

1.1 Prior Work

Zero-Knowledge Contingent Payments (ZKCP). The Zero-Knowledge Contingent Payment (ZKCP) construction [4] (by Gregory Maxwell) requires a zero-knowledge proof system capable over the relation induced by the predicate, a semantically secure encryption scheme (**Enc**) and a collision resistant hash function (**CRH**). The original formulation is in terms of a seller (acting as the prover), selling a witness for a predicate P to a buyer (acting as the auditor) in exchange for financial compensation. The scheme operates as follows: for a public o, C (chosen by the seller), the seller proves to the buyer in zero-knowledge that he knows w, k st.

$$o = \text{CRH}(k), P(w) = 1, C = \text{Enc}(k, w) \quad (1)$$

The seller then sends o, C and the proof π to the buyer, who aborts the protocol in case π is invalid. Otherwise the buyer posts a transaction (acting as the judge) to the blockchain, which can only be spend by revealing a preimage of o . The seller claims the funds of the transaction using k , whereby the buyer learns k and is able to decrypt C to obtain the witness. Variations of this scheme has been considered[2][10] in applications where supplying π itself leaks information about the witness, e.g. whenever π itself constitutes a ‘witness’².

FairSwap. The FairSwap[1] protocol (by Stefan Dziembowski, Lisa Eckey, Sebastian Faust) avoids the need for a Zero-Knowledge proof system at the cost of transmitting the entire encrypted computation trace. Additionally FairSwap requires that the predicate be computed using a straight-line program. The execution model is a computational circuit: an acyclic graph wherein every vertex/gate applies an operation to its children/inputs. The scheme operates by having the prover evaluate and encrypt the full execution trace (initial inputs and outputs of every gate), then the prover computes a Merkle commitment to the encrypted execution trace and sends this to the judge. The encrypted execution trace is transferred to the auditor, who recomputes the Merkle tree and verifies that it is consistent with the one held by the judge. Then the decryption key is sent by the prover to the judge and the auditor decrypts the execution trace. If any gate is applied incorrectly (or the output of the computation is not accepting), the auditor can prove Merkle paths to the inputs of the erroneously applied gate and convince the judge that the prover is malicious. The FairSwap protocol (as formulated) assumes that the full predicate description is available to the judge, which makes it best suited for applications where the predicate is has a small description but potentially a long running time: the example in the paper being the computation of a Merkle hash which allows the purchasing of files, where the linear communication complexity of FairSwap in the length of the trace is optimal. FastSwap is inspired by the FairSwap protocol.

² An example being Proofs-of-Storage, where a Proof-of-Knowledge for a Proof-of-Storage on a given challenge is itself a Proof-of-Storage.

Comparison. Let $|w|$ be the size of the witness, let $|P|$ be the size of the predicate description, let n be the length of the computation of $P(w)$. We compare the complexity and computational models of the two prior works to FastSwap in Figure 1 and Figure 2. Note that ZKCP does not have a ‘dispute resolution’ phase. FastSwap reduces the communication complexity compared to FairSwap, additionally FastSwap provides more freedom in the choice of computational model, in particular allowing efficient execution of branching RAM machines, which enables relatively easy and efficient compilation of existing imperative smart contract languages.

Name	Computational Model	Comm. (off-chain)	Comm. (on-chain)
ZKCP (zk-SNARK)	Arithmetic Circuit	$\Theta(P + w)$	$\Theta(1)$, 2 rounds
FairSwap	Computational Circuit	$\Theta(P + w + n)$	$\Theta(P)$, 2 rounds
FastSwap	RAM Machines	$\Theta(P + w)$	$O(1)$, 2 rounds

Fig. 1. Complexity of honest execution.

Name	Communication (on-chain)	Rounds (on-chain)
ZKCP (zk-SNARK)	–	–
FairSwap	$\Theta(\log n)$	$\Theta(1)$
FastSwap	$\Theta(\log n)$	$\Theta(\log n)$

Fig. 2. Complexity of dispute resolution.

Recall that generic transformation of a program in the RAM model running in n steps, requires a circuit of size $n^3 \log n$. Hence efficient compilation of e.g. existing smart contract language to FairSwap (or e.g. ZKCP with zk-SNARKs) predicates is unlikely, while this is one of the envisioned applications of FastSwap. One heuristic argument (without rigors game theoretic backing), as to why we believe the dispute resolution complexity is less crucial than the honest execution for many real-world applications is that both systems, FairSwap and FastSwap, allows the judge to discern which party is malicious. Hence a penalty can be enforced by having both parties deposit collateral with the judge prior to the swap, which can be seized / send to the honest party in case of malicious behavior.

1.2 Features of FastSwap

Simple & efficient primitives. The FastSwap protocol does not rely on ‘heavy weight’ primitives like zero-knowledge proof systems, a central goal of FastSwap is to provide concrete efficiency for a wide class of very large predicates.

Constant communication in the honest case. The communication complexity during honest execution is the size of the program, the size of the witness and a small constant. The communication complexity is independent of the length of the execution for the predicate.

Logarithmic communication for dispute resolution. In case of a malicious prover or auditor, dispute resolution for an execution trace of n steps is completed within $O(\log n)$ rounds and $O(\log n)$ communication with small constants.

Flexible execution model. Previous work require that the predicate is implemented via straight-line program, FastSwap additionally supports efficient branching execution and RAM machines. One possible application is to enable efficient compilation of existing smart contract languages to predicates for contingent payments.

Efficient for large program descriptions. The program description of the predicate need only be available to the prover and auditor, this allows executing program with large descriptions. This also allows deployment of a generic ‘interpreter & dispute resolution’ judge contract, which can be reused for selling different witnesses to different predicates by different parties.

2 Notation

Symbols enclosed in angle brackets $\langle \cdot \rangle$ represents unique symbols (‘atoms’), e.g. $\langle Identifier \rangle$ is simply a symbol recognized by all participants in the protocol. The length of a bit string s is denoted by $|s|$. Throughout the article κ will denote a security parameter.

3 Primitives

3.1 Symmetric Encryption

Definition 1 (Symmetric Key Encryption). *A symmetric encryption schemes is a family two algorithms running on 1^κ (omitted for brevity):*

- A PPT algorithm, which samples uniformly from the key space $k \xleftarrow{\$} \mathcal{K}_\kappa$
- A PPT algorithm ‘encryption’ $Enc : \mathcal{K}_\kappa \times \mathcal{M} \rightarrow \mathcal{C}_\kappa$
- A PPT algorithm ‘decryption’ $Dec : \mathcal{K}_\kappa \times \mathcal{C}_\kappa \rightarrow \mathcal{M}$

Satisfying perfect completeness:

$$\forall m \in \mathcal{M} : 1 = \mathbb{P}[Dec(k, Enc(k, m)) = m : k \xleftarrow{\$} \mathcal{K}_\kappa]$$

Definition 2 (One-Time Semantic Security). A family of symmetric encryption schemes (Definition 1) is said to be one-time semantically secure if for all pairs of PPT algorithms $(\mathcal{A}_1, \mathcal{A}_2)$, there exists a negligible function negl st.

$$\begin{aligned} 1/2 + \text{negl}(\kappa) &\geq \mathbb{P}[b' = b \wedge |m_1| = |m_2| : (m_1, m_2) \leftarrow \mathcal{A}_1(1^\kappa), \\ &\quad b \xleftarrow{\$} \{0, 1\}, k \xleftarrow{\$} \mathcal{K}_\kappa, b' \leftarrow \mathcal{A}_2(1^k, \text{Enc}(k, m_b))] \end{aligned}$$

Note that unlike the ordinary IND-CPA definition, we do not require the encryption scheme to be indistinguishable across multiple encryption queries. In particular $\text{Enc} : \mathcal{K}_\kappa \times \mathcal{M} \rightarrow \mathcal{C}$ can be deterministic.

3.2 Collision Resistant Hashes

Definition 3 (Cryptographic Hash). A family of cryptographic hash functions consists of an efficient deterministic algorithm running on 1^κ :

- A polynomial time algorithm ‘hash’ $\text{CRH} : \{0, 1\}^* \rightarrow \mathcal{H}_\kappa$

Where $\forall h \in \mathcal{H}_\kappa : |h| = \kappa$

Definition 4 (Collision Resistance). A hash function family (Definition 3) is said to be collision resistant if for every PPT algorithm \mathcal{A} , there exists a negligible function negl st.

$$\text{negl}(\kappa) \geq \mathbb{P}[m \neq m' \wedge \text{CRH}(m) = \text{CRH}(m') : (m, m') \leftarrow \mathcal{A}(1^\kappa)]$$

3.3 Binding & Hiding Commitments

Definition 5 (Commitment). A commitment scheme is a family of two efficient algorithms running on 1^κ (omitted for brevity):

- A PPT algorithm ‘commit’ $\text{Comm} : \mathcal{R}_\kappa \times \mathcal{M} \rightarrow \mathcal{C}_\kappa$
- A PPT algorithm ‘open’ $\text{Open} : \mathcal{R}_\kappa \times \mathcal{M} \times \mathcal{C}_\kappa \rightarrow \{0, 1\}$

Satisfying perfect completeness:

$$\forall m \in \mathcal{M} : 1 = \mathbb{P}[\text{Open}(r, m, c) = 1 : r \xleftarrow{\$} \mathcal{R}_\kappa, c \leftarrow \text{Comm}(r, m)]$$

Definition 6 (Computationally Binding Commitment). A commitment scheme (Definition 5) is said to be computationally binding if for all PPT algorithm \mathcal{A} there exists a negligible function negl st.

$$\begin{aligned} \text{negl}(\kappa) &\geq \mathbb{P}[m_1 \neq m_2 \wedge \text{Open}(r_1, m_1, c) = 1 \wedge \text{Open}(r_2, m_2, c) = 1 : \\ &\quad (c, r_1, r_2, m_1, m_2) \leftarrow \mathcal{A}(1^\kappa)] \end{aligned}$$

Definition 7 (Computationally Hiding Commitment). A commitment scheme (Definition 5) is said to be computationally hiding if for all pairs of PPT algorithms $(\mathcal{A}_1, \mathcal{A}_2)$ there exists a negligible function negl st.

$$\begin{aligned} 1/2 + \text{negl}(\kappa) &\geq \mathbb{P}[b' = b : (m_1, m_2) \leftarrow \mathcal{A}_1(1^\kappa), \\ &b \xleftarrow{\$} \{0, 1\}, r \xleftarrow{\$} \mathcal{R}_\kappa, b' \leftarrow \mathcal{A}_2(1^k, \text{Comm}(r, m_b))] \end{aligned}$$

4 Authenticated Computation Structures

Definition 8 (Authenticated Data Structure). An authenticated data structure scheme consists of a set of possible states \mathcal{S}_κ , a set of tags \mathcal{T}_κ , a set of possible operations \mathcal{O} , a set of results \mathcal{R} , a set of descriptions of initial states \mathcal{I} and four deterministic polynomial time algorithms:

- *Initial* : $\mathcal{I} \rightarrow \mathcal{S}$. Construct an initial state from a description.
- *Tag* : $\mathcal{S} \rightarrow \mathcal{T}_\kappa$. Compute a succinct ‘tag’ of the state.
- *Apply* : $\mathcal{S} \times \mathcal{O} \rightarrow \mathcal{S} \times \mathcal{R} \times \mathcal{P}_\kappa$. Apply an operation to the state, optionally yielding a result. Produce a proof of correct application of the operation, which can be verified using only the tags of the original and resulting state.
- *Verify* : $\mathcal{T}_\kappa \times \mathcal{T}_\kappa \times \mathcal{O} \times \mathcal{R} \times \mathcal{P}_\kappa \rightarrow \{1, 0\}$. Verifies the execution of an operation on the state corresponding to the tag of the previous state and tag of the resulting state after application of the operation.

Satisfying perfect completeness:

$$\begin{aligned} \forall S \in \mathcal{S}, O \in \mathcal{O} : \text{Verify}(T, T', R, O, \pi) = 1 \text{ where} \\ (S', R, \pi) \leftarrow \text{Apply}(S, O) \\ T \leftarrow \text{Tag}(S), T' \leftarrow \text{Tag}(S') \end{aligned}$$

Computation is formulated in terms of ‘Authenticated Computation Structures’, which can be seen as an authenticated data structure scheme, wherein the operation is uniquely defined by the current state of the data structure and an immutable ‘environment’.

Definition 9 (Authenticated Computation Structure). An authenticated computation structure scheme consists of an input space \mathcal{I} containing descriptions of initial computation states, a space of possible computation structures \mathcal{S} , a space of possible ‘environments’ \mathcal{E} , a set of ‘tag’ values \mathcal{T}_κ , a set of proofs \mathcal{P}_κ and five deterministic polynomial time algorithms:

- *Initial* : $\mathcal{I} \rightarrow \mathcal{S}$. Construct an initial state from a description.
- *Tag* : $\mathcal{S} \rightarrow \mathcal{T}_\kappa$. Produce a succinct tag corresponding to the structure.
- *Step* : $\mathcal{E} \times \mathcal{S} \rightarrow \mathcal{S}$. Progresses the computation by ‘a single step’.

- *Prove* : $\mathcal{E} \times \mathcal{S} \rightarrow \mathcal{P}_\kappa$. Produce a succinct proof of correct execution.
- *Verify* : $\mathcal{E} \times \mathcal{T}_\kappa \times \mathcal{T}_\kappa \times \mathcal{P}_\kappa \rightarrow \{1, 0\}$. Verify the execution of a step.

Satisfying perfect completeness:

$$\begin{aligned} \forall e \in \mathcal{E}, S \in \mathcal{S} : \text{Verify}(e, T, T', \pi) = 1 \text{ where} \\ S' \leftarrow \text{Step}(e, S), \pi \leftarrow \text{Prove}(e, S), \\ T \leftarrow \text{Tag}(S), T' \leftarrow \text{Tag}(S') \end{aligned}$$

i.e. verification succeeds for every pair of successive computation structures.

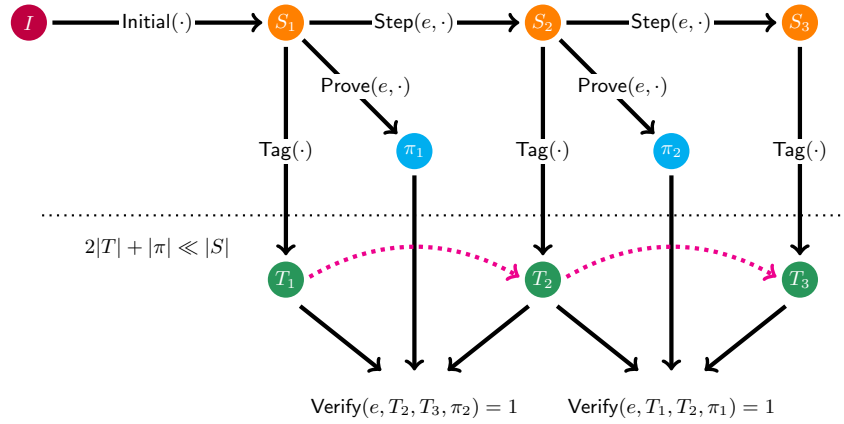


Fig. 3. Relation between the algorithms for auth. computation structures

The primitive is directly related to authenticated data structures (Defintion 8) and can be generically constructed from such schemes by defining a function $\text{Operation} : \mathcal{S} \rightarrow \mathcal{O} \times \mathcal{P}_\kappa$ which takes the state of the data structure and returns the next operation to apply and a proof, then deriving an implementation of the algorithms above in the obvious way. A concrete example of this pattern is provided in Section 7.

Definition 10 (Computational Integrity). An authenticated computation structure scheme is said to provide computational integrity, if $\text{Tag} : \mathcal{S} \rightarrow \mathcal{T}_\kappa$ is a collision resistant hash (Defintion 3) and for every PPT algorithm \mathcal{A} , there exists a negligible function negl such that:

$$\begin{aligned} \text{negl}(\kappa) \geq \mathbb{P}[T' \neq \text{Tag}(\text{Step}(e, S)) \wedge \text{Verify}(e, T, T', \pi) = 1 : \\ (e, S, T', \pi) \leftarrow \mathcal{A}(1^\kappa), T \leftarrow \text{Tag}(S)] \end{aligned}$$

i.e. it is intractable to compute a tag and a proof that passes verification, except when the tag corresponds to the successor state.

For later convenience we define some simple functions which are derived from any authenticated computational structure scheme:

Definition 11 (Terminate : $\mathcal{E} \times \mathcal{S} \rightarrow \mathbb{N}_+$). *Terminate repeatedly applies Step and returns the number of steps before an accepting or rejecting state is reached. Formally, with the patterns being matched by preference from top to bottom:*

$$\begin{aligned} \text{Terminate}(e, S) &:= 1 \text{ where } S \in \{\langle \text{Accept} \rangle, \langle \text{Reject} \rangle\} \\ \text{Terminate}(e, S) &:= 1 + \text{Terminate}(e, S') \text{ where } S' \leftarrow \text{Step}(e, S) \end{aligned}$$

Where $\langle \text{Accept} \rangle$ and $\langle \text{Reject} \rangle$ is uniquely recognized accepting and rejecting terminal states respectively.

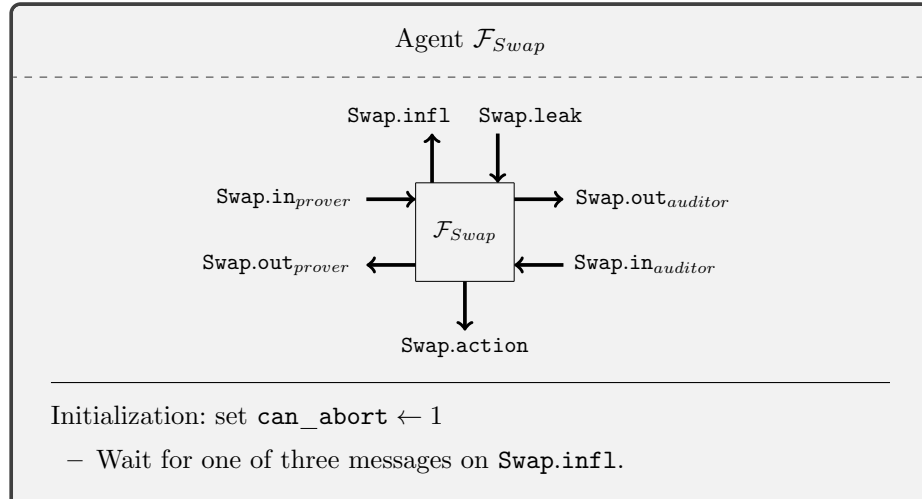
Definition 12 (StepN : $\mathcal{E} \times \mathcal{S} \times \mathbb{N}_+ \rightarrow \mathcal{S}$). *StepN applies Step a specified number of times and returns the resulting state. Formally, with the patterns being matched by preference from top to bottom:*

$$\begin{aligned} \text{StepN}(e, S, 1) &:= S \\ \text{StepN}(e, S, *) &:= S \text{ where } S \in \{\langle \text{Accept} \rangle, \langle \text{Reject} \rangle\} \\ \text{StepN}(e, S, n) &:= \text{StepN}(e, S', n-1) \text{ where } S' \leftarrow \text{Step}(e, S) \end{aligned}$$

Where $\langle \text{Accept} \rangle$ and $\langle \text{Reject} \rangle$ is uniquely recognized accepting and rejecting terminal states respectively. One can think of StepN as returning the n 'th step of the computation right-padded by the final accepting/rejecting state.

5 Ideal Functionalities

We formulate the behavior of FastSwap using the universal composability (UC) framework with the style and notation of Cramer, et al. [8]. The $\mathcal{F}_{\text{Swap}}$ functionality captures the desired behavior of a contingent exchange protocol:

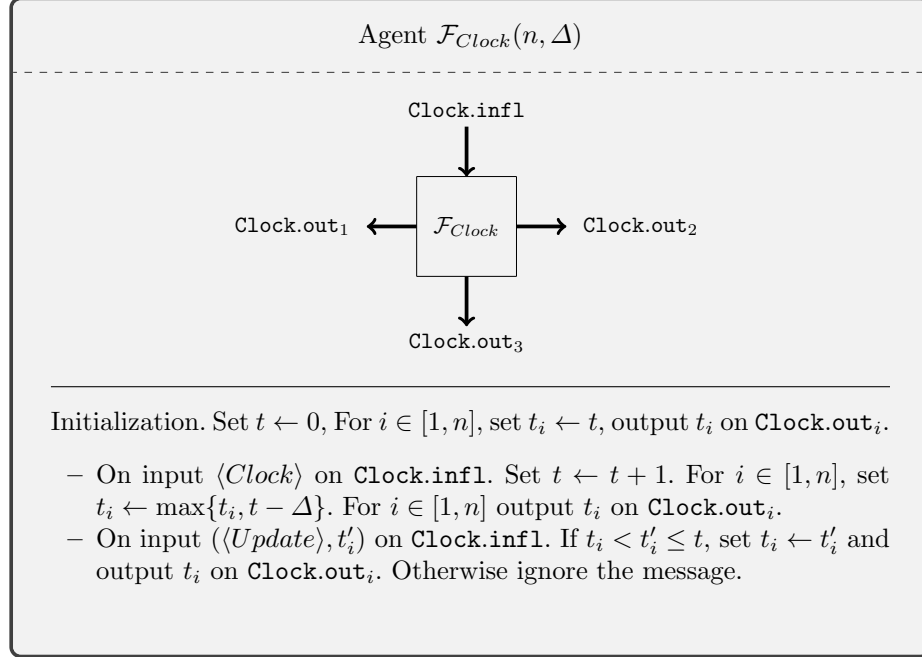


- ◊ $\langle \text{Auditor} \rangle$: Mark the prover as corrupted, by ignoring any message on $\text{Swap.in}_{\text{auditor}}$. Whenever a message of the form $(\langle \text{Send} \rangle, m)$ is received on Swap.infl act as if m was received on $\text{Swap.in}_{\text{auditor}}$. Whenever a message m is output on $\text{Swap.out}_{\text{auditor}}$, also output m on Swap.leak
 - ◊ $\langle \text{Prover} \rangle$: Mark the prover as corrupted, by ignoring any message on $\text{Swap.in}_{\text{prover}}$. Whenever a message of the form $(\langle \text{Send} \rangle, m)$ is received on Swap.infl act as if m was received on $\text{Swap.in}_{\text{prover}}$. Whenever a message m is output on $\text{Swap.out}_{\text{prover}}$, also output m on Swap.leak
 - ◊ $\langle \text{Honest} \rangle$. Indicating no corruption.
- Ignore any subsequent corruption messages.
- Any time, on input $\langle \text{Abort} \rangle$ on Swap.infl , $\text{Swap.in}_{\text{auditor}}$ or $\text{Swap.in}_{\text{prover}}$ and if $\text{can_abort} = 1$, then abort the protocol:
 - Output \perp on $\text{Swap.out}_{\text{prover}}$.
 - Output \perp on $\text{Swap.out}_{\text{auditor}}$.
 - Output \perp on Swap.leak .
 - Ignore any further messages on any in port.
 - On input P on $\text{Swap.in}_{\text{auditor}}$:
 - Store P .
 - Output P on $\text{Swap.out}_{\text{prover}}$.
 - Output $|P|$ on Swap.leak .
 - On input w on $\text{Swap.in}_{\text{prover}}$:
 - Store w .
 - Output $|w|$ on Swap.leak .
 - On input $\langle \text{Swap} \rangle$ on $\text{Swap.in}_{\text{prover}}$, when both P, w has been set:
 - Set $\text{can_abort} \leftarrow 0$.
 - Output w on $\text{Swap.out}_{\text{auditor}}$.
 - If either party is corrupted leak the entire state of the functionality on Swap.leak : every message sent and received by the functionality.
 - On input $\langle \text{Action} \rangle$ on Swap.infl , when $\text{can_abort} = 0$:
 - Interpret P as a description of a computable function.
 - Output $P(w)$ on Swap.action and Swap.leak .

The $\mathcal{F}_{\text{Swap}}$ functionality leaks its entire state after $\text{can_abort} = 0$ whenever a corrupted party is present. Intuitively we can accept to leak the witness to the world in case of corruption after the protocol cannot be aborted, since after $\text{can_abort} = 0$ the corrupted party will possess the witness and could publish this (outside the scope of the protocol) regardless. Hiding of the witness must only be ensured as long as $\text{can_abort} = 1$ or whenever both parties are honest.

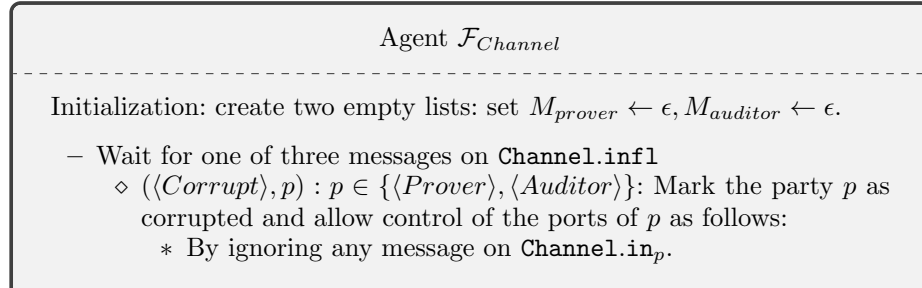
The separation of the $\langle Swap \rangle$ and $\langle Action \rangle$ messages, enables the implementation to run some ‘dispute’ protocol in case one of the parties is corrupted, before delivering the output on **Swap.action**. The leaked state after $\langle Swap \rangle$ can be used to simulate the leakage of this ‘dispute’ protocol.

The \mathcal{F}_{Clock} functionality models n monotonically increasing clocks, where the drift between any pair of clocks is bounded by a constant Δ :



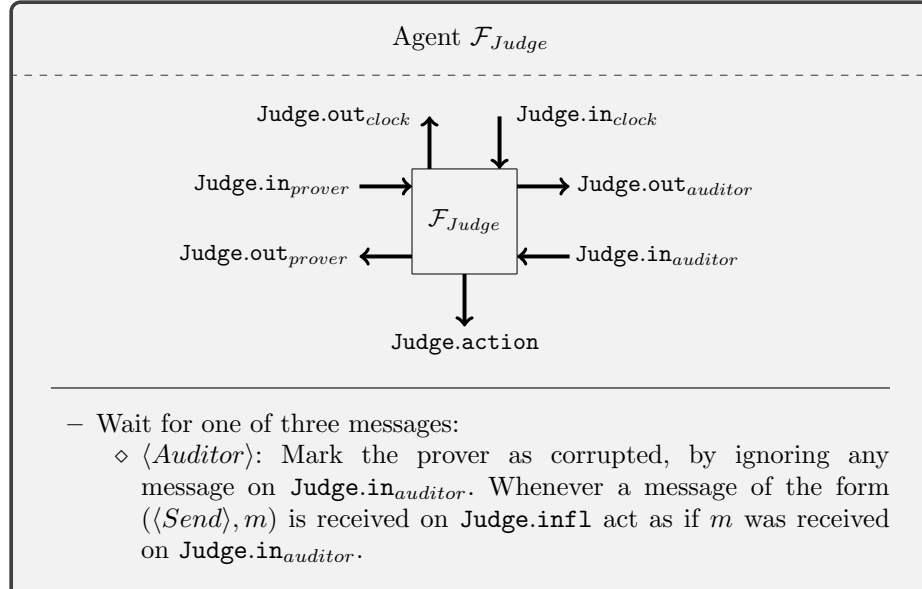
This formulation allows instantiation of the functionality using a blockchain which offers ‘finality’ guarantees; ensuring that the view of the honest parties cannot be rolled back past finalized blocks. Furthermore, one needs to assume that the view of any node is at most Δ blocks behind the most recently finalized block.

The $\mathcal{F}_{Channel}$ functionality models an authenticated and encrypted channel between the prover and auditor, which guarantees in-order delivery of messages:



- * Whenever a message of the form $(\langle Send \rangle, m)$ is received on **Channel.infl** act as if m was received on **Channel.in_p**.
 - * Whenever a message of the form $(\langle Recv \rangle, m)$ is received on **Channel.infl** output m on **Channel.out_p**.
 - * Whenever a message m is output on **Channel.out_p** output m on **Channel.leak** instead of **Channel.out_p**.
 - ◇ $\langle Honest \rangle$. Indicating no corruption.
- Ignore any subsequent corruption messages.
- On input m on **Channel.in_{prover}**:
 - Push m to the back of $M_{auditor}$
 - Output $(\langle Auditor \rangle, |m|)$ on **Channel.leak**.
 - On input m on **Channel.in_{auditor}**:
 - Push m to the back of M_{prover}
 - Output $(\langle Prover \rangle, |m|)$ on **Channel.leak**.
 - On input $(\langle Deliver \rangle, p)$ on **Channel.infl**:
 - If M_p is not empty, pop the front-most element m and output m on **Channel.in_p**.

The judge is instantiated with a description D of its transition function, which both parties must agree upon. Whenever the judge receives input, this is provided to all parties and leaked, reflecting that the state of the judge is completely public. The judge furthermore has access to a clock functionality and an ‘action’ port, which will later correspond to the action port of the \mathcal{F}_{Swap} functionality:



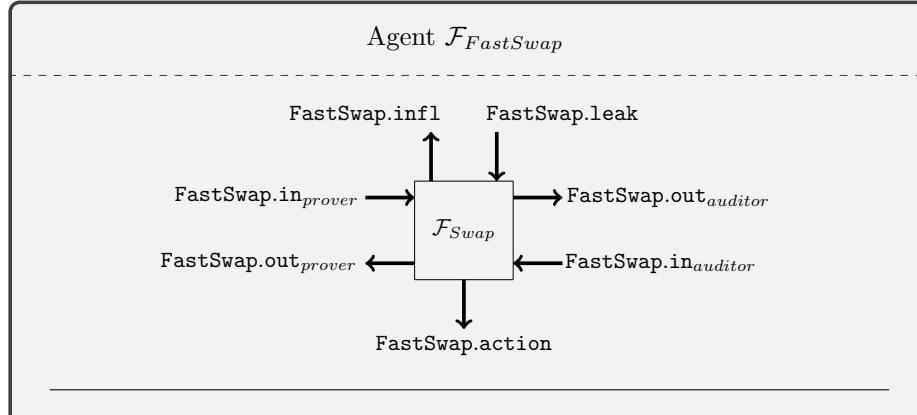
- ◊ $\langle Prover \rangle$: Mark the prover as corrupted, by ignoring any message on $Judge.in_{prover}$. Whenever a message of the form $\langle Send \rangle, m$ is received on $Judge.infl$ act as if m was received on $Judge.in_{prover}$.

- ◊ $\langle Honest \rangle$. Indicating no corruption.

Ignore any subsequent corruption messages.

- Whenever t_{new} is received on $Judge.in_{clock}$, store $t \leftarrow t_{new}$.
- On input D on $Judge.in_{auditor}$: output D on $Judge.leak$, output D on $Judge.out_{auditor}$, store D .
- On input D' on $Judge.in_{prover}$: if $D \neq D'$, output \perp on $Judge.out_{prover}$, output \perp on $Judge.out_{auditor}$ and abort the protocol, by ignoring any subsequent messages on all in ports. Otherwise set $S \leftarrow \epsilon$ and begin processing input messages.
- On input m on $Judge.in_{auditor}$: output (m, t) on $Judge.leak$, output (m, t) on $Judge.out_{auditor}$, output (m, t) on $Judge.out_{prover}$, update the state $(S, r) \leftarrow D(S, \langle Auditor \rangle, m, t)$, if $r \neq \epsilon$ output r on $Judge.action$.
- On input m on $Judge.in_{prover}$: output (m, t) on $Judge.leak$, output (m, t) on $Judge.out_{auditor}$, output (m, t) on $Judge.out_{prover}$, update the state $(S, r) \leftarrow D(S, \langle Prover \rangle, m, t)$, if $r \neq \epsilon$ output r on $Judge.action$.

The *FastSwap* functionality enables the two parties to agree on the initial state of a authenticated computation scheme, then allows the prover to input an environment. If repeated application of Step on the initial state with the given environment terminates in an accepting state the functionality outputs 1 on $FastSwap.action$, otherwise the functionality outputs 0. When both parties are honest the functionality leaks only the environment and the accepting/rejecting outcome of the computation, in particular it does not leak the initial state:

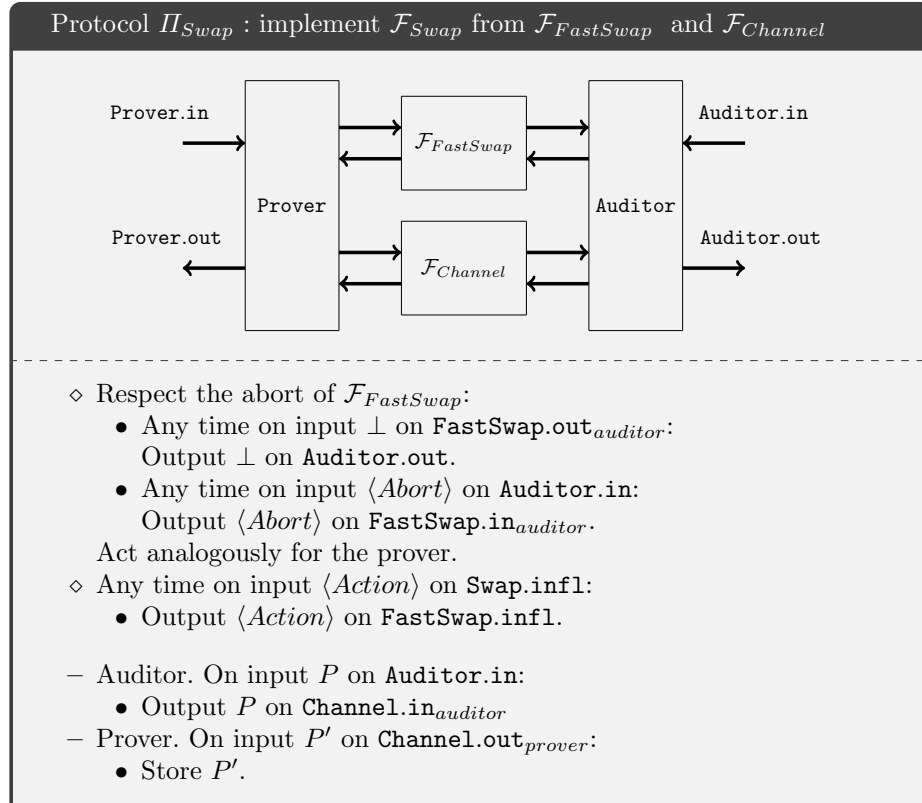


Initialization: set $\text{can_abort} \leftarrow 1$

- Wait for one of three messages on FastSwap.infl .
 - ◊ $\langle \text{Auditor} \rangle$: Mark the prover as corrupted, by ignoring any message on $\text{FastSwap.in}_{\text{auditor}}$. Whenever a message of the form $\langle \text{Send} \rangle, m$ is received on FastSwap.infl act as if m was received on $\text{FastSwap.in}_{\text{auditor}}$. Whenever a message m is output on $\text{FastSwap.out}_{\text{auditor}}$, also output m on FastSwap.leak
 - ◊ $\langle \text{Prover} \rangle$: Mark the prover as corrupted, by ignoring any message on $\text{FastSwap.in}_{\text{prover}}$. Whenever a message of the form $\langle \text{Send} \rangle, m$ is received on FastSwap.infl act as if m was received on $\text{FastSwap.in}_{\text{prover}}$. Whenever a message m is output on FastSwap.out , also output m on FastSwap.leak
 - ◊ $\langle \text{Honest} \rangle$. Indicating no corruption.
- Ignore any subsequent corruption messages.
- Any time, on input $\langle \text{Abort} \rangle$ on FastSwap.infl , $\text{FastSwap.in}_{\text{auditor}}$ or $\text{FastSwap.in}_{\text{prover}}$ and if $\text{can_abort} = 1$, then abort the protocol:
 - Output \perp on $\text{FastSwap.out}_{\text{prover}}$.
 - Output \perp on $\text{FastSwap.out}_{\text{auditor}}$.
 - Output \perp on FastSwap.leak .
 - Ignore any further messages on any in port.
- On input I' on $\text{FastSwap.in}_{\text{auditor}}$:
 - Store I' .
 - If the prover is corrupted, output I' on FastSwap.leak .
 - Output $\langle \text{Input} \rangle$ on FastSwap.leak
- On input I on $\text{FastSwap.in}_{\text{prover}}$, when I' has been set:
 - Compute $S \leftarrow \text{Initial}(I)$.
 - Compute $S' \leftarrow \text{Initial}(I')$.
 - If $S \neq S'$ then abort the protocol (as if $\langle \text{Abort} \rangle$ was received).
- On input e on $\text{FastSwap.prover}_{\text{in}}$:
 - Set $\text{can_abort} \leftarrow 0$
 - Output e on $\text{FastSwap.out}_{\text{auditor}}$.
 - Output e on FastSwap.leak .
 - If either party is corrupted leak the entire state of the functionality on FastSwap.leak : every message sent and received by the functionality.
- On input $\langle \text{Action} \rangle$ on FastSwap.infl , when $\text{can_abort} = 0$:
 - Compute $n \leftarrow \text{Terminate}(e, S)$.
 - Output $\text{StepN}(e, S, n) \stackrel{?}{=} \langle \text{Accept} \rangle$ on FastSwap.action and FastSwap.leak .

We implement the \mathcal{F}_{Swap} functionality using: $\mathcal{F}_{FastSwap}$, $\mathcal{F}_{Channel}$, a semantically secure encryption scheme (Definition 2)³ and a sufficiently expressive authenticated computational structure scheme (Definition 9):

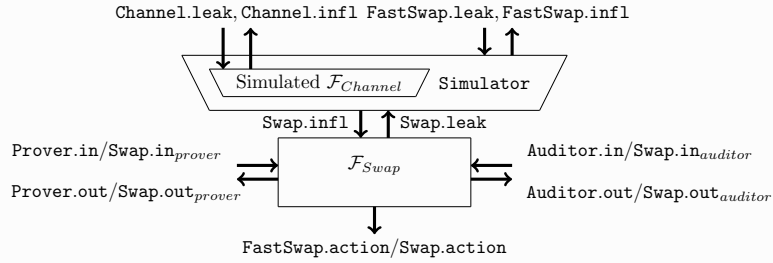
The authenticated computational structure scheme must allow expression of the $\text{Dec} : \mathcal{K}_\kappa \times \mathcal{C}_\kappa \rightarrow \mathcal{M}$ function as well as the set of predicates. The set of environments for the computational structure scheme must contain \mathcal{K}_κ . Furthermore we assume that input descriptions \mathcal{I} can be provided in the form (P, W) where P is the description of a predicate and W is the input to the predicate, st.: repeated application of the $\text{Step} : \mathcal{E} \times \mathcal{S} \rightarrow \mathcal{S}$ functions computes $P(e, W)$, where $e \in \mathcal{E}$ is the environment. We can therefore transform the problem in \mathcal{F}_{Swap} of evaluating the predicate P on w , into the problem of repeatedly applying the Step function to the initial state described by $I = (P \circ \text{Dec}(e, \cdot), W)$ where $W \leftarrow \text{Enc}(e, w)$, with $e \in \mathcal{K}_\kappa$ provided as the environment of the Step function. Intuitively this enables us to swap a constant size key in place of the actual witness, which additionally provides semantic hiding of the witness from the auditor while the protocol can still be aborted and from the environment in case of honest execution.



³ For which we require a computable description, hence the application of the IND-CPA encryption scheme is non-blackbox.

- Output P' on Prover.out .
- Prover. On input w on Prover.in :
 - Sample $k \xleftarrow{\$} \mathcal{K}_\kappa$.
 - Compute $E \leftarrow \text{Enc}(k, w)$.
 - Define $I = (P' \circ \text{Dec}, E)$.
 - Output E on $\text{Channel.in}_{\text{prover}}$.
 - Output I on $\text{FastSwap.in}_{\text{prover}}$.
- Auditor. On input E' on $\text{Channel.out}_{\text{auditor}}$:
 - Define $I' = (P \circ \text{Dec}, E')$
 - Output I' on $\text{FastSwap.in}_{\text{auditor}}$.
- Prover. On input $\langle \text{Swap} \rangle$ on Prover.in :
 - Output k on $\text{FastSwap.in}_{\text{prover}}$
- Auditor. On input k on $\text{FastSwap.out}_{\text{auditor}}$.
 - Output $\text{Dec}(k, E')$ on Auditor.out

Simulator $\mathcal{S}_{\text{Swap}}$: simulate Π_{Swap} using $\mathcal{F}_{\text{Swap}}$



Respect the abort: any time, on input $\langle \text{Abort} \rangle$ on FastSwap.infl , $\text{FastSwap.in}_{\text{auditor}}$ or $\text{FastSwap.in}_{\text{prover}}$: Output $\langle \text{Abort} \rangle$ on Swap.infl , $\text{Swap.in}_{\text{auditor}}$ or $\text{Swap.in}_{\text{prover}}$ respectively. On \perp on Swap.leak , output \perp on FastSwap.leak

Wait for the corruption pattern for both $\mathcal{F}_{\text{FastSwap}}$ and $\mathcal{F}_{\text{Channel}}$ (the class of environments is assumed corruption respecting: corrupting the same parties for every functionality):

Case 1. Neither party is corrupted:

- On input $|P|$ on Swap.leak :
 - Simulate sending $0^{|P|}$ on $\text{Channel.in}_{\text{auditor}}$.
- On input $|w|$ on Swap.leak :
 - Sample $k \xleftarrow{\$} \mathcal{K}_\kappa$.

- Compute $E \leftarrow \text{Enc}(k, 0^{|w|})$.
- Simulate sending E on $\text{Channel.in}_{\text{prover}}$.
- On simulated input E on $\text{Channel.in}_{\text{auditor}}$:
 - Output $\langle \text{Input} \rangle$ on FastSwap.leak .
- On input $\langle \text{Action} \rangle$ on FastSwap.infl :
 - Output $\langle \text{Action} \rangle$ on Swap.infl .
- On $P(w)$ on Swap.leak :
 - Output k on FastSwap.leak (as ‘ e ’).
 - Output $P(w)$ on FastSwap.leak .

Case 2. Auditor is corrupted, Prover is honest:

(**Note:** not simulatable without random oracles, see proof section.)

- On input P on Swap.leak :
 - Simulate sending P on $\text{Channel.in}_{\text{auditor}}$.
- On simulated input P' on $\text{Channel.out}_{\text{prover}}$:
 - Store P'
- On input $|w|$ on Swap.leak :
 - Sample $k \xleftarrow{\$} \mathcal{K}_\kappa$.
 - Compute $E \leftarrow \text{Enc}(k, 0^{|w|})$.
 - Simulate sending E on $\text{Channel.in}_{\text{prover}}$.
- On simulated input E' on $\text{Channel.in}_{\text{auditor}}$:
 - Output $\langle \text{Input} \rangle$ on FastSwap.leak .
- On $w, P(w)$ on Swap.leak :
 - Reprogram Dec using the RO, such that $\text{Dec}(k, E) = w$.
 - Output k on FastSwap.leak (as ‘ e ’).
 - Output $w, P(w)$ on FastSwap.leak .
- On input $\langle \text{Action} \rangle$ on FastSwap.infl :
 - Output $\langle \text{Action} \rangle$ on Swap.infl .

Case 3. Auditor is honest, Prover is corrupted:

- On input P on Swap.leak :
 - Simulate sending P on $\text{Channel.in}_{\text{auditor}}$.
- On simulated input P' on $\text{Channel.out}_{\text{prover}}$:
 - Store P'
- On input w on Swap.leak :
 - Sample $k \xleftarrow{\$} \mathcal{K}_\kappa$.
 - Compute $E \leftarrow \text{Enc}(k, w)$.
 - Simulate sending E on $\text{Channel.in}_{\text{prover}}$.
- On simulated input E' on $\text{Channel.in}_{\text{auditor}}$:
 - Output $\langle \text{Input} \rangle$ on FastSwap.leak .
- On $w, P(w)$ on Swap.leak :
 - Output k on FastSwap.leak (as ‘ e ’).
 - Output $w, P(w)$ on FastSwap.leak .
- On input $\langle \text{Action} \rangle$ on FastSwap.infl :
 - Output $\langle \text{Action} \rangle$ on Swap.infl .

Lemma 1 ($\Pi_{\text{Swap}} \diamond \mathcal{F}_{\text{FastSwap}} \diamond \mathcal{F}_{\text{Channel}} \geq_{\text{comp}} \mathcal{F}_{\text{Swap}}$). Π_{Swap} implements $\mathcal{F}_{\text{Swap}}$ using $\mathcal{F}_{\text{FastSwap}}$ and $\mathcal{F}_{\text{Channel}}$ with respect to all computationally bounded (PPT) environments.

Proof. By case analysis on the corruption pattern of the environment:

Case 1. Neither party is corrupted:

Consider the hybrid $\mathcal{H}_{\text{Swap}}$ which is equal to $\mathcal{S}_{\text{Swap}}$, except where w is extracted from $\mathcal{F}_{\text{Swap}}$ and E is derived as $E \leftarrow \text{Enc}(k, w)$. The difference in the distributions is the leakage on **FastSwap.leak**: $\mathcal{S}_{\text{Swap}}$ leaks $E' \leftarrow \text{Enc}(k, 0^{|w|})$, since $|0^{|w|}| = |w|$ the distributions must be computationally indistinguishable by the assumption that $\text{Enc} : \mathcal{K}_\kappa \times \mathcal{M} \rightarrow \mathcal{C}_\kappa$ is a CPA secure encryption scheme (Definition 2).

Case 2. Auditor is corrupted, Prover is honest:

We assume that Enc, Dec is non-committing and implemented using a random oracle (e.g. using a construction from [7]). Consider again a hybrid \mathcal{H}_w which is equal to $\mathcal{S}_{\text{Swap}}$, except where w is extracted from $\mathcal{F}_{\text{Swap}}$ and E is replaced with $E_w \leftarrow \text{Enc}(k, w)$ since $|0^{|w|}| = |w|$, E_w and E must be computationally indistinguishable by the assumption that $\text{Enc} : \mathcal{K}_\kappa \times \mathcal{M} \rightarrow \mathcal{C}_\kappa$ is a CPA secure encryption scheme (Definition 2). Furthermore since $k \stackrel{\$}{\leftarrow} \mathcal{K}_\kappa$ the probability that the environment has queried the oracle on any of the queries made during $\text{Dec}(k, E)$ prior to receiving k is negligible, hence reprogramming is successful with overwhelming probability. Hence \mathcal{H}_w and $\mathcal{S}_{\text{Swap}}$ are computationally indistinguishable.

A simulatable alternative in the standard model is to deploy non-committing symmetric encryption, e.g. encrypting with a one-time pad, however this significantly impedes efficiency since k must have the same size as the witness and hence the communication with the judge would be linear in the size of the witness.

Case 3. Auditor is honest, Prover is corrupted:

Since a corrupted prover leaks the secret witness of the protocol (before `can_abort` $\leftarrow 0$), this simulation is trivial and the distributions are equal.

The inability to simulate this protocol in the standard model whenever $|k| < |w|$ is inherent to the structure of the scheme: When the auditor is corrupt we need to output to the environment a message E which is indistinguishable from an encryption of the witness, however since the prover is honest only $|w|$ is leaked, hence E must be uncorrelated with w . Later we must output k to the environment st. $\text{Dec}(k, E) = w$ (except with negligible probability), however this implies communication at rates greater than channel capacity: since E is uncorrelated with the message w it could be sampled the receiver directly, then w is transmitted by sending k .

In practical terms this means that the auditor can obtain an encryption of the witness and then abort the protocol without paying. We note that the prior works mentioned earlier (would) also require such non-committing encryption to achieve simulation security. This is due to the similarity between all these scheme of exchanging a decryption key which enables decryption of the witness, which has been encrypted and exchanged ‘off-chain’ priorly.

6 The FastSwap Protocol

6.1 Protocol

The protocol is parameterized by a timeout Δ_{action} . The judge maintains a timer D_{action} , when D_{action} expires the judge outputs the current value of the *result* variable on the action port as the output of the protocol⁴. To simplify the description we assume that the transition function of the judge is sent to the judge functionality by both players at the start of the protocol and that upon receiving \perp the honest party aborts the protocol. This allows us to treat the judge as a third party in the protocol.

The overall idea of FastSwap is to have both parties agree on a commitment of the initial state, with both parties knowing the opening of the commitment. In case of contingent payments the auditor/buyer would then deposit funds at the judge. Subsequently the prover reveals the environment by sending it directly to the judge, at this point a unique⁵ execution trace is now defined by the environment and the initial state inside the commitment. In the honest case, where the trace is accepting, the auditor simply lets the timer D_{action} expire, after which the action is assumed complete:

FastSwap : Honest Execution

- Auditor:
 - Sample $R' \xleftarrow{\$} \mathcal{R}_\kappa$.
 - Compute $S'_1 \leftarrow \text{Initial}(I')$.
 - Compute $T'_1 \leftarrow \text{Tag}(S'_1)$.
 - Compute $C' \leftarrow \text{Comm}(R', T'_1)$.
 - Send R' to the prover.
 - Send C' to the judge.
- Judge:
 - Receive C' from the auditor.

⁴ In blockchain applications for contingency payments, the judge contract can be converted into a wallet contract after the expiry of D_{action} where *result* denotes which party is allowed to withdraw the funds.

⁵ By ‘unique’, we mean that neither party can break the binding property of the commitment scheme and **Tag** function, hence can only possess one such trace. Since the state is significantly larger than the commitments it is clearly not unique in the strict sense.

- Set $result \leftarrow 0$.
- Start D_{action} with timeout Δ_{action} .
- Prover:
 - Receive R from the auditor.
 - Compute $S_1 \leftarrow \text{Initial}(I)$.
 - Compute $T_1 \leftarrow \text{Tag}(S_1)$.
 - Compute $C \leftarrow \text{Comm}(R, T_1)$.
 - If $C \neq C'$ (from the judge) abort the protocol.
 - Send e to the judge.
- Judge:
 - Receive e from the prover.
 - Set $result \leftarrow 1$.
 - Reset D_{action} with timeout Δ_{action} .
- Auditor:
 - Compute $m \leftarrow \text{Terminate}(e, S'_1)$.
 - Compute $S'_m \leftarrow \text{StepN}(e, S'_1, m)$.
 - If $S'_m = \langle \text{Accept} \rangle$ terminate the protocol.
 - Otherwise proceed to dispute resolution (see below).

The intuition for the dispute resolution protocol is to maintain two pointers l and r into the computation trace of the prover. The pointer l will always point to a computation step that both parties agree on (initially S_1 , the state inside the commitment). The pointer r (when defined), will point to a computation step where $S_r \neq S'_r$. We then search for the greatest value of l and the smallest value of r , by using an interactive binary search mediated by the judge to ensure message delivery. Eventually $r - l = 1$ and the prover uses the authenticated computation structure scheme to show correct transition from S_l to S_r , with a succinct proof:

FastSwap : Dispute Resolution

- Auditor:
 - Send $\langle \text{Dispute} \rangle$ to the judge.
- Judge:
 - Set $result \leftarrow 0$
 - Set $l \leftarrow 1, r \leftarrow \perp$. Define $m = (r - l)/2$ (initially $m = \perp$).
 - Reset D_{action} with timeout Δ_{action}
- While $r = \perp$ or $r - l > 1$:
 - Prover:
 - * If $r = \perp$ (first iteration):
 - Compute $n \leftarrow \text{Terminate}(e, S_1)$.
 - Locally set $r \leftarrow n$.
 - Send n to the judge.
 - * Compute $S_m \leftarrow \text{StepN}(e, S_1, m)^a$.

- * Compute $T_m \leftarrow \text{Tag}(S_m)$.
- * Send T_m to the judge.
- Judge:
 - * If $r = \perp$ (first iteration), set $r \leftarrow n$.
 - * Store T_m .
 - * Set $result \leftarrow 1$
 - * Reset D_{action} with timeout Δ_{action}
- Auditor:
 - * Compute $S'_m \leftarrow \text{StepN}(e, S'_1, m)$
 - * If $T_m = \text{Tag}(S'_m)$, send $\langle Left \rangle$ to the judge.
 - * Otherwise, send $\langle Right \rangle$ to the judge.
- Judge:
 - * If received $\langle Left \rangle$, set $l \leftarrow m$, set $T_l \leftarrow T_m$.
 - * If received $\langle Right \rangle$, set $r \leftarrow m$, set $T_r \leftarrow T_m$.
 - * Set $result \leftarrow 0$
 - * Reset D_{action} with timeout Δ_{action}
- Prover:
 - Compute $S_l \leftarrow \text{StepN}(e, S_1, l)$.
 - Compute $\pi_{l \rightarrow r} \leftarrow \text{Prove}(e, S_l)$.
 - If $l = 1$, send $(\pi_{l \rightarrow r}, R, T_1)$ to the judge.
 - Otherwise send $\pi_{l \rightarrow r}$ to the judge.
- Judge:
 - Set $result \leftarrow 1$.
 - If $l = 1$ and $\text{Open}(R, T_1, C') = 0$, set $result \leftarrow 0$.
 - If $r = n$ and $T_r \neq \text{Tag}(\langle Accept \rangle)$, set $result \leftarrow 0$
 - If $\text{Verify}(e, T_l, T_r, \pi_{l \rightarrow r}) = 0$, set $result \leftarrow 0$.


^a Note that m is defined at this point.

The dispute resolution protocol additionally guarantees that if the output is 1, the auditor is corrupted, if the output is 0, the prover must be corrupted. This allows the judge to optionally trigger penal action towards the dishonest party (e.g. in a smart contract environment, this might be seizing collateral added to the contract during the start of the protocol) in the cases where the dispute resolution protocol is triggered.

6.2 Security Proof

We simulate $\Pi_{FastSwap}$ using $\mathcal{F}_{FastSwap}$ as follows:

Simulator $\mathcal{S}_{FastSwap}$: simulate $\Pi_{FastSwap}$ using $\mathcal{F}_{FastSwap}$



Case 1. Neither party is corrupted:

- On input $\langle Input \rangle$ on `FastSwap.leak`:
 - Sample $R' \xleftarrow{\$} \mathcal{R}_\kappa$.
 - Compute $C' \leftarrow \text{Comm}(R', \epsilon)$.
 - Simulate output R' on `Channel.in_auditor`.
 - Simulate output C' on `Judge.in_auditor`.
- On input e on `FastSwap.leak`:
 - Wait for simulated input R on `Channel.out_prover`.
 - Simulate output e on `Judge.in_prover`.
- On expiry of D_{action} (inside judge simulation):
 - Output $\langle Action \rangle$ on `FastSwap.infl`.
 - Output 1 on `Judge.leak`.

Case 2. Auditor is corrupted, prover is honest:

- On input I' (auditors initial state) on `FastSwap.leak`:
 - Sample $R' \xleftarrow{\$} \mathcal{R}_\kappa$.
 - Compute $S'_1 \leftarrow \text{Initial}(I')$.
 - Compute $T'_1 \leftarrow \text{Tag}(S'_1)$.
 - Compute $C' \leftarrow \text{Comm}(R', T'_1)$.
 - Simulate output R' on `Channel.in_auditor`.
 - Simulate output C' on `Judge.in_auditor`.
- On $e, S_{FastSwap}$ on `FastSwap.leak ($S_{FastSwap}$ is the leaked state):

 - Store $S_{FastSwap}$.
 - Simulate output e on Judge.in_prover.`
- On $\langle Dispute \rangle$ on `Judge.leak`.
 - Simulate the dispute resolution protocol using $S_{FastSwap}$, by observing the messages from the corrupted auditor using `Judge.leak` and simulating the messages on `Judge.in_prover` of the honest prover according to the dispute resolution protocol.
- On expiry of D_{action} :
 - Output $\langle Action \rangle$ on `FastSwap.infl`.
 - Obtain res on `FastSwap.leak`: output res on `Judge.leak`.

Case 3. Auditor is honest, prover is corrupted:

Due to $\mathcal{F}_{FastSwap}$ leaking the auditors initial state when the prover is corrupted the simulation is very similar to the case of a corrupted auditor:

- On input I' (auditors initial state) on **FastSwap.leak**:
 - Sample $R \xleftarrow{\$} \mathcal{R}_\kappa$.
 - Compute $S'_1 \leftarrow \text{Initial}(I')$.
 - Compute $T'_1 \leftarrow \text{Tag}(S'_1)$.
 - Compute $C' \leftarrow \text{Comm}(R, T'_1)$.
 - Simulate output R on **Channel.in_{prover}**.
 - Simulate output C on **Judge.in_{auditor}**.
- On $e, S_{FastSwap}$ on **FastSwap.leak** ($S_{FastSwap}$ is the leaked state):
 - Store $S_{FastSwap}$.
 - Output e on **Judge.leak**.
- On $\langle \text{Dispute} \rangle$ on **Judge.leak**.
 - Simulate the dispute resolution protocol using $S_{FastSwap}$, by observing the messages from the corrupted auditor using **Judge.leak** and simulating the messages on **Judge.in_{prover}** of the honest auditor according to the dispute resolution protocol.
- On expiry of D_{action} :
 - Output $\langle \text{Action} \rangle$ on **FastSwap.infl**.
 - Obtain res on **FastSwap.leak**: output res on **Judge.leak**.

Lemma 2 ($\Pi_{FastSwap} \diamond \mathcal{F}_{Judge} \diamond \mathcal{F}_{Channel} \diamond \mathcal{F}_{Clock} \geq_{comp} \mathcal{F}_{FastSwap}$). $\Pi_{FastSwap}$ implements $\mathcal{F}_{FastSwap}$ using \mathcal{F}_{Judge} , $\mathcal{F}_{Channel}$ and \mathcal{F}_{Clock} with respect to all computationally bounded (PPT) environments.

Proof. By case analysis on the corruption pattern:

Case 1. Neither party is corrupted:

The prover poses a valid witness and $\langle \text{Dispute} \rangle$ is not sent to the judge by the auditor. Hence the leakage in the real execution is comprised solely of the leakage in the honest execution part of the protocol. The output on **FastSwap.action** is always 1, if neither party aborts and the output 1 on **Judge.leak** is consistent with the final value of result outputted on **Judge.action** in the real execution.

Case 2. Auditor is corrupted, prover is honest:

The leakage from the simulation of the honest part of the protocol has exactly the same distribution as the real protocol. We therefore focuses on the simulation of the dispute resolution (recall that we obtain the entire state of $\mathcal{F}_{FastSwap}$), in particular that the leakage is consistent with the output on **FastSwap.action**.

Since the prover is honest it follows that $\langle \text{Accept} \rangle = \text{StepN}(e, S_1, n)$ where $n \leftarrow \text{Terminate}(e, S_1)$. Except with negligible probability $S_1 = S'_1$ by computational integrity of the authenticated computation scheme (Definition 10) and binding of the commitment scheme (Definition 6). We claim an invariant of the loop in the protocol:

$$l < r \leq n \text{ and } T_l = \text{Tag}(S_l) \text{ and } T_r = \text{Tag}(S_r)$$

This is immediately obvious from inspection of the dispute protocol. Upon termination of the loop $r - l = 1$ and $\text{Verify}(e, T_l, T_r, \pi_{l \rightarrow r}) = 1$ with probability 1 (by completeness of the authenticated computation scheme), furthermore whenever $r = n$, we have $S_r = \langle \text{Accept} \rangle$ hence $T_r = \text{Tag}(\langle \text{Accept} \rangle)$ also with probability 1. Therefore the simulated judge always outputs 1, which is consistent with **FastSwap.action**.

Case 3. Auditor is honest, prover is corrupted:

The leakage from the simulation of the honest part of the protocol has exactly the same distribution as the real protocol. We therefore focuses on the simulation of the dispute resolution (recall that we obtain the entire state of $\mathcal{F}_{\text{FastSwap}}$), in particular that the leakage is consistent with the output on **FastSwap.action**.

Since the auditor is honest it follows that $\langle \text{Accept} \rangle \neq \text{StepN}(e, S'_1, m)$ where $m \leftarrow \text{Terminate}(e, S'_1)$, hence the judge should output 0. We first establishes an invariant of the loop in the protocol: $T_l = \text{Tag}(S'_l)$ and at least one of the following holds:

- ◇ $l < r \leq n$ and $T_r \neq \text{Tag}(S'_r)$
- ◇ $l < r = n$ and $T_r \neq \text{Tag}(\langle \text{Accept} \rangle)$

The invariant holds initially where $r = n$ and $l = 1$, since $T_1 = \text{Tag}(S'_1)$ is established during the honest part of the protocol and $\forall i \in [1, n] : S'_i \neq \langle \text{Accept} \rangle$ (otherwise $\langle \text{Accept} \rangle = \text{StepN}(e, S'_1, m)$ as well). During the protocol the corrupted auditor provides T_w with $l < w < r$ and the invariant is maintained:

- If $T_w = \text{Tag}(S'_w)$, then $l \leftarrow w$.
Hence $T_l = \text{Tag}(S'_l)$ is maintained and r, T_r is unchanged.
- If $T_w \neq \text{Tag}(S'_w)$, then $r \leftarrow w$.
Hence $T_r \neq \text{Tag}(S'_r)$ is established and l, T_l is unchanged.

Upon termination of the loop: $r - l = 1$, $T_l = \text{Tag}(S'_l)$ and:

- If $r = n$ and $T_r \neq \text{Tag}(\langle \text{Accept} \rangle)$, the output is always 0.
- If $T_r \neq \text{Tag}(S'_r)$, then $\text{Verify}(e, T_l, T_r, \pi_{l \rightarrow r}) = 0$ except with only negligible probability, by computational integrity (Definition 10) of the authenticated computation scheme. Hence the output is 0.

7 Instantiation of FastSwap

In this section we propose a simple ‘Ethereum-like’ instantiation of the FastSwap protocol, based on an authenticated Patricia trie over a sparse memory space. The state is a tuple $(pc, I, \mathcal{R}_{reg}, S)$ consisting of:

- An instruction pointer $pc \in \mathbb{N}_+$ pointing to a cell.
- An optional word-sized instruction I (which might be ϵ).
- A register bank \mathcal{R}_{reg} containing word-sized registers r_1, \dots, r_n .
- An authenticated data structure S over a memory space of M words.

The memory space is provided by simply ameliorating a Patricia trie⁶ with a superimposed Merkle tree (see e.g. [6] appendix D for details), which allows proving memory lookups by providing at most $2 \cdot \log(M)$ hashes of size κ , where M is the size of the memory space. We let $\text{Prove}_{\text{Patricia}}$, $\text{Verify}_{\text{Patricia}}$, $\text{Apply}_{\text{Patricia}}$ be the associated algorithms of the authenticated Patricia trie. We let $\mathcal{R}_{reg}[r_i]$ denote the looking up the value of the register r_i and $\mathcal{R}_{reg}[r_i \leftarrow v]$ denote a new register bank, where the value v is assigned to the register r_i .

Tag function. We define $\text{Tag}((pc, I, \mathcal{R}_{reg}, S)) = \text{CRH}((\text{Tag}_{\text{Patricia}}(S), pc, I, \mathcal{R}_{reg}))$. Meaning the full register bank, Merkle root, current instruction and program counter is provided during the verification.

Step function. For efficiency and simplicity reasons the instantiation limits the number of operations on the memory space during every step to at most one, this is done by using a ‘2-cycle’ register machine, where every instruction in the instruction set takes two applications of **Step** to execute. The **Step** function operates as follows, with the state being matched occurring on the left:

$$\begin{aligned}
 \text{Step}(e, (pc, \epsilon, \mathcal{R}_{reg}, S)) &:= (pc, I, \mathcal{R}_{reg}, S) \\
 &\quad \text{where } (*, I, *) \leftarrow \text{Apply}_{\text{Patricia}}(S, \text{load}(pc)) \\
 \text{Step}(e, (pc, \text{load}(i, j), \mathcal{R}_{reg}, S)) &:= (pc + 1, \epsilon, \mathcal{R}_{reg}[r_i \leftarrow M], S) \\
 &\quad \text{where } (*, M, *) \leftarrow \text{Apply}_{\text{Patricia}}(S, \text{load}(\mathcal{R}_{reg}[r_j])) \\
 \text{Step}(e, (pc, \text{store}(i, j), \mathcal{R}_{reg}, S)) &:= (pc + 1, \epsilon, \mathcal{R}_{reg}, S') \\
 &\quad \text{where } (S', *, *) \leftarrow \text{Apply}_{\text{Patricia}}(S, \text{store}(\mathcal{R}_{reg}[r_i], \mathcal{R}_{reg}[r_j])) \\
 \text{Step}(e, (pc, \text{jump}(i, j), \mathcal{R}_{reg}, S)) &:= (\Delta, \epsilon, \mathcal{R}_{reg}, S) \\
 &\quad \text{where if } \mathcal{R}_{reg}[r_j] > 0 \text{ then } \Delta = \mathcal{R}_{reg}[r_i] \text{ else } \Delta = (pc + 1) \\
 \text{Step}(e, (pc, \text{mult}(i, j), \mathcal{R}_{reg}, S)) &:= (pc + 1, \epsilon, \mathcal{R}_{reg}[r_i \leftarrow a \cdot b], S) \\
 &\quad \text{where } a = \mathcal{R}_{reg}[r_i], b = \mathcal{R}_{reg}[r_j] \\
 \text{Step}(e, (pc, \text{add}(i, j), \mathcal{R}_{reg}, S)) &:= (pc + 1, \epsilon, \mathcal{R}_{reg}[r_i \leftarrow a + b], S) \\
 &\quad \text{where } a = \mathcal{R}_{reg}[r_i], b = \mathcal{R}_{reg}[r_j]
 \end{aligned}$$

⁶ Radix tree with a radix of 2.

$$\text{Step}(e, (pc, \text{env}(i), \mathcal{R}_{reg}, S)) := (pc + 1, \epsilon, \mathcal{R}_{reg}[r_i \leftarrow e], S)$$

Additionally there are two predefined values of pc corresponding to an accepting and a rejecting state. If either of these addresses are reached, **Step** replaces the state with some predefined canonical $\langle \text{Accept} \rangle$ or $\langle \text{Reject} \rangle$ state not otherwise reachable, regardless of the contents of the register bank or memory space:

$$\begin{aligned} \text{Step}(e, (pc_{\text{accept}}, \epsilon, \mathcal{R}_{reg}, S)) &:= \langle \text{Accept} \rangle \\ \text{Step}(e, (pc_{\text{reject}}, \epsilon, \mathcal{R}_{reg}, S)) &:= \langle \text{Reject} \rangle \end{aligned}$$

The **Step** function can always be made complete by mapping any non-conforming state to $\langle \text{Reject} \rangle$.

Prove function. The prove function outputs the register bank and a proof for the authenticated Patricia trie in case of a memory operation:

$$\begin{aligned} \text{Prove}(e, (pc, \epsilon, \mathcal{R}_{reg}, S)) &:= (\epsilon, pc, \mathcal{R}_{reg}, \text{Tag}_{\text{Patricia}}(S), I, \pi_{\text{lookup}}) \\ &\text{where } (*, I, \pi_{\text{lookup}}) \leftarrow \text{Apply}_{\text{Patricia}}(S, \text{lookup}(pc)) \end{aligned}$$

$$\begin{aligned} \text{Prove}(e, (pc, \text{load}(i, j), \mathcal{R}_{reg}, S)) &:= (\text{load}(i, j), pc, \mathcal{R}_{reg}, \text{Tag}_{\text{Patricia}}(S), R, \pi_{\text{lookup}}) \\ &\text{where } (*, R, \pi_{\text{lookup}}) \leftarrow \text{Apply}_{\text{Patricia}}(S, \text{lookup}(\mathcal{R}_{reg}[r_j])) \end{aligned}$$

$$\begin{aligned} \text{Prove}(e, (pc, \text{store}(i, j), \mathcal{R}_{reg}, S)) &:= (\text{store}(i, j), pc, \mathcal{R}_{reg}, \text{Tag}_{\text{Patricia}}(S), \text{Tag}_{\text{Patricia}}(S'), \pi_{\text{store}}) \\ &\text{where } (S', *, \pi_{\text{store}}) \leftarrow \text{Apply}_{\text{Patricia}}(S, \text{store}(\mathcal{R}_{reg}[r_i], \mathcal{R}_{reg}[r_j])) \end{aligned}$$

$$\text{Prove}(e, (pc, \text{jump}(i, j), \mathcal{R}_{reg}, S)) := (\text{jump}(i, j), pc, \mathcal{R}_{reg}, \text{Tag}_{\text{Patricia}}(S))$$

$$\text{Prove}(e, (pc, \text{mult}(i, j), \mathcal{R}_{reg}, S)) := (\text{mult}(i, j), pc, \mathcal{R}_{reg}, \text{Tag}_{\text{Patricia}}(S))$$

$$\text{Prove}(e, (pc, \text{add}(i, j), \mathcal{R}_{reg}, S)) := (\text{add}(i, j), pc, \mathcal{R}_{reg}, \text{Tag}_{\text{Patricia}}(S))$$

$$\text{Prove}(e, (pc, \text{env}(i), \mathcal{R}_{reg}, S)) := (\text{env}(i), pc, \mathcal{R}_{reg}, \text{Tag}_{\text{Patricia}}(S))$$

Verify function. The verify function follows the approach of computing the resulting tag from the proof directly. Then verifies that the proof corresponds to the current tag and that the new tag is equal to the one provided:

$$\begin{aligned} \text{Verify}(e, T, T', \pi) &:= T = T_{\text{before}} \wedge T' = T_{\text{after}} \wedge \text{Validate}(e, \pi) = 1 \\ &\text{where } T_{\text{after}} \leftarrow \text{TagAfter}(e, \pi), T_{\text{before}} \leftarrow \text{TagBefore}(e, \pi) \end{aligned}$$

With **TagBefore** extracting the ‘previous’ tag from the proof:

$$\text{TagBefore}(e, (\epsilon, pc, \mathcal{R}_{reg}, T, I, \pi_{\text{lookup}})) := \text{CRH}((T, pc, \epsilon, \mathcal{R}_{reg}))$$

$$\begin{aligned}
\text{TagBefore}(e, (\text{load}(i, j), pc, \mathcal{R}_{reg}, T, R, \pi_{lookup})) &:= \text{CRH}((T, pc, \text{load}(i, j), \mathcal{R}_{reg})) \\
\text{TagBefore}(e, (\text{store}(i, j), pc, \mathcal{R}_{reg}, T, T', \pi_{store})) &:= \text{CRH}((T, pc, \text{store}(i, j), \mathcal{R}_{reg})) \\
\text{TagBefore}(e, (\text{jump}(i, j), pc, \mathcal{R}_{reg}, T)) &:= \text{CRH}((T, pc, \text{jump}(i, j), \mathcal{R}_{reg})) \\
\text{TagBefore}(e, (\text{mult}(i, j), pc, \mathcal{R}_{reg}, T)) &:= \text{CRH}((T, pc, \text{mult}(i, j), \mathcal{R}_{reg})) \\
\text{TagBefore}(e, (\text{add}(i, j), pc, \mathcal{R}_{reg}, T)) &:= \text{CRH}((T, pc, \text{add}(i, j), \mathcal{R}_{reg})) \\
\text{TagBefore}(e, (\text{env}(i), pc, \mathcal{R}_{reg}, T)) &:= \text{CRH}((T, pc, \text{env}(i), \mathcal{R}_{reg}))
\end{aligned}$$

With `TagAfter` extracting the ‘resulting’ tag from the proof, by simulating the step function using the data provided in the proof string:

$$\begin{aligned}
\text{TagAfter}(e, (\epsilon, pc, \mathcal{R}_{reg}, T, I, \pi_{lookup})) &:= \text{CRH}((T, pc, I, \mathcal{R}_{reg})) \\
\text{TagAfter}(e, (\text{load}(i, j), pc, \mathcal{R}_{reg}, T, R, \pi_{lookup})) &:= \text{CRH}((T, pc + 1, \epsilon, \mathcal{R}_{reg}[r_i \leftarrow R])) \\
\text{TagAfter}(e, (\text{store}(i, j), pc, \mathcal{R}_{reg}, T, T', \pi_{store})) &:= \text{CRH}((T', pc + 1, \epsilon, \mathcal{R}_{reg})) \\
\text{TagAfter}(e, (\text{jump}(i, j), pc, \mathcal{R}_{reg}, T, \pi_{store})) &:= \\
&\quad \text{CRH}((T, \text{if } \mathcal{R}_{reg}[r_j] > 0 \text{ then } \mathcal{R}_{reg}[r_i] \text{ else } pc + 1, \epsilon, \mathcal{R}_{reg})) \\
\text{TagAfter}(e, (\text{mult}(i, j), pc, \mathcal{R}_{reg}, T, \pi_{store})) &:= \text{CRH}((T, pc + 1, \epsilon, \mathcal{R}_{reg}[r_i \leftarrow r_i \cdot r_j])) \\
\text{TagAfter}(e, (\text{add}(i, j), pc, \mathcal{R}_{reg}, T, \pi_{store})) &:= \text{CRH}((T, pc + 1, \epsilon, \mathcal{R}_{reg}[r_i \leftarrow r_i + r_j])) \\
\text{TagAfter}(e, (\text{env}(i), pc, \mathcal{R}_{reg}, T, \pi_{store})) &:= \text{CRH}((T, pc + 1, \epsilon, \mathcal{R}_{reg}[r_i \leftarrow e]))
\end{aligned}$$

With `Validate` : $\mathcal{E} \times \mathcal{P}_\kappa \rightarrow \{1, 0\}$ validating the memory operations by applying the verification of the authenticated data structure used to emulate a large memory space:

$$\begin{aligned}
\text{Validate}(e, (\epsilon, pc, \mathcal{R}_{reg}, T, I, \pi_{lookup})) &:= \text{Verify}_{\text{Patricia}}(T, T, \text{lookup}(pc), I, \pi_{lookup}) \\
\text{Validate}(e, (\text{load}(i, j), pc, \mathcal{R}_{reg}, T, R, \pi_{lookup})) &:= \\
&\quad \text{Verify}_{\text{Patricia}}(T, T, \text{lookup}(\mathcal{R}_{reg}[r_j]), R, \pi_{lookup}) \\
\text{Validate}(e, (\text{store}(i, j), pc, \mathcal{R}_{reg}, T, T', \pi_{store})) &:= \\
&\quad \text{Verify}_{\text{Patricia}}(T, T', \text{store}(\mathcal{R}_{reg}[r_i], \mathcal{R}_{reg}[r_j]), \epsilon, \pi_{store}) \\
\text{Validate}(e, *) &:= 1
\end{aligned}$$

8 Concrete Efficiency Considerations

In this section we cover a few simple optimizations which are of less theoretical interest, but can improve the concrete efficiency of FastSwap greatly. This section is aimed at potential implementors.

Reusing the judge. Often deploying the code of a smart contract has significant cost of its own. However, note that the functionality of the judge does not depend on the predicate, but only on the authenticated computation structure scheme. Hence the code can be reused between swaps or separated into a library which can be shared by multiple independent contracts.

High-level execution language. Rather than applying the **Step** function of the authenticated computation structure directly the prover and auditor can execute a more efficient higher level language where each instruction decomposes into a sequence of simpler low-level instructions from authenticated computation structure scheme. In case of a dispute the offending high-level instruction must be unpacked into its lower-level instructions and dispute resolution carried out at the lower layer. For instance this enables the use of hardware acceleration for cryptographic primitives in the high-level language while using a function call to a naive implementation in the low-level language.

‘Just-In-Time’ authenticated data structures. Rather than apply operations directly to the authenticated data structure used in the authenticated computation structure, concrete efficiency can often be gained by representing the data more efficiently during applications of the **Step** function and only ameliorate the data structure with the authentication data at states revealed during dispute resolution. An example of this is executing a higher level language, where each instruction corresponds to a long sequence of instructions on the **For** instance, when the memory space is represented as a Patricia trie, then when calling during **Prove** and **Tag** a Merkle tree is temporarily imposed over the data structure. This enables application of authenticated data structures which would otherwise inhibit concrete efficiency, e.g. RSA or CDH based vector commitments [5], which would only have to be computed over logarithmically many snapshots of the vector representing the memory space in case of dispute, rather than updated at every step of the computation during the honest execution.

Reduce computational complexity during dispute resolution. Rather than naively recomputing S_w from S_1 during dispute resolution, resulting in $n \log n$ computation steps, this can easily be reduced to n steps, by simply storing the state S_l corresponding to the left pointer and computing S_w from S_l whenever $l < w$ and from S_1 otherwise.

Efficient language. Language designers are likely to want a language close to the that of the underlying smart contract language in which the verifier is implemented. This is due to the verifier essentially being an interpreter for the source language, the size of which is directly proportional to the cost of deploying the judge contract. Additionally high-level instructions of the underlying smart contract language (like signature verification and cryptographic hash function evaluation) can be provided in the source language. Application of such high-level functions might greatly simplify the implementation of the decryption of the witness inside the predicate.

Send multiple tags during dispute resolution. The number of rounds during dispute resolution can be reduced by a constant $\log_2 c$, by having the prover send 2 tags T_{w_1}, \dots, T_{w_c} , then having the auditor send the index of the last match l and first mismatch r . For a computation of 2^{30} steps, letting $c = 2^5$, this reduces the number of interactions with the judge during dispute from 62 to 14.

Limit storage in the judge contract. The previous optimization introduces a significantly increased storage requirement on the judge (e.g. 32 hashes stored every iteration during dispute resolution). Some smart contract execution environments, in particular the Ethereum virtual machine, sets the price of storage very high (20000 ‘gas’ per 256 bits[6]⁷), compared to the price of memory (e.g. call arguments) or computation. In particular the cost of:

- Sending 32 words of 256 bits to the contract is ≈ 100 gas[6].
- Computing a Merkle tree over 32 words of 256 bits is ≈ 3000 gas[6]⁸.
- Storing 32 words is 640000 gas[6].

Hence it is significantly cheaper⁹ to have the judge compute a Merkle tree over the arguments (tags) and store the root. Then having the auditor prove a path to the (at most) two leafs which corresponds to updated l and r values. This is possible because every input to the judge, not only its current state, is public and therefore available to the auditor.

9 Further Research

9.1 Constructions of authenticated computation structures.

Unlike authenticated data structures where a proof must prove the correct execution of a full operation, the proofs for authenticated computation structures need only prove a single step of computation which can be arbitrarily small. In some cases this might enable significantly more efficient proofs than those for authenticated data structures under the same cryptographic assumptions:

In Section 7, we have described a concrete instantiation wherein the map lookup is a single instruction in the language. For our concrete instantiation this results in proofs of size $\log M$, with M being the size of the memory space. Alternatively low level operations for walking the authenticated data structure can be provided by the language and smaller atomic steps in the lookup can be proved instead. As a simple example consider lookups (*load* instructions) in the authenticated Patricia trie of Section 7, but where the state additionally contains a cryptographic hash digest for an ‘authenticated’ node inside the Patricia tree. Hence the proof becomes an instance of:

- An instruction pointer pc .
- An instruction I (which might be ϵ)
- A finite number of fixed-sized registers r_1, \dots, r_n .
- A tag for an authenticated Patricia tree T .
- A node pointer H_{node} .

⁷ Of which 15000 can be recouped by later clearing the memory.

⁸ Using 64 invocations of the **SHA3** instruction.

⁹ Our estimates for $c = 2^5$ is a 80 - 90 % ‘gas’ saving

Whenever $I \neq \text{load}(i, j)$, the verifier operates as in Section 7. Whenever $I = \text{load}(i, j)$ and $H_{\text{node}} = \epsilon$, the verifier checks that $H_{\text{node}} \leftarrow T$ in the subsequent tag. Whenever $I = \text{load}(i, j)$ and $H_{\text{node}} \neq \epsilon$, the proof additionally consists of a node in the Patricia tree, $\text{Node}(\text{prefix}, \text{len}, H_{\text{left}}, H_{\text{right}})$, and the verifier checks that $H_{\text{node}} = \text{CRH}(\text{Node}(\text{prefix}, \text{len}, H_{\text{left}}, H_{\text{right}}))$ and that $H_{\text{node}} \leftarrow H_{\text{left}}$ or $H_{\text{node}} \leftarrow H_{\text{right}}$ in the subsequent tag, depending on whether the lookup in the tree progresses left/right based on $r_j[\text{len}]$. When the leaf is reached, verify it similarly, set $I \leftarrow \epsilon$, set $H_{\text{node}} \leftarrow \epsilon$. For updates, where the new hash is propagated up through the tree, a similar process must be repeated in the opposite direction, then $T \leftarrow H_{\text{node}}$ at the leaf. Using this approach, the proof size can be made constant in M while the number of rounds during dispute grows by at most $\log \log M$ times.

References

1. Author, Stefan Dziembowski, Author, Lisa Eckey, Author, Sebastian Faust. Fair-Swap: How To Fairly Exchange Digital Goods. In Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, Series: CCS '18. ACM, New York (2018), <https://doi.org/10.1145/3243734.3243857>.
2. Author, Matteo Campanelli, Author, Rosario Gennaro, Author, Steven Goldfeder, Author, Luca Nizzardo. Zero-Knowledge Contingent Payments Revisited: Attacks and Payments for Services. <https://doi.org/10.1145/3133956.3134060>
3. Author, Henning Pagnia, Author, Felix C. Gartner. On the impossibility of fair exchange without a trusted third party. 1999.
4. BitcoinWiki, Zero Knowledge Contingent Payment, 2016, https://en.bitcoin.it/wiki/Zero_Knowledge_Contingent_Payment
5. Author, Dario Catalano, Author, Dario Fiore. Vector Commitments and Their Applications. Public-Key Cryptography – PKC 2013 https://doi.org/10.1007/978-3-642-36362-7_5
6. Author, Gavin Wood, Title, Ethereum: A secure decentralized generalized transaction ledger, 2018-08-16 (version e7515a3).
7. Author, Jesper Buus Nielsen. Separating Random Oracle Proofs from Complexity Theoretic Proofs: The Non-committing Encryption Case. Advances in Cryptology — CRYPTO 2002. https://doi.org/10.1007/3-540-45708-9_8
8. Author, Ronald Cramer, Author, Ivan Bjerre Damgård, Author, Jesper Buus Nielsen. Secure Multiparty Computation and Secret Sharing (1st edition). **ISBN-13: 978-1107043053**.
9. Author, Wacław Banasik, Author, Stefan Dziembowski, Author, Daniel Malinowski. Title, Efficient Zero-Knowledge Contingent Payments in Cryptocurrencies Without Scripts. Computer Security – ESORICS 2016. ESORICS 2016. Lecture Notes in Computer Science, vol 9879. Springer. https://doi.org/10.1007/978-3-319-45741-3_14
10. Author, Georg Fuchsbauer. Title, WI Is Not Enough: Zero-Knowledge Contingent (Service) Payments Revisited. CCS '17 Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, Pages 229-243. <https://doi.org/10.1145/3133956.3134060>