



MONASH University

Layer-2 Solutions for Bitcoin Scalability

by

Arash Mirzaei,

Thesis

Submitted by Arash Mirzaei

for fulfillment of the Requirements for the Degree of

Doctor of Philosophy (0190)

Supervisor: Dr. Amin Sakzad

Associate Supervisor: A/Prof. Ron Steinfeld, Dr. Jiangshan Yu

**Faculty of Information Technology
Monash University**

July, 2023

© Copyright

by

Arash Mirzaei

2023

Layer-2 Solutions for Bitcoin Scalability

Arash Mirzaei,

Monash University, 2023

Supervisor: Dr. Amin Sakzad

Associate Supervisor: A/Prof. Ron Steinfeld, Dr. Jiangshan Yu

Abstract

Payment channels are an effective layer-2 approach to improve the scalability of blockchain-based cryptocurrencies. A payment channel allows two parties to execute transactions off-chain. But since each party's share of coins in a channel changes over time, one party might attempt to close the channel with an old state to maximise her profit. In practice, when a channel party publishes a channel state on the blockchain, a period called *dispute period* starts, in which their counter-party can provide some evidence that proves the invalidity of the published state and hence prevents the channel from getting finalised with an old state. Several payment channels have been proposed for Bitcoin, each with its own limitations. Some issues with existing payment channels include: (1) *state duplication* meaning each party has its own version of transactions leading to complex transaction flows, (2) lack of a penalisation mechanism against dishonest parties and hence incentive incompatibility, (3) incompatibility of the payment channel with some important digital signature algorithms such as BLS which limits the payment channel usage in blockchains with such underlying digital signature algorithms, and (4) increase in the storage requirements of the channel parties over time.

Since the dispute period length is limited, channel parties must be always online and synced with the blockchain. Otherwise, they might fail to react to their counter-parties' misbehaviour in time. However, it is impossible for ordinary users to stay always online. Thus, parties might delegate this task to a third party, called the *watchtower*. Multiple watchtower schemes have been proposed in the literature, each with its own limitations; Some are unfair against channel parties or the watchtower (i.e. an honest channel party or a responsive watchtower might lose some funds in the channel); Some do not provide privacy preserving properties and some cannot be deployed on Bitcoin.

Therefore, there are some limitations in existing payment channels and their watchtower schemes, particularly in those that have been designed for blockchains with limited scripting languages. Thus, we focus on analysing the mentioned limitations and designing payment channels and watchtower schemes that outperform the existing ones. Hence, our contributions are as follows:

- We formalise watchtowers and their properties including *agility*, *privacy*, *fairness* and *coverage*. Furthermore, we prove some trade-offs between the above-mentioned properties. We also focus on different properties of a payment channel and examine the limitations of existing payment channels in achieving the required properties.
- We design a new watchtower scheme for Bitcoin, called FPPW, which is fair with respect to both channel parties and the watchtower. Therefore, the honest party's funds in the channel are safe and the honest watchtower's rewards are also guaranteed. Furthermore, FPPW provides privacy against the watchtower as the watchtower obtains no data on the way channel funds are redistributed among channel parties. Focusing on storage costs, we also design a second watchtower scheme for Bitcoin, called Garrison, where the storage costs of channel parties and their watchtower increase logarithmically with the number of channel updates. Additionally, Garrison avoids state duplication. Both FPPW and Garrison can be implemented without any update in Bitcoin blockchain.
- Relying on the deployment of the ANYPREVOUT signature type in Bitcoin, we also present a payment channel with an unlimited lifetime for Bitcoin, called Daric, that achieves optimal storage. Daric is the first payment channel that provides penalisation against the dishonest channel party and simultaneously avoids state duplication without relying on any particular property for the underlying digital signature. We also prove the security of Daric in the Universal Composability model.

Layer-2 Solutions for Bitcoin Scalability

Declaration

This thesis is an original work of my research and contains no material which has been accepted for the award of any other degree or diploma at any university or equivalent institution and that, to the best of my knowledge and belief, this thesis contains no material previously published or written by another person, except where due reference is made in the text of the thesis.

Arash Mirzaei

July 9, 2023

Publications

Published works (included in the thesis):

- Arash Mirzaei, Amin Sakzad, Jiangshan Yu, Ron Steinfeld. FPPW: A Fair and Privacy Preserving Watchtower For Bitcoin. In *International Conference on Financial Cryptography and Data Security*, pages 151-169. Springer, 2021.
- Arash Mirzaei, Amin Sakzad, Jiangshan Yu, Ron Steinfeld. Garrison: A Novel Watchtower Scheme for Bitcoin. In *Australasian Conference on Information Security and Privacy*, pages 489-508. Springer, 2022.
- Arash Mirzaei, Amin Sakzad, Jiangshan Yu, Ron Steinfeld. Daric: A Storage Efficient Payment Channel With Penalization Mechanism. In *International Conference on Information Security*, pages 229-249. Springer, 2022.

Other works during my PhD (not included in the thesis):

- Arash Mirzaei, Amin Sakzad, Ron Steinfeld, Jiangshan Yu. Algorand blockchain, in *Blockchains - A Handbook on Fundamentals, Platforms and Applications*, Lecture Notes in Computer Science (LNCS).
- Xianrui Qin, Shimin Pan, Arash Mirzaei, Oguzhan Ersoy, Amin Sakzad, Muhammed Esgin, Joseph K. Liu, Jiangshan Yu, Tsz Hon Yuen. BlindHub: Bitcoin-Compatible Privacy-Preserving Payment Channel Hubs Supporting Variable Amounts, In *IEEE Symposium on Security and Privacy*, 2023.

Acknowledgments

I would like to express my heartfelt gratitude to my main supervisor Amin Sakzad for all his endless support, guidance and encouragement in various aspects both on a professional and personal level. His support started even before the commencement of my PhD and went beyond my expectations throughout my research. I also thank my co-supervisors Ron Steinfeld and Jiangshan Yu for their help and support. Their expertise, insights, and valuable feedback helped me to achieve my academic goals. I am also grateful to my wife Lili for her unwavering support, patience, and love throughout this journey. Her understanding, encouragement, and sacrifice have been invaluable in enabling me to balance my academic pursuits with my personal life. I could not have achieved this without her support and I am forever thankful.

Moreover, I am thankful to my panel committee members Carsten Rudolph, Joseph Liu, Rafael Baiao Dowsley and Shifeng Sun for their constructive feedback. I thank the Faculty of Information Technology and the Graduate Research Office at Monash University for providing me the scholarship to cover the tuition fee and my cost of living during my PhD study. I also owe great thanks to my fellow PhD students Ahmad, Sara and Rumpa for their moral support.

Finally, I would like to take this opportunity to acknowledge the bravery and resilience of Iranian men and women who have fought for their rights and freedom. During my PhD study, I was moved by the story of Mahsa Amini, a young Iranian woman who was killed while in police custody. Her tragic death, along with the deaths of others including kids who were killed during the “*Women Life Freedom*” protests, the blinding of many individuals, and the imprisonment of countless others, is a stark reminder of the courage and determination of their unwavering commitment to the principles of human rights, democracy, and gender equality. May their sacrifice and courage never be forgotten, and may their struggle continue to inspire and guide us in the fight for justice, freedom, and equality.

Arash Mirzaei

Monash University
July 2023

Contents

Abstract	iii
Acknowledgments	vii
List of Figures	xii
List of Tables	xiii
1 Introduction	1
1.1 Contributions	4
1.2 Thesis Structure	5
2 Background and Literature Review	7
2.1 Background	7
2.2 Payment Channel	9
2.3 Watchtower	11
2.4 Channel Synchronisation	13
3 Preliminaries and Notations	17
3.1 Preliminaries	17
3.1.1 Digital Signature	17
3.1.2 Hard relation	18
3.1.3 Adaptor Signature	18
3.2 Notations	19

4	Formal Treatment of Watchtower	23
4.1	Introduction	23
4.2	Formalisation of Payment Channel and Watchtower	23
4.3	Watchtower Service Properties	27
4.3.1	Agility	27
4.3.2	Privacy	28
4.3.3	Fairness and Coverage	31
4.4	Conclusion	33
5	FPPW: a fair and privacy preserving Bitcoin watchtower	35
5.1	Introduction	35
5.3	FPPW Overview	37
5.3.1	System Model	37
5.3.2	Overview	37
5.3.2.1	NVG: A New Variant of the Generalized Channel	38
5.3.2.2	Adding a Watchtower Service with Fairness w.r.t. the Hiring Party to NVG	39
5.3.2.3	Allowing Watchtower to Terminate its Service	40
5.4	FPPW Protocol Description	41
5.4.1	Create	42
5.4.2	Update	45
5.4.3	Close	48
5.4.4	React	49
5.4.5	Watchtower Terminate	50
5.5	Security Analysis	50
5.6	Fee Handling	57
5.7	FPPW Protocol	59
5.8	Temporarily unavailable watchtower	65
5.9	FPPW Transactions Scripts	70
5.10	FPPW with One Hiring Party	73
5.11	Conclusion	73

6	Garrison: a storage efficient Bitcoin watchtower	75
6.1	Introduction	75
6.2	Notations	76
6.3	Garrison Overview	77
6.3.1	System Model	77
6.3.2	Garrison Overview	77
6.3.2.1	Reducing the Storage Requirements of the Watchtower	77
6.3.2.2	Reducing the Storage Requirements of channel parties	80
6.4	Garrison Protocol Description	81
6.4.1	Create	81
6.4.2	Update	84
6.4.3	Close	84
6.4.4	Punish	85
6.5	Security Analysis	86
6.6	Garrison Transactions Scripts	89
6.7	Garrison Protocol	91
6.8	Conclusion	95
7	Daric: a storage efficient channel with penalisation	97
7.1	Introduction	97
7.2	Notations and Background	99
7.2.1	Notations	99
7.2.2	Background	100
7.2.2.1	Floating Transactions	100
7.2.2.2	eltoo [1]	100
7.3	Daric Overview	100
7.3.1	Revocation Per State	101
7.3.2	Revocation Per Channel	101
7.3.3	Avoiding State Duplication	102
7.3.4	State Ordering	102

7.3.5 Putting Pieces Together	103
7.4 Daric Protocol Description	103
7.4.1 Create	103
7.4.2 Update	106
7.4.3 Close	108
7.4.4 Punish	109
7.5 Security Analysis Overview	109
7.5.1 Notation and Security Model	109
7.5.2 Ideal Functionality Properties	110
7.6 Daric Versus Eltoo	111
7.6.1 HTLC Security	111
7.6.2 Punishment Mechanism	112
7.7 Performance Analysis	114
7.8 Daric Transactions Scripts	115
7.9 UC Framework	117
7.10 Ideal Functionality	119
7.10.1 Functionality Wrapper	123
7.11 Daric Protocol	124
7.11.1 Protocol Wrapper	130
7.12 Security Analysis	132
7.13 Discussions	146
7.14 Conclusion and Future Work	147
8 Conclusion and future work	149
Conclusion	149
Future Work	150
References	153

List of Figures

2.1	Lifetime of a payment channel.	9
3.1	A sample transaction flow.	20
5.1	NVG Channel Transactions Flow	39
5.2	Adding a Fair Watchtower to NVG	40
5.3	FPPW Channel Transaction flow	41
5.4	A summary of FPPW Channel Create.	46
5.5	FPPW Channel Update.	49
5.6	An FPPW Bitcoin Channel with only A Being the Hiring Party.	74
6.1	A sample transaction flow.	77
6.2	Reducing the Storage Requirements of the Watchtower	78
6.3	Adding Y and R Values to Commit Transactions	79
6.4	Reducing Storage Requirements of Channel Parties	81
6.5	Summary of Garrison channel creation phase.	82
6.6	Summary of Garrison channel update phase from state i to $i + 1$	85
7.1	A sample transaction flow.	101
7.2	Transaction flows for a Lightning channel with punish-then-split mechanism.	102
7.3	Transaction flows for state i of a Daric channel.	103

List of Tables

3.1	Summary of notations	21
4.1	Summary of the Condition $\mathcal{C}_S$ for Different Watchtower Schemes.	26
4.2	Summary of the Condition $\mathcal{C}_U$ for Different Watchtower Schemes.	26
4.3	Summary of the Condition $\mathcal{C}_T$ for Different Watchtower Schemes.	27
4.4	Comparison of Different Watchtower Schemes.	34
5.1	Different Properties of the FPPW Scheme.	36
5.2	Different Phases of an FPPW Channel	42
6.1	Comparison of different dispute period-based payment channels with n channel updates, m HTLC outputs on average per state and k channel splits on top of each other.	76
6.2	Different Properties of the Garrison Scheme.	96
7.1	Comparison of different payment channels with n channel updates and k recursive channel splitting.	99
7.2	On-chain cost of different closure scenarios and the number of operations performed by each party for a channel update for different payment channels with m HTLC outputs ($0 \leq m \leq 966$).	116

Chapter 1

Introduction

Payment Channel: Bitcoin [2], the pioneering cryptocurrency, has emerged as a significant innovation with far-reaching implications. Its decentralized nature, secure transactions, and potential to reshape the financial landscape highlight its importance as a transformative force in the digital era. Scalability is of paramount importance in the context of Bitcoin as it directly affects its ability to function as a global, mainstream currency. With the growing adoption and transactional demands, ensuring a scalable infrastructure becomes crucial for Bitcoin to maintain efficiency, low fees, and widespread usability, ultimately shaping its potential for broader acceptance and impact.

However, Bitcoin scalability is still a grand open challenge limiting the adoption of blockchain technologies [3, 4, 5]. For instance, Bitcoin can only process about 10 transactions per second on average [6, 7], which is much lower than the figures for centralised payment systems such as Visa [8]. This issue has led to several solutions: (1) modifications in consensus protocols [9, 10], (2) sharding [11, 12], (3) side-chains [13], and (4) layer-2 or off-chain protocols [14, 4, 15], where this thesis targets the last solution.

Layer-2 protocols are so named because they operate on top of the base layer of the Bitcoin blockchain. These solutions aim to alleviate scalability issues by conducting off-chain transactions, reducing the burden on the main blockchain. By implementing layer-2 solutions, Bitcoin can achieve faster and more cost-effective transactions, increased capacity, and improved scalability while maintaining the security and decentralization features of the underlying blockchain. Notably, layer-2 protocols do not require any modifications in the blockchain and its consensus mechanism. Layer-2 protocols are deployed using the scripting language of the underlying blockchain and rely on two main properties of the first layer blockchain [16]: *integrity* and *eventual synchronicity with an upper time-bound*, where the former means only valid transactions can be published on

the blockchain and the latter implies every submitted valid transaction is finally added to the blockchain before a limited timeout.

There are various types of off-chain solutions: (1) Payment channel [15, 4, 14] allow participants to create a secure off-chain channel for conducting multiple transactions without each transaction being recorded on the main blockchain (2) State channels [17] expand upon the idea of payment channels to encompass the execution of arbitrary applications, extending their functionality beyond simple payment transactions, and (3) Commit-chains [18, 19] allow transactions to be facilitated through the involvement of a centralized yet untrusted intermediary, enabling communication between participating parties. Given their simplicity, widespread adoption and significance in addressing scalability challenges, payment channels take center stage in this thesis.

A payment channel between two parties Alice (or A) and Bob (or B) allows them to perform a number of transactions without publishing every single transaction on the blockchain. To create a channel, A and B respectively deposit a and b coins into a joint address that is controlled by both parties. Parties can privately update their balance in the channel by exchanging off-chain transactions and agreeing on a new distribution of channel funds. Each party can close the channel at any time by enforcing the latest channel state on the blockchain. Payment channels can also be linked to form a Payment Channel Network (PCN), where payments between users with no direct payment channel can be routed via intermediaries.

Since the channel parties are generally untrustworthy and blockchain miners are unaware of the off-chain transactions, a mechanism must be adopted to prevent potentially cheating parties from publishing an old state. To achieve that, Lightning Network [15], as the most popular payment channel network, adopts a punishment mechanism to prevent parties from acting dishonestly. So upon authorising a new state, channel parties exchange some revocation secrets to revoke the previous state. Then, if a party publishes a revoked state, her counterparty who is supposed to be always online uses the corresponding revocation secrets and reacts within a limited time period, called *dispute period*, to take all the channel funds.

Although elegantly designed, the Lightning Network has some shortcomings. Firstly, since channel parties must store all the revocation secrets, received from their counterparties, their storage amount increases with the number of channel updates. Moreover, to detect and punish the misbehaving party, the channel state is duplicated meaning each party has its own copy of the state. To solve the former issue, eltoo [1] removes the punishment mechanism causing incentive incompatibility. To solve the second issue, Generalized channel [14] uses a dedicated design of adaptor signatures which introduces compatibility issues with BLS [20] or most post-quantum digital signatures.

Watchtower The dispute process works based on the assumption that the parties are always online and synced with the blockchain to detect malicious behaviour. This requirement has always been a drawback for payment channels because it can be practically violated due to crash failures or by performing DoS attacks on channel parties. *Watchtower* was introduced to relax this strong assumption by allowing users to delegate the watching tasks to watchtower services [15]. Watchtowers are always-online services that monitor the blockchain and act on behalf of their clients to secure their funds.

Monitor [21] is the first watchtower scheme for the Lightning Network which mainly focuses on channel privacy against watchtower. However, Monitor has two main issues, both of which are related to *fairness*. Firstly, honest watchtowers might be rewarded for fraud (i.e. broadcast of an old state on-chain), which is unfair with respect to (w.r.t.) the watchtower. Secondly, honest parties cannot penalise the unresponsive watchtower, which is unfair towards an honest hiring party. Moreover, the storage costs of the watchtower in Monitor increase linearly with each channel update. Outpost [22] solves the issue of fairness towards the watchtower by paying her per channel update. It also improves the storage requirements of the watchtower.

Cerberus [23] and PISA [24] elegantly provide fairness w.r.t. the hiring party by enforcing the watchtower to lock some collateral that is given to their clients given that the watchtower is unresponsive upon dispute. However, PISA fails to be deployed on Bitcoin and Cerberus sacrifices privacy. In particular, the Cerberus watchtower learns the distribution of funds in the channel. Moreover, since watchtowers in PISA and Cerberus have to lock some collateral per channel, their *coverage* or equivalently their capability in watching all the existing payment channels on a fixed Blockchain is limited. Therefore, each watchtower scheme focuses on some particular properties among privacy, storage requirements of the watchtower or fairness. However, none of the existing Bitcoin-compatible watchtower schemes can achieve all the desired features.

Therefore, the main objective of this thesis is the analysis of different limitations of existing payment channels and their watchtowers and then moving towards improving different aspects of the existing schemes. Accordingly, this thesis answers the following research questions (RQ):

RQ 1 How to formally define different properties of a watchtower scheme? Why do existing watchtower schemes fail to meet all the required properties?

RQ 2 How to design a watchtower scheme for existing payment channels that mitigates the limitations of existing schemes?

RQ 3 What are other limitations of the existing payment channels? How to design a payment channel solution that mitigates the limitations of existing payment channel proposals? How to analyse the security of the proposed payment channel?

1.1 Contributions

In this thesis, we analyse the existing payment channel and watchtower schemes. Then, focusing on the shortcomings of the current schemes, we move towards design of new payment channel and watchtower schemes with better properties. In more detail, the contributions of this thesis are as follows:

- **Formal Definition of a Watchtower and its Required Properties:** For RQ 1, we formalise the definition of a watchtower and its different properties (See Chapter 4). Those properties are as follows:
 - Agility: flexibility of a watchtower to start and terminate its service,
 - Privacy: the amount of knowledge the watchtower (or respectively any third party) obtains about the payment channel (or respectively about the hiring status of the watchtower).
 - Fairness: the level of guarantee that the watchtower (or respectively the hiring party) provides to the hiring party (or respectively to the watchtower) on its service (or respectively on its payment), and
 - Coverage: the capability of a watchtower (on a scale between 0 to 1) in watching all the existing payment channels on a fixed Blockchain.

We also evaluate the existing schemes regarding these properties and show that existing proposals fail to provide acceptable results in all aspects. Furthermore, we prove that there is a trade-off between the level of fairness that a watchtower provides to its clients and the coverage it can achieve.

- **Design and Analysis of Two New Watchtower Schemes:** For RQ 2, we design two watchtower schemes: FPPW (See Chapter 5) and Garrison (See Chapter 6). FPPW focuses on the properties we introduced earlier and provides fairness w.r.t. all channel participants including both channel parties and the watchtower. It means that the funds of any honest channel participant are safe even assuming that the other two channel participants are corrupted and/or collude with each other. Furthermore, the watchtower in FPPW learns no information about the off-chain transactions and hence the channel privacy is preserved. Furthermore, we

show that FPPW's coverage is higher than other existing schemes. FPPW can be implemented without any update in the Bitcoin script.

While FPPW offers desirable properties, it is essential to consider that the storage costs for channel parties and watchtowers increase linearly with each channel update. Thus, we design Garrison to achieve lower storage costs. Notably, the storage cost for both channel parties and their watchtower in Garrison would be $\mathcal{O}(\log N)$ with N being the number of channel updates. Furthermore, using properties of the adaptor signature, Garrison avoids state duplication. It means both parties store the same version of transactions for each state and hence the number of off-chain transactions does not exponentially increase with the number of applications built on top of each other in the channel. Moreover, independent of the complexity of the published revoked state, the honest party or his watchtower in a Garrison channel publishes a single transaction to react upon dispute. Garrison can also be implemented without any update in the Bitcoin script.

- **Design and Analysis of a New Payment Channel Scheme:** For RQ 3, we discuss the required properties of a payment channel and analyse the limitations of existing payment channels in achieving the required properties, i.e. unlimited lifetime, fixed storage costs, punishment mechanism, compatibility with any digital signature algorithms, state duplication avoidance and *bounded closure* where the latter enables channel parties to close the channel within a bounded time. Then, we introduce Daric (See Chapter 7), a payment channel with an unlimited lifetime for Bitcoin that achieves optimal storage (constant storage for both channel parties and their watchtower) and bounded closure. Moreover, Daric implements a punishment mechanism and simultaneously avoids the methods other schemes commonly used to enable punishment: (1) state duplication which leads to an exponential increase in the number of transactions with the number of applications on top of each other or (2) dedicated design of adaptor signatures which introduces compatibility issues with BLS or most post-quantum digital signatures. We also formalise Daric and prove its security in the Universal Composability model.

1.2 Thesis Structure

In Chapter 2, we review the existing payment channel and watchtower schemes. Chapter 3 presents the background, preliminaries and notations. Chapter 4 formalises the watchtower and its different properties. Two proposed watchtower schemes, i.e. FPPW and Garrison, will be presented in Chapter 5 and Chapter 6, respectively. Chapter 7 presents our proposed payment channel scheme, i.e. Daric. We conclude this thesis and discuss the potential future works in Chapter 8.

Chapter 2

Background and Literature Review

2.1 Background

Bitcoin [2], as the first cryptocurrency, is a collection of concepts and technologies that can be used to make payments without relying on any central authorities like banks or financial institutes. Bitcoin users who are connected over the Internet can transfer units of currency, called bitcoin, in exchange for some goods or services. However, unlike traditional payment systems, bitcoins are fully virtual. It means that there are not any physical coins to be transferred between the sender and the recipient. So units of value in the Bitcoin system are transferred among users by exchanging some digital transactions that transfer bitcoin values from the sender to the recipient. To prevent users from spending each other's coins, each coin in this system corresponds with a key pair of a public key cryptosystem whose private key is only possessed by the coin owner. So, to create a transaction and send some coins to a new owner, the current owner needs to prove the ownership of his coins by signing the transaction using his corresponding private key. This transaction transfers the coins to a new owner with his own key pair whose public key is recorded in the transaction.

Bitcoin is a fully distributed, peer-to-peer system. Thus, there is not any central authority to validate and store the users' transactions. Therefore, the important question is what can prevent a user from double spending, i.e. sending the same coins more than once. For fiat currencies, notes or coins carry some special features that can only be made by a central authority. So counterfeited money would be detectable. Also, when we use our mobile banking app to transfer money over a digital medium, one or multiple central authorities like banks record the corresponding transactions. So they prevent double spending. However, due to its decentralised nature, Bitcoin cannot use these mechanisms to avert double spending. As will be explained in the following, blockchain technology is the solution to this issue [2].

The Bitcoin blockchain [2] is an immutable public distributed ledger that records Bitcoin transactions. Let's explain these terms further. Each block is a limited-size list of Bitcoin transactions that are added to the blockchain data structure one at a time. The term public distributed ledger refers to the fact that rather than storing in a central database, blockchain data is distributed to a large number of independent nodes (also known as miners) and everyone is also capable to join the node community. Transactions in the blockchain are immutable meaning that once a transaction is added to the blockchain, it is computationally impossible to change its value or remove it from the ledger unless the block containing the transaction and all subsequent blocks are altered. This property is achieved as each block, along with its list of transactions, contains a cryptographic hash of the previous block. In this way, each block in the blockchain is chained to all previous blocks up to the first block in the blockchain.

Now that the ledger in the Bitcoin blockchain is distributed to many nodes, these nodes need a way to reach an agreement on the data that is going to be added to the ledger. In more detail, Bitcoin users constantly send their transactions to a publicly available peer-to-peer network. Then, miners validate these transactions and list them to form a block. However, due to the fact that these nodes are geographically distributed, they receive these transactions with different orders. Sometimes, due to network issues, some transactions might not be received by some blockchain nodes. Thus, their view of the submitted transactions and the latest created block would be different from each other. Therefore, another important building block in blockchain technology is how blockchain nodes reach a consensus. The consensus mechanism that is used in Bitcoin is called Proof-of-Work (PoW) [2, 25]. In this method, after creating the block, each node solves a computation-heavy cryptographic puzzle. The first node that succeeds in solving this puzzle would win and add the corresponding block to the blockchain. In order to incentivise the nodes to perform such heavy computations, the winner is also rewarded with some newly generated Bitcoins as well as some transaction fees.

The Bitcoin blockchain cryptographic puzzle is set such that each block is added to the blockchain every 10 minutes. Due to the fact that the size of each block is also limited to 1 MB, Bitcoin is significantly slower than centralised payment systems such as Visa [15]. A payment channel [15, 26, 4] is a promising solution to this issue where two channel parties open a channel on the blockchain by publishing a transaction on the blockchain through which they send their coins to a joint account that needs signatures by both parties to be spent. Then, they exchange off-chain transactions outside of the Bitcoin blockchain where those transactions send the coins in the joint account to each party. Each party can submit the latest channel transaction to the Bitcoin network to close the channel. Channel parties can also use their channels to create a network of payment channels, called Payment Channel Network (PCN), to perform indirect payments

through intermediaries [15, 14]. In practice, some routing algorithms are required to find the appropriate intermediaries for this purpose. In more detail, the indirect payment would be feasible if all channels along the route have enough balances in the desired direction. Otherwise, the payment fails. Thus, the routing algorithms intend to find the best existing paths from the payer to the payee for a given payment.

2.2 Payment Channel

A payment channel is a state machine run by two participants and its lifetime consists of three main phases including *create*, *update* and *close* (See Fig. 2.1). In the channel create phase, parties lock their funds in a joint account whose value can only be spent upon both parties' agreement. In the channel update phase, parties agree on a new channel state. Finally, in the channel close phase, parties close the channel by recording the latest channel state on-chain. Different payment channel proposals use different techniques to replace the channel state. In the following, we will explain these techniques [16].




Figure 2.1: Lifetime of a payment channel.

- *Replace by Incentive (RbI)*. This mechanism was proposed by Hearn and Spilman [26] and led to the first payment channel. This channel is funded by one of the channel parties who is the payer in the channel. Then, for each channel update, a partially signed transaction is given from the payer to the payee. The channel can be closed before a predetermined deadline by the payee who is incentivised to do that using the latest channel state because it gives the payee the highest deserved value. This type of channel has two limitations: (1) The lifetime of the channel must be determined at the time the channel is created and cannot be extended, and (2) The channel is unidirectional meaning that payment can only be done in one direction.

- *Replace by Time-Lock (RbT)*. This mechanism allows both channel parties to pay each other [4]. Each agreed-on state in such a channel has an absolute time-lock whose value decreases with each channel update. Thus, the latest state can be published on the blockchain earlier than all other states. But due to its decrementing time-lock, the number of channel updates is limited.
- *Replace by Revocation (RbR)*. If upon authorising a new state, the previous state is revoked, the payment channel is categorised as an RbR type. To revoke the previous state, parties exchange some revocation secrets that can be used for fraud. In other words, if a party tries to close the channel with a revoked state, her counterparty uses the corresponding revocation secret and publishes a transaction, called *revocation* transaction and penalises the dishonest party by taking all her funds in the channel. This idea was first proposed by Poon and Dryja in [15] for Bitcoin and has been actively used in Lightning Network with a current network capacity of around 5,000 BTC (as of November 2022).

Although elegantly designed, the Lightning channel has some shortcomings. Firstly, since all the revocation secrets must be stored by channel parties, their storage requirements increase with each channel update. Secondly, each party has its own version of the channel state. This *state duplication* is required to distinguish a dishonest party from her counterparty. To add an application (e.g. *Virtual channel* [27]) on top of the channel, parties have to split their channel into sub-channels. The state of each sub-channel is duplicated and it must propagate on both duplicates of the parent channel. This causes the number of transactions to exponentially rise with the number of applications built on top of each other [14].

Recently, Aumayr et al. [14] proposed a new design called *Generalized channel* which uses *adaptor signatures* to distinguish the publisher of a revoked state from her counterparty, i.e. once the revoked state is published, a *publishing secret* is revealed that can be used to penalise the dishonest party. In this way, state duplication is avoided. Towards a different direction, Aumayr et al. [28] proposed a channel, called “sleepy channel”, that allows channel parties to go offline for prolonged periods.

- *Replace by Version (RbV)*. Another method for replacing an old state is the usage of a monotonic counter that represents the version of each state. Channel parties store the channel state with the highest version and use it upon fraud. This idea has been used in several proposals on Turing complete blockchains (e.g. Ethereum) [29, 30, 17]. To extend this idea to Bitcoin, Decker et al. [1] introduced ANYPREVOUT as a new signature type for Bitcoin whose deployment requires a soft fork in the Bitcoin protocol [31]. This new signature type allows Bitcoin to support *floating*

transactions, i.e. transactions that can spend any Unspent Transaction Output (UTXO) with matching scripts. The proposal eltoo [1] uses this idea to deploy the concept of version.

Channels of type RbI and RbT are closed once the latest state is broadcast on-chain. However, for RbR and RbV channels, once a channel state is published by a channel party, a period, called the *dispute period*, starts. If the published state is invalid, the honest party proves its invalidity within the dispute period. For RbR channels, the invalidity of the published state is proved by publishing another transaction called the *revocation transaction* which can only be done by the honest party who knows the value of the corresponding revocation secret. For RbV channels, invalidity proof is done by broadcasting a state with a higher version.

Since the length of the dispute period is limited, RbR and RbV channels are based on the assumption that channel parties are always online. However, since this strong requirement cannot be met by most ordinary users, this task is delegated to a third-party service provider called the *watchtower*.

2.3 Watchtower

Monitor [21] is a watchtower scheme for Lightning Network. In this scheme, the hiring party provides the watchtower with IDs of revoked transactions as well as their corresponding revocation transactions. Then, the watchtower is supposed to look for a transaction with the matching ID on the blockchain. When a match is found, the watchtower can immediately publish its corresponding revocation transaction. To improve the efficiency, just some parameters such as addresses, signatures, and other metadata which are required for constructing the revocation transaction can be given to the watchtower [21]. Then, a full revocation transaction can be built by the watchtower if it is necessary.

To improve privacy, the hiring channel party encrypts the revocation transaction (or its corresponding parameters) using the second half of the revoked transaction ID and provides the watchtower with the first half of the ID along with the encrypted version of the revocation transaction. Then, the watchtower monitors the blockchain looking for every transaction whose ID matches the IDs received from the hiring party. If the watchtower finds a match, it decrypts its corresponding encrypted transaction using the second half of the found transaction ID and if the result is a valid transaction, transmits it on the blockchain. To reward the watchtower, the hiring party pays the watchtower through the revocation transactions, meaning that the watchtower is paid given that a revoked transaction is broadcast and its corresponding revocation transaction is published by the watchtower. To reduce the storage costs and hence operational costs of

the watchtower, Khabbazian et al. [22] proposed Outpost which proposes storing each revocation transaction as part of its corresponding channel state. Therefore, the watchtower just needs to watch for specific revoked transaction IDs on the blockchain. Then, the watchtower extracts the signed revocation transaction directly from the published revoked channel state that has appeared on the blockchain.

One issue with Monitor and Outpost is that the watchtower's client cannot completely trust that the watchtower remains online or does not collude with their counterparties. To address this concern specifically for Turing complete blockchains, DCWC and for Bitcoins, DCWC* [32] propose the implementation of a network of watchtowers that are financially motivated to faithfully cooperate with each other. This incentivised cooperation increases the probability that at least one honest watchtower prevents the channel from getting finalised with an old state. However, even for these schemes, channel parties cannot still be completely certain that they do not lose any money in their channels. The reason is that watchtowers are unaccountable or unfair to their clients. In other words, watchtowers are not penalised if they crash or deliberately do not act upon malicious behaviour. To encourage watchtowers to well-behave, [33] proposes a reputation system that works based on a cryptographic hash-based proof-of-work algorithm [34], called hashcash. This algorithm was originally designed to mitigate the email spamming issue.

Fail-safe [35] is a watchtower scheme for Turing complete blockchains and tries to resolve the issue of watchtower failure. A Fail-safe watchtower verifies all the off-chain channel updates and stores the latest state. Whenever a channel party, let's say *A*, attempts to close the channel, the watchtower can immediately reject or validate the committed state and receives its reward. However, if for any reason the watchtower is offline, a long timeout is triggered. Party *B* is supposed to get online before this long timeout is expired. Then, if the published state is an old one, *B* has this opportunity to prove the invalidity of the published state.

Another issue with Monitor, DCWC and DCWC* is that the watchtower is paid only if it observes a revoked channel state on the blockchain and broadcasts its corresponding revocation transaction. However, given that both parties of a given channel act honestly, the honest watchtower does not receive any reward although it has consumed some resources to remain online and monitor that channel. Such a reward mechanism is unfair to the watchtower and might disincentivise entities to run such services.

One simple solution to this issue (lack of fairness with respect to the watchtower), proposed by Outpost [22], is unconditional payment to the watchtower upon each channel update. However, this reward mechanism somehow exacerbates the consequences of the first mentioned problem (lack of fairness with respect to the hiring party) as it is probable for a channel party not only to lose some funds in the channel but also to pay

an unconditional reward to one or more watchtowers which actually did not prevent from fraudulent channel closure.

Therefore, having a watchtower scheme with a fair reward mechanism that also guarantees the safety of its channel parties' funds in the channel is desired. This is the main motivation behind the design of PISA which is a watchtower scheme proposed by McCorry et al. [24]. The watchtower's client in PISA receives a signed receipt from the watchtower using which they can prove the watchtower's wrongdoing. In such situations, the large security deposit of the watchtower is forfeited. The watchtower is also paid per channel update.

However, PISA cannot be used for more limited scripting languages, and hence cannot be utilised for Bitcoin payment channels. Avarikioti et al. [23] proposed Cerberus which is a fair watchtower scheme for Bitcoin. In particular, the watchtower in Cerberus locks some collateral per channel that is taken by the channel party given that the watchtower is unresponsive upon fraud. However, it is not privacy-preserving as the watchtower learns how channel funds are redistributed amongst channel parties at each channel update.

Furthermore, there are two other schemes whose security assumptions are different from the ones mentioned earlier. Unlike other watchtower schemes, TEE Guard [36] relies on features of Trusted Execution Environments to build watchtowers. TEE Guard can be deployed for Lightning Network and its storage costs are constant per channel. Towards a different direction, Brick [37] is actually a state channel construction in which the dispute period is replaced with a committee of $n = 3f + 1$ members with at most f byzantine members where signatures of $t = 2f + 1$ members of the committee are required for each channel update. Each committee member in Brick locks some funds as collateral and cheating committee members are penalised by losing their collateral. Also, unlike other watchtower schemes, if one of the channel parties wants to unilaterally close the channel, the involvement of $t = 2f + 1$ committee members is required. This mitigates the synchrony requirement of other schemes which work based on a dispute period. In other words, watchtowers are no longer necessary in Brick.

2.4 Channel Synchronisation

To establish a payment channel, parties must lock some funds in the channel. If a user wants to perform transactions with many other users, a naive idea would be to establish payment channels with all of them. This idea clearly does not scale well. Thus, in practice payment channels are linked to form a Payment Channel Network (PCN) where each payment can be routed via multiple intermediaries [15]. Thus, if Alice wants to

send some funds to David whom does not share any channels, Alice finds a path to David where intermediate nodes (Bob and then Charlie) relay the payment from Alice to David. As an alternative, a Payment Channel Hub (PCH) [38, 39, 29, 30, 40] deploys a star topology where users can pay each other via a single intermediary (called the *tumbler*).

Hashed Time-Lock Contracts (HTLC) [41], is the main method to synchronise payments in different channels in PCN or PCH deployments. Now the most prominent PCNs, i.e. Lightning Network [15] on Bitcoin [2] and Raiden network [42] on Ethereum [43] are working based on HTLC. HTLC in a transaction output locks x coins in a contract with two parameters: a timeout t and a hash value $y := \mathcal{H}(R)$ where $\mathcal{H}$ is a collision-resistant hash function and R is a randomly selected value. The HTLC contract's condition can be fulfilled as follows: (1) If Bob provides a pre-image R^* with $y = \mathcal{H}(R^*)$ before time t , then Alice pays Bob x coins, (2) Otherwise, Alice is refunded with x coins after time t . Now assume that Alice wants to pay David via Bob and Charlie who agree to forward this payment in exchange for f coins as the fee. Then, Alice conditionally pays Bob $x + 2f$ coins with $3t$ and y as HTLC parameters. Similarly, Bob conditionally pays Charlie $x + f$ coins with $2t$ and y as HTLC parameters. Finally, Charlie conditionally pays David x coins with t and y as HTLC parameters. Now if David provides R before time t , he is paid x coins by Charlie. Having the value of R , Charlie and then Bob are also paid $x + f$ and $x + 2f$ coins by Bob and Alice, respectively. If R is not provided by David, then Charlie, Bob and finally Alice get back their coins after time t , $2t$ and $3t$, respectively.

The timeout value in HTLC mechanism is linear to the length of the path. Sprites [17] focuses on making the timeout value independent of the path length. However, it can only be deployed in Turing complete blockchains such as Ethereum. AMCU [44] resolves the same issue but in a way that is deployable on Bitcoin. Additionally, it extends the functionality by enabling concurrent payments in channels that are not necessarily along a path from the sender to the receiver. This could be useful for some applications such as crowdfunding. Payment Trees [45] presents an attack against AMCU and also provides a secure solution similar to Sprites but compatible to Bitcoin.

HTLC mechanism also suffers from two other problems: (1) The hash value y can be used by an adversary to compromise privacy by finding out who is paying to whom and (2) Malicious intermediaries might steal the fees from honest intermediaries in the same payment path [46]. [47] proposes a Multi-hop Hash Time-Lock Contract (MHTLC) protocol to improve HTLC privacy. However, its proposed protocol is quite expensive in terms of computation and communication. [48] proposes a less complex solution, called Chameleon Hash Time-Lock Contract (CHTLC) to solve the same privacy problem. However, it is incompatible with Bitcoin. AMHL [46] resolves both issues by replacing the HTLC contract with a novel cryptographic lock. However, since this novel idea is

incompatible with some important digital signatures such as BLS [20], [49] presents a new primitive called *lockable signature* which is compatible with any digital signature algorithm.

HTLC, MHTLC, CHTLC, AMHL and lockable signature follow a 2-phase-commit paradigm where in the first phase, HTLC contracts are set up one by one from the payer to the payee (i.e. from Alice to David in the previously mentioned scenario). Then in the second scenario, locks are released by passing the pre-image value (i.e. R) from the payee to the payer via intermediaries before HTLC time-locks are expired. Blitz [50] provides a 1-phase protocol with provable security for multi-hub payments.

In all previously introduced solutions, all intermediate nodes along the path between the payer and the payee should be actively involved in each and every single payment. This makes the whole PCN network less reliable as offline users cannot contribute to payments. Also, this active involvement adds to each intermediary's service fee. To mitigate this issue, [29] provides a novel solution for Ethereum called virtual channels. A virtual channel is like a bridge between two users who do not share a direct payment channel. Intermediaries become involved in the process of virtual channel creation but once the virtual channel is created, end users can perform arbitrarily many off-chain transactions without requiring any involvement by intermediaries. Aumayr et al. [51] also extended this idea to a Bitcoin-compatible virtual channel.

Chapter 3

Preliminaries and Notations

3.1 Preliminaries

This section introduces the different cryptographic primitives used in this thesis. For all the cryptographic primitives defined in this section, κ is the security parameter. Also, a negligible function is defined as below.

Definition 3.1. *A function $v : \mathbb{N} \rightarrow \mathbb{R}$ is negligible in κ if for every $n \in \mathbb{N}$, there exists $n_0 \in \mathbb{N}$ such that for every $\kappa \geq n_0$, $|v(\kappa)| \leq 1/\kappa^n$ holds.*

3.1.1 Digital Signature

A digital signature scheme Π includes three algorithms as following:

- **Key Generation.** $(pk, sk) \leftarrow \text{Gen}(1^\kappa)$ on input 1^κ , outputs the public/private key pair (pk, sk) .
- **Signing.** $\sigma \leftarrow \text{Sign}_{sk}(m)$ on inputs the private key sk and a message $m \in \{0, 1\}^*$ outputs the signature σ .
- **Verification.** $b \leftarrow \text{Vrfy}_{pk}(m; \sigma)$ takes the public key pk , a message m and a signature σ as input and outputs 1 if σ is a valid signature on message m created with the private key corresponding to pk . Otherwise, it outputs 0.

Correctness of a digital signature guarantees that for any honestly generated signature σ on the message m w.r.t. a public key pk , $\text{Vrfy}_{pk}(m; \sigma)$ outputs 1. In this work, we assume that the utilised signature schemes are existentially unforgeable under an adaptive chosen-message attack. It guarantees that the probability that an adversary who has access to a signing oracle outputs a valid signature on any new message is negligible

in κ . We call such signature schemes secure. ECDSA [52] is a secure signature scheme that is currently being used in Bitcoin. Schnorr signature [53] is another important secure signature scheme that has been proposed to be introduced in Bitcoin due to its key aggregation and signature aggregation properties.

3.1.2 Hard relation

A relation $\mathcal{R}$ with statement/witness pairs (Y, y) is called a hard relation if (1) There exists a polynomial time generating algorithm $(Y, y) \leftarrow \text{GenR}(1^\kappa)$ that on input 1^κ outputs a statement/witness pair $(Y, y) \in \mathcal{R}$; (2) The relation between Y and y can be verified in polynomial time, and (3) For any polynomial-time adversary $\mathcal{A}$, the probability that $\mathcal{A}$ on input Y outputs y is negligible. We also let $L_{\mathcal{R}} := \{Y \mid \exists y \text{ s.t. } (Y, y) \in \mathcal{R}\}$. Statement/witness pairs of $\mathcal{R}$ can be public/private key of a signature scheme generated by Gen algorithm.

3.1.3 Adaptor Signature

Adaptor signatures appeared first in [14]. Adaptor signature is used in the Generalized channel to tie together the authorisation of a commit transaction and the leakage of a secret value. In what follows, we recall how an adaptor signature works. Given a hard relation $\mathcal{R}$ and a signature scheme Π , an adaptor signature protocol Ξ includes four algorithms as follows:

- **Pre-Signing.** $\tilde{\sigma} \leftarrow \text{pSign}_{sk}(m, Y)$ is a probabilistic polynomial time (PPT) algorithm that on input a private key sk , message $m \in \{0, 1\}^*$ and statement $Y \in L_{\mathcal{R}}$, outputs a pre-signature $\tilde{\sigma}$.
- **Pre-Verification.** $b \leftarrow \text{pVrfy}_{pk}(m, Y; \tilde{\sigma})$ is a deterministic polynomial time (DPT) algorithm that on input a public key pk , message $m \in \{0, 1\}^*$, statement $Y \in L_{\mathcal{R}}$ and pre-signature $\tilde{\sigma}$, outputs a bit b .
- **Adaptation.** $\sigma \leftarrow \text{Adapt}(\tilde{\sigma}, y)$ is a DPT algorithm that on input a pre-signature $\tilde{\sigma}$ and witness y , outputs a signature σ .
- **Extraction,** $\text{Ext}(\sigma, \tilde{\sigma}, Y)$ is a DPT algorithm that on input a signature σ , pre-signature $\tilde{\sigma}$, and statement $Y \in L_{\mathcal{R}}$, outputs $\perp$ or a witness y such that $(Y, y) \in \mathcal{R}$.

Correctness of an adaptor signature guarantees that for an honestly generated pre-signature $\tilde{\sigma}$ on the message m w.r.t. a statement $Y \in L_{\mathcal{R}}$, we have $\text{pVrfy}_{pk}(m, Y; \tilde{\sigma}) = 1$.

Furthermore, when $\tilde{\sigma}$ is adapted to the signature σ , we have $\text{Vrfy}_{pk}(m; \sigma) = 1$ and $\text{Ext}(\sigma, \tilde{\sigma}, Y)$ outputs y such that $(Y, y) \in \mathcal{R}$.

An adaptor signature scheme is secure if it is existentially unforgeable under chosen message attack (aEUF–CMA security), pre-signature adaptable and witness extractable. The aEUF–CMA security guarantees that it is of negligible probability that any PPT adversary who has access to signing and pre-signing oracles outputs a valid signature for any arbitrary new message m even given a valid pre-signature and its corresponding Y on m . Pre-signature adaptability guarantees that every pre-signature (possibly generated maliciously) w.r.t. Y can adapt to a valid signature using the witness y with $(Y, y) \in \mathcal{R}$. Witness extractability guarantees that it is of negligible probability that any PPT adversary who has access to signing and pre-signing oracles outputs a valid signature and a statement Y for any new message m such that the valid signature does not reveal a witness for Y even given a valid pre-signature on m w.r.t. Y . The ECDSA-based and Schnorr-based adaptor signature schemes were constructed and analysed in [14].

3.2 Notations

Throughout this work, we define different attribute tuples. Let U be a tuple of multiple attributes including the attribute attr . To refer to this attribute, we use $U.\text{attr}$. Our focus in this work is on Bitcoin or any other blockchains with *Unspent Transaction Output* (UTXO) model. In this model, units of value—which we call coins—are held in *outputs*. An output θ is a tuple of two attributes, $\theta = (\text{cash}, \varphi)$, where $\theta.\text{cash}$ denotes the number of coins held in this output and $\theta.\varphi$ denotes the condition that needs to be fulfilled to spend the output θ . The condition $\theta.\varphi$ is encoded using any script supported by the underlying blockchain. If the condition $\theta.\varphi$ contains a user P 's public key, we say that P controls or owns the output θ . If satisfying a condition requires authorisations by multiple parties, such a condition contains public keys of all the involved parties separated by $\wedge$ operation(s). Relative time-lock of T rounds in an output condition is denoted by T^+ and means the output cannot be spent unless at least T rounds passed since the output was recorded on the blockchain. A condition might also have several sub-conditions, one of which must be satisfied to spend the output. Different sub-conditions of output are separated by $\vee$ operation(s).

A transaction changes ownership of coins, meaning it takes a list of existing outputs and transfers their coins to a list of new outputs. To distinguish between these two lists, we refer to the list of existing outputs as *inputs*. A transaction TX is formally defined as the tuple $(\text{txid}, \text{Input}, \text{nLT}, \text{Output}, \text{Witness})$. The identifier $\text{TX.txid} \in \{0, 1\}^*$ is computed as $\text{TX.txid} := \mathcal{H}([\text{TX}])$, where $[\text{TX}]$ is called the *body* of the transaction defined as $[\text{TX}] := (\text{TX.Input}, \text{TX.nLT}, \text{TX.Output})$ and $\mathcal{H}$ is a hash function, which is modelled as a random

oracle. The attribute TX.nLT denotes the value of the parameter $nLockTime$, where TX is invalid unless its $nLockTime$ is in the past. The attribute TX.Input is a list of identifiers for all inputs of TX. The attribute TX.Output is a list of new outputs. The attribute $\text{TX.Witness} = (W_1, \dots, W_m)$ is a list of tuples where its i^{th} tuple authorises spending the output that is taken as the i^{th} input of TX. The tuple $W_i = (\eta, \zeta)$ of the witness Tx.Witness contains two attributes where $W_i.\zeta$ denotes the data, e.g. the signature(s), that is (are) required to meet the $W_i.\eta^{\text{th}}$ sub-condition of the output that is taken as the i^{th} input of TX. The signature and pre-signature of party P on TX for $\text{TX.Witness}.W_j.\zeta$ is denoted by $\sigma_{\text{TX}}^{P,j}$ and $\tilde{\sigma}_{\text{TX}}^{P,j}$, respectively, where j can be removed if TX has one input. The i^{th} entry of a list L is denoted by $L[i]$ with $i > 0$.

We use charts to illustrate transaction flows. As Fig. 3.1 shows, double-edge and single-edge rectangles represent published and unpublished transactions, respectively. Also, dotted rectangles represent transactions that are still unprepared to be propagated in the blockchain network. In other words, dotted rectangles denote transactions that lack some required elements (e.g. some signatures in the witness). Considering that TX contains two inputs, each with a value of a and b respectively, then since the output of TX with the value of $a + b$ has two sub-conditions, it is denoted by a diamond shape with two arrows. One of the sub-conditions can be fulfilled by both A and B and is relatively time-locked by T rounds. Another sub-condition can be fulfilled by C where rather than the public key pk_C , the user identity, i.e. C , has been used in the chart. The $nLockTime$ parameter for TX and TX' is 0, so it is not shown inside these transactions.




Figure 3.1: A sample transaction flow.

Table 3.1 summarises the mentioned notations.

Table 3.1: Summary of notations

Notation	Description
TX	Transaction $TX = (\text{txid}, \text{Input}, \text{nLT}, \text{Output}, \text{Witness})$
[TX]	Tuple $(TX.\text{Input}, TX.\text{nLT}, TX.\text{Output})$
TX.txid	Identifier of the transaction TX
TX.nLT	Parameter <i>nLockTime</i> of the transaction TX
TX.Input	List of identifiers for all inputs of TX
TX.Output	List of new outputs for TX
TX.Witness	List of witnesses for TX
$\theta = (\text{cash}, \varphi)$	Output with monetary value cash and condition φ
$W = (\eta, \zeta)$	Witness that meets η^{th} sub-condition of an output using data ζ
$\sigma_{TX}^{P,j}$ (or $\tilde{\sigma}_{TX}^{P,j}$)	Signature (or pre-signature) of P for j^{th} input of TX
T^+	The relative time-lock of T rounds
$L[i]$	i^{th} entry of a list L with $i \geq 1$

Chapter 4

Formal Treatment of Watchtower

4.1 Introduction

In this chapter, we formally define a payment channel and a watchtower and then work on our first research question, RQ1, which aims to formally define different properties of a watchtower scheme and analyse why existing watchtower schemes that fail to meet all the required properties. Monitor [15], DCWC¹ and DCWC* [32], Outpost [22], Cerberus [23], PISA [24] and Fail-safe [35] are the watchtower schemes we will analyse in this chapter.

We construct some parts of this chapter based on our published paper, “FPPW: A Fair and Privacy Preserving Watchtower For Bitcoin” [54].

4.2 Formalisation of Payment Channel and Watchtower

A payment channel is defined as follows:

Definition 4.1 (Payment Channel). *A payment channel γ of blockchain $\mathbb{B}$ is a state machine run by two participants. Let S_i with $i \geq 0$ be the i^{th} channel state. γ consists of three phases, namely create, update and close, as follows:*

- Channel create (state S_0): *Channel participants commit to the initial state S_0 and record it on $\mathbb{B}$.*
- Channel update (from S_i to S_{i+1} for $i \in [0, n)$): *Channel participants update the channel state from S_i to S_{i+1} and commit to S_{i+1} .*

¹Disclose Cascade Watch Commit

- Channel close (state S_n): Any channel participants can close the channel by recording the latest state S_n on $\mathbb{B}$.

Remark 4.1. Definition 4.1 is a generic one for a bi-directional payment channel. Moreover, to analyse the payment channel security in the universal composability framework [55], some ideal functionalities of payment channels can also be found in the literature [14, 28]. We will also present an ideal functionality for our designed payment channel, Daric, in Chapter 7.

The capacity of a payment channel γ is defined below.

Definition 4.2 (Channel Capacity). For a payment channel γ with two participants A and B , let $x_{P,i}$ with $P \in \{A, B\}$ denote the balance of P at the channel state S_i . Then, the channel capacity is denoted by $X_{A,B}$ and equals $x_{A,0} + x_{B,0}$.

Definition 4.3 (Payment Channel Security). A payment channel γ of blockchain $\mathbb{B}$ with two channel participants A and B is μ -secure if and only if given that γ with latest state of S_n is closed by recording the state S_i with $i < n$ on $\mathbb{B}$, an honest party $P \in \{A, B\}$ can claim at least $x_{P,n}$ by recording some evidence proving invalidity of S_i on $\mathbb{B}$ within μ blocks.

The above definition states that when a μ -secure payment channel is closed using a transaction that corresponds with an old state, the honest party has still μ blocks time to get refunded or equivalently to invalidate the published channel state. After μ blocks, the honest party's funds are not guaranteed anymore. To invalidate the published channel state, the honest party must record some evidence on $\mathbb{B}$ which proves the published state is not the latest channel state that parties have committed to. For Replace by Revocation or RbR payment channels (e.g. [14, 15]), the evidence is knowledge of the revocation secret corresponding with the published channel state. This knowledge is practically proved by publishing a revocation transaction that gives all the funds in the channel to the honest party. For Replace by Version or RbV type channels, the evidence is a state with higher version [1, 17].

One or both parties of a payment channel can outsource the task of publishing the evidence to a third-party called watchtower defined as follows.

Definition 4.4 (Watchtower). For a μ -secure payment channel γ of $\mathbb{B}$ with two participants A and B , a watching contract $C = (W, \gamma, P, S_s, \mathcal{C})$ with $\mathcal{C} = \{\mathcal{C}_S, \mathcal{C}_U, \mathcal{C}_T\}$ is a contract between a watchtower W and a party $P \in \{A, B\}$ in which W commits to follow start, update, termination and alarm, as follows:

- Watchtower start (state S_s , condition $\mathcal{C}_S$): Watchtower starts when the condition $\mathcal{C}_S$ is met and P provides the evidence e to W . Let e invalidate a set of channel states $\mathcal{S}$.

The start state S_s is defined by

$$S_s := \begin{cases} S_{\max\{i: S_i \in \mathcal{S}\}+1} & \mathcal{S} \neq \emptyset, \\ S_0 & \mathcal{S} = \emptyset. \end{cases}$$

- Watchtower update (from S_k to S_{k+1} , condition $\mathcal{C}_U$): With each channel update from S_k to S_{k+1} , watchtower updates if the condition $\mathcal{C}_U$ is met and party P provides W with some evidence proving the invalidity of S_k .
- Watchtower terminate (condition $\mathcal{C}_T$): Watchtower terminates once the latest state S_n is recorded on $\mathbb{B}$ or once the condition $\mathcal{C}_T$ is met.
- Watchtower react: Let γ with latest state of S_n be closed by recording the state S_i with $i \neq n$ on $\mathbb{B}$. If this event occurs before the condition $\mathcal{C}_T$ is met, then watchtower W records some evidence proving the invalidity of S_i on $\mathbb{B}$ within μ blocks.

The condition set $\mathcal{C}$ could be a watchtower scheme specific condition, possibly relying on different parameters. In the following, we will discuss the condition set for the existing watchtower schemes.

For DCWC, DCWC* [32] and monitor [15], $\mathcal{C}_S$ and $\mathcal{C}_U$ are $\emptyset$. However, these schemes have not discussed their watchtower termination conditions although channel closure by the hiring party P is definitely included in $\mathcal{C}_T$.

The conditions $\mathcal{C}_S$ and $\mathcal{C}_U$ for Outpost are payments to the watchtower. However, condition $\mathcal{C}_T$ is somehow complicated. An outpost watchtower stores a pre-defined number of evidence per user and deletes older items based on a FIFO order. Thus, when a user has several channels, it is possible that while a channel is still open, its older states are not monitored by the watchtower anymore. Thus, $\mathcal{C}_T$ for Outpost is met once the channel is closed by the hiring party P or with an old state whose corresponding evidence has been removed from the watchtower's evidence list.

The watchtower in Cerberus [23] locks some collateral at the watchtower start phase. This collateral is taken by the hiring party given that the watchtower does not follow the watchtower reaction rules. So, $\mathcal{C}_S$ is locking some collateral by the watchtower. The watchtower is rewarded upon each watchtower update. Also, as part of each watchtower update, the watchtower must provide the channel parties with signatures on two *penalty* transactions that are used by the cheated hiring party to penalise the unresponsive watchtower for not following the watchtower react rules. So, $\mathcal{C}_U$ is payment to the watchtower and signing the corresponding penalty transactions by the watchtower. $\mathcal{C}_T$ is met when the watchtower's collateral is redeemed or the channel is closed by the hiring party P .

Table 4.1: Summary of the Condition $\mathcal{C}_S$ for Different Watchtower Schemes.

	$\mathcal{C}_S$
Monitor [21]	$\emptyset$
DCWC [32]	$\emptyset$
Outpost [22]	$\emptyset$
Cerberus [23]	collateral
PISA [24]	custodian contract with security deposit, payment, signed receipt
Fail-safe [35]	tower contract, payment, signed receipt

Table 4.2: Summary of the Condition $\mathcal{C}_U$ for Different Watchtower Schemes.

	$\mathcal{C}_U$
Monitor [21]	$\emptyset$
DCWC [32]	$\emptyset$
Outpost [22]	payment
Cerberus [23]	payment, signatures of penalty transactions
PISA [24]	payment, signed receipt
Fail-safe [35]	payment, signed receipt

The watchtower in PISA first locks a large deposit in a *smart contract*² called *custodian contract*. This large deposit is forfeited if the watchtower does not follow the watchtower reaction rules. To hire the watchtower, the watchtower is paid by the hiring party in exchange for a signed receipt that can be used to prove the watchtower's wrongdoing before an agreed expiry time. So, $\mathcal{C}_S$ is setting up the custodian contract and locking a large security deposit in this contract by the watchtower, payment to the watchtower and providing the corresponding signed receipt to the hiring party. $\mathcal{C}_U$ is a new payment to the watchtower and providing the corresponding new signed receipt to the hiring party. $\mathcal{C}_T$ is met at an agreed expiry time.

For Fail-safe, $\mathcal{C}_S$ is having a smart contract called the *tower contract* set up by the watchtower on the blockchain, payment to the watchtower, and providing the corresponding signed receipt to the hiring party. $\mathcal{C}_U$ is a new payment to the watchtower and providing the corresponding new signed receipt to the hiring party. The condition $\mathcal{C}_T$ is $\emptyset$, meaning that the watchtower has to monitor the channel and be responsive as long as the channel is open. Otherwise, the watchtower might lose its reward.

Tables 4.1, 4.2 and 4.3 respectively summarise the conditions $\mathcal{C}_S$, $\mathcal{C}_U$ and $\mathcal{C}_T$ for different watchtower schemes.

²A self-executing, programmable code that runs on a blockchain network, enabling automated and trustless execution of predefined actions and agreements between parties.

Table 4.3: Summary of the Condition $\mathcal{C}_T$ for Different Watchtower Schemes.

	$\mathcal{C}_T$
Monitor [21]	at least channel closure by the hiring party P
DCWC [32]	at least channel closure by the hiring party P
Outpost [22]	(1) channel closure by the hiring party P or (2) channel closure using a transaction with removed evidence
Cerberus [23]	(1) redeemed collateral or (2) channel closure by the hiring party P
PISA [24]	expiry time
Fail-safe [35]	$\emptyset$

4.3 Watchtower Service Properties

In this section, we formally define different new properties of a watchtower service.

4.3.1 Agility

This section defines *agility* for a watchtower. The agility of a watchtower shows how flexibly a watching process starts and terminates.

Definition 4.5 (Agility). *A payment channel γ between two parties A and B with watchtower provides agility iff*

- *two distinct contracts $C_1 = (W_1, \gamma, P, S_s, \mathcal{C})$ and $C_2 = (W_2, \gamma, P, S_s, \mathcal{C})$ may be formed, where $P \in \{A, B\}$ and $W_1 \neq W_2$, and*
- *Any watching contract can start and terminate at any arbitrary state of the channel γ .*

According to the above definition, Cerberus and Fail-safe do not achieve agility as the watchtower in both schemes should be specified in the channel create phase. Moreover, watching for both schemes must start and terminate when the channel is created and closed, respectively. [23] discusses that if the watching contract for Cerberus terminates before channel closure, the hiring party, e.g. B , might continue using the channel given that he is always online. However, this statement is correct with this assumption that after the watchtower termination phase, the watchtower will not be involved in this channel any more. Since this assumption might be violated in the real world, B would be at risk of losing some funds in the channel. In more detail, once A publishes an old channel state on the blockchain, the dispute period starts and B must publish a revocation transaction within this period to take all the channel funds. However, if

the dishonest party A has colluded with the watchtower (whose collateral has already been redeemed), they can cooperatively claim the channel funds before the revocation transaction is published by B . So, even staying always online would not guarantee the safety of the honest party's funds in the channel. In other words, once the watchtower redeems its collateral, the channel for the hiring party transforms to a μ -Secure channel with $\mu = 0$. So, the hiring party must close the channel before the watchtower's collateral redemption (or equivalently before the time when the watchtower contract terminates). The watching contract for other schemes can start and terminate at any time with any arbitrary watchtower. So, Monitor, DCWC, DCWC*, Outpost and PISA achieve agility.

4.3.2 Privacy

One of the properties of payment channels is that off-chain transactions are not broadcast on the blockchain and hence only the channel parties know how the channel funds are redistributed between them. Thus, in the following, to evaluate how the usage of watchtowers can affect this benefit, we define a *privacy game* and then based on this privacy game we define *Weak Privacy Against Watchtower*.

Challenge. Let there exist two payment channels where the first one is between honest channel parties A and B and the second one is between honest channel parties A' and B' and both channels have the same number of channel updates n and the same channel setup, i.e. $x_{A,0} = x_{A',0}$, $x_{B,0} = x_{B',0}$, $x_{A,n} = x_{A',n}$ and $x_{B,n} = x_{B',n}$. Let $\mathbf{x}_{P,[i,j]}$ show the sequence of balance values of party P between i^{th} to j^{th} states of the payment channel that P is involved in. Assume that $\mathcal{A}$ is any passive PPT adversarial watchtower excluding A , B , A' and B' which watches these two channels. To challenge $\mathcal{A}$, the challenger selects a random bit b and gives the sequence $(\mathbf{x}_{P,[1,n-1]}, \mathbf{x}_{\bar{P},[1,n-1]})$ to $\mathcal{A}$ where $P = A$ and $\bar{P} = B$ if $b = 0$ and $P = A'$ and $\bar{P} = B'$ otherwise.

Output. The adversary $\mathcal{A}$ outputs a bit b' to guess that the received sequence belongs to the first or the second channel. The adversary wins the game if and only if $b = b'$.

Definition 4.6 (Weak Privacy Against Watchtower). *A payment channel with a watchtower provides weak privacy against the watchtower if according to the privacy game $|\Pr[b = b'] - 1/2|$ is negligible in κ .*

Remark 4.2. *For any multihop payment routed via the channel between A and B or the channel between A' and B' , we assume that the passive adversary is not involved as a channel party in routing such payments.*

The privacy guarantee would be stronger, if in the defined privacy game, (1) the sequence $(\mathbf{x}_{P,[0,n]}, \mathbf{x}_{\bar{P},[0,n]})$ is given to the passive PPT adversarial watchtower and (2) the watchtower determines if the channel setup for two channels are the same or not. This strong privacy implies weak privacy because the channel setup in the weak privacy game is a special case of that in the strong privacy game. Moreover, the watchtower in the strong privacy game receives more input data in the challenge phase than the watchtower in the weak privacy game. Thus, the probability of her winning the strong privacy game would be at least equal to the corresponding probability in the weak privacy game.

Monitor, DCWC, DCWC* and Outpost achieve strong privacy against the watchtower because the watchtower receives at most the encrypted version of channel transactions and if the channel parties behave honestly, these encrypted transactions are not decrypted at all. PISA and Fail-safe achieve weak privacy against the watchtower because for these schemes, states are invisible to the watchtower and only state hash values are given to the watchtower and since a large random value is also used in the computation of the hash value, finding the pre-image by exhaustive search is also infeasible. However, The first and the last states of the channel are revealed. Cerberus does not achieve privacy against the watchtower as the watchtower learns the balances of both channel parties in the channel.

Another important privacy-related subject regarding a payment channel and its corresponding watchtower is the knowledge of a third party (i.e. any party other than the hiring party and the watchtower) regarding the hiring status of the watchtower for a given channel. For example in the channel γ with two parties A and B , such knowledge can be of importance to a channel party, let's say party A , given that he is malicious and is looking for an appropriate time (i.e. when party B is not using the watchtower with a significant probability) to broadcast an old state or to a third party (other than party A) who has some other channels with party B . Such information can help him to conduct a behavioural analysis on party B to find the best time to attack him on his other channels. This type of privacy is also defined below.

Definition 4.7 (Weak Watchtower Privacy against Third-Party). *Let W be a watchtower. A challenger A establishes a payment channel γ with B , samples $b \in \{0, 1\}$ randomly and hires (does not hire) W if $b = 1$ ($b = 0$, respectively). The watchtower W achieves weak privacy against third-party if for all PPT adversaries B , we have that*

$$\left| \Pr[1 \leftarrow B(\cdot) | b = 1] - \frac{1}{2} \right|$$

is negligible in κ .

Definition 4.8 (Strong Watchtower Privacy against Third-Party). *Let W be a watchtower. A challenger A establishes a payment channel γ with B , samples $b \in \{0, 1\}$ randomly and*

hires (does not hire) W if $b = 1$ ($b = 0$, respectively). The watchtower W achieves strong privacy against third-party if for all PPT adversaries $\mathcal{A}$ including B , we have that

$$\left| \Pr[1 \leftarrow \mathcal{A}^{BAW\text{-view}(\cdot)} | b = 1] - \frac{1}{2} \right|$$

is negligible in κ , where $BAW\text{-view}$ is an oracle in which $\mathcal{A}$ has access to what B sees on γ and sees all public keys of A and W and all transactions related to them on $\mathbb{B}$.

The strong watchtower privacy against third-party implies the weak watchtower privacy against third-party. The reason is that the PPT adversary in the former also includes B . Moreover, she accesses more information than the adversary in the weak privacy definition.

According to Definition 4.8, Monitor, DCWC and DCWC* provide strong privacy against third-party because there is not any transaction between the hiring party and the watchtower. Thus, even knowledge of all the public keys and transactions of the hiring party and the watchtower cannot help the third party to guess whether there is a watchtower contract between the hiring party and the watchtower.

However, PISA, Outpost and Cerberus do not provide strong privacy against third-party as for all these schemes payment to the watchtower is done through a one-way payment channel per update and hence $\mathcal{A}^{BAW\text{-view}(\cdot)}$ for these watchtowers can output the correct result with non-negligible probability. For instance, no transaction including public keys of A and W implies no one-way channel between A and W and hence no watching contract between them.

PISA, Outpost and Cerberus potentially achieve weak privacy against third-party given that any relationship between transactions in γ and watchtower-related transactions are invisible to B (as defined in Definition 4.7). To achieve that, public keys used in watchtower-related transactions must be independent of the ones used in γ . This condition can be simply met. Moreover, watchtower-related transactions must not change the distribution of different transaction types on the underlying blockchain. For example, if transactions with multi-signature outputs (and in particular 2-of-2 multi-signature class) are rarely exchanged between the users, usage of such transactions for watchtower purposes might be distinguishable from other transactions. In such a case, a straightforward option is aggregating public keys of the watchtower and the hiring party in watchtower-related transactions. Also, there must be some random differences between the time when the channel is created and when the watchtower-related transactions are published on the blockchain.

Fail-safe does not provide weak privacy against third-party as B knows the hiring status of W upon establishment of the channel γ .

4.3.3 Fairness and Coverage

In this section, we formalise the concepts of fairness and coverage in a watchtower scheme. Fairness with respect to the hiring party (watchtower) is actually a factor to evaluate the level of guarantee that the watchtower (hiring party) provides to the hiring party (watchtower) on its service (payment).

Definition 4.9 (Channel party α -Fairness). *A payment channel with watchtower is α -party-fair, if the following holds for an honest channel party P :*

- *P can close the channel at any time and*
- *α is the largest real number such that regardless of the reward that P pays to the watchtower, P loses at most $(1 - \alpha) \cdot x_P$ coins in the channel where x_P denotes balance of P in the latest channel state.*

Note that $0 \leq \alpha \leq 1$, where $\alpha = 1$ implies that the honest party P will not lose any funds in the channel and $\alpha = 0$ means that P might lose all of his funds. Since for Monitor, DCWC, Outpost and Fail-safe, the hiring party P might lose all his funds in the channel, these schemes are unfair w.r.t. the channel party (i.e. they are α -party-fair with $\alpha = 0$). The value of the watchtower's collateral for Cerberus is around the channel capacity. So the hiring party would not lose any funds in the channel. In other words, $\alpha = 1$ for Cerberus. Although for PISA α is adjustable, the hiring party must not logically accept a PISA watchtower with $\alpha < 1$. Otherwise, he might lose some funds in the channel. Thus, the value of α for PISA must be 1 in practice or equivalently the watchtower must lock some collateral per channel where the value of the collateral must be equal to the channel capacity. This collateral will be taken by the hiring party given that the watchtower is unresponsive upon fraud.

Definition 4.10 (Watchtower Fairness). *A payment channel with a watchtower is watchtower-fair if the following holds for an honest watchtower $\mathcal{W}$:*

- *$\mathcal{W}$ is rewarded with some non-zero amounts of coins and*
- *given that $\mathcal{W}$ has locked some collateral as part of the watching service, it is of negligible probability that the honest watchtower cannot redeem all the collateral once watching terminates according to the watching agreement.*

Monitor [21] and DCWC [32] are called unfair w.r.t. the watchtower because, for these schemes, it is possible that the watchtower is not rewarded. In more detail, the watchtower in these schemes is rewarded if and only if (1) the counterparty of the hiring party P publishes an old channel state on the blockchain and (2) the watchtower succeeds in

broadcasting the corresponding evidence and wins the race against other watchtowers, which are simultaneously monitoring the same channel for P . On the other hand, due to the punishment mechanism of the Lightning Network, fraudulent channel closures and hence payments to the watchtower would be rare and this can dis-incentivise entities to run such services. PISA, outpost, Cerberus and Fail-safe are watchtower-fair as the watchtower in these schemes is rewarded upon each channel update and the watchtower in Cerberus can redeem its collateral at any time.

Next, we define β -coverage, which basically measures the capability of a watchtower (on a scale between 0 to 1) in watching all the existing payment channels on a fixed Blockchain.

Definition 4.11 (Coverage). *For a blockchain $\mathbb{B}$ with N payment channels, a watchtower $\mathcal{W}$ provides β -coverage with $\beta := \frac{\mathcal{X}}{\mathcal{C} + \mathcal{X}}$, where $\mathcal{C}$ is the total collateral required by $\mathcal{W}$ to watch all payment channels for both channel parties and $\mathcal{X}$ is the total capacity of all channels.*

The parameter β can take any value in the interval $[0, 1]$. For Cerberus and PISA (with $\alpha = 1$), β equals $\frac{1}{3}$ because, for these schemes, collateral of the watchtower must be twice the channel capacity if the watchtower is going to be hired by both channel parties. Although PISA allows lower values of collateral, such values cannot provide channel party α -fairness with $\alpha = 1$ and hence cannot guarantee that the honest party does not lose any funds. The watchtower in other schemes does not need to lock any collateral per channel. So they provide β -coverage with $\beta := 1$.

The following Lemma shows a trade-off between watchtower fairness and coverage.

Lemma 4.1. *For an (α, R) -fair watchtower W with β -coverage, we have:*

1. $\mathcal{X} \leq \beta \cdot MC$, where MC denotes the market cap of the used cryptocurrency and $\mathcal{X}$ is defined in Definition 4.11,
2. $\mathcal{C} \geq \alpha \cdot \mathcal{X}$, where $\mathcal{C}$ is defined in Definition 4.11 and
3. $\beta \leq \frac{1}{1+\alpha}$.

Proof. We prove each item in the following:

1. Since $\mathcal{X} + \mathcal{C} \leq MC$ and W has β -coverage we have:

$$\mathcal{X} = \beta \cdot (\mathcal{C} + \mathcal{X}) \leq \beta \cdot MC.$$

2. For the channel γ_i with channel parties A_i and B_i , as defined in Definition 4.11, we have two watching contracts $C_i = (W, \gamma_i, A_i, S_i, \mathcal{C}_i)$ and $C'_i = (W, \gamma_i, B_i, S'_i, \mathcal{C}'_i)$. Assume that γ_i is closed by a party, let's say party B , at state S_n by committing to state S_j with $j < n$ before the condition $\mathcal{C}_i$ being met. For the case $x_{A_i,n} = X_{A_i,B_i}$ and $x_{A_i,j} = 0$, honest party A is cheated out of its total funds which equals X_{A_i,B_i} . However, since W is (α, R) -fair, party A does not lose more than $(1 - \alpha) \cdot x_{A_i,n} = (1 - \alpha) \cdot X_{A_i,B_i}$ if he is honest. Thus, capital of W for this channel cannot be less than

$$X_{A_i,B_i} - (1 - \alpha) \cdot X_{A_i,B_i} = \alpha \cdot X_{A_i,B_i}$$

Otherwise, party A loses more than $(1 - \alpha) \cdot x_{A_i,n}$, which contradicts with definition of fairness. Therefore, the total capital of W must satisfy

$$\mathcal{C} \geq \sum_{i=1}^N \alpha \cdot X_{A_i,B_i} = \alpha \cdot \mathcal{X}. \quad (4.1)$$

3. Based on the definition of coverage and (4.1), we have:

$$\mathcal{X} = \beta \cdot (\mathcal{C} + \mathcal{X}) \geq \beta \cdot (\alpha \cdot \mathcal{X} + \mathcal{X})$$

Since $\mathcal{X} \neq 0$ and $\alpha \neq -1$, we have:

$$\beta \leq \frac{1}{1 + \alpha}.$$

□

The above Lemma shows that although an increase in α for a watchtower promotes its fairness with regard to the hiring party, it raises the required capital of the watchtower and negatively affects its coverage. In other words, if all payment channel parties seek to achieve full guarantee by their watchtowers, in the best case at most half of the market cap of a given cryptocurrency can be used for transactions between channel parties and another half must be locked as collateral by their watchtowers.

The Table 4.4 summarises the comparison results for the existing watchtower schemes. The table illustrates that the current watchtower schemes face challenges in being deployed on Bitcoin, ensuring privacy, or providing channel-party 1-fairness.

4.4 Conclusion

In this chapter, we formalised payment channels and watchtowers. We also formally defined different properties of a watchtower scheme including agility, privacy, fairness

Table 4.4: Comparison of Different Watchtower Schemes.

	Bitcoin Support	Agility	Priv. ag. Watch.	Priv. ag. 3 rd Party	Watch. Fairness	α	β
Monitor [21]	Yes	Yes	Strong	Strong	No	0	1
Outpost [22]	Yes	Yes	Strong	Weak	No	0	1
DCWC [32]	No	Yes	Strong	Strong	No	0	1
DCWC* [32]	Yes	Yes	Strong	Strong	No	0	1
Cerberus [23]	Yes	No	-	Weak	Yes	≈ 1	$\approx \frac{1}{3}$
PISA [24]	No	Yes	Weak	Weak	Yes	Adj.	$\frac{1}{1+2\alpha}$
Fail-safe [35]	No	No	Weak	-	Yes	0	1

and coverage and compared the existing watchtower schemes against these properties. Furthermore, we proved a trade-off between fairness and coverage of a watchtower. This trade-off shows although an offline channel party with a watchtower satisfying $\alpha < 1$ would be at the risk of losing some funds, the fully fair schemes (Cerberus and PISA with $\alpha = 1$) cannot achieve acceptable coverage. Moreover, upon examining existing watchtower schemes on Bitcoin, we observed that none of them can simultaneously achieve channel party 1-fairness and privacy against watchtowers.

Chapter 5

FPPW: a fair and privacy preserving Bitcoin watchtower

In this chapter, our focus is on addressing our second research question, RQ2. Our objective is to develop a watchtower scheme specifically tailored for Bitcoin, aiming to overcome the limitations of existing schemes. We concentrate on enhancing the fairness and privacy properties of the watchtower in this investigation. We construct this chapter based on the full version of our published paper, “FPPW: A Fair and Privacy Preserving Watchtower For Bitcoin” [54] (The full version paper is available at <https://eprint.iacr.org/2021/117.pdf>).

5.1 Introduction

Monitor [21] is the first watchtower scheme for Lightning Network, which mainly focuses on channel privacy against watchtower. However, Monitor has two main issues, both of which are related to fairness. Firstly, honest watchtowers might be rewarded for fraud (i.e. broadcast of an old state on the blockchain), which is unfair with respect to (w.r.t.) the watchtower. Secondly, honest parties cannot penalise the unresponsive watchtower, which is unfair towards an honest hiring party.

DCWC* [32] proposes the usage of a network of watchtowers that must cooperate to maximise their interest. This reduces the probability that the channel gets finalised with an old state. However, watchtowers might still crash or get unresponsive without being penalised by the hiring party. Also, the reward mechanism is still unfair w.r.t. the watchtower. Outpost [22] solves the issue of fairness towards the watchtower by paying her per channel update.

Table 5.1: Different Properties of the FPPW Scheme.

	Bitcoin Support	Agility	Priv. ag. Watch.	Priv. ag. 3 rd Party	Watch. Fairness	α	β
Cerberus [23]	Yes	No	-	Weak	Yes	≈ 1	$\approx \frac{1}{3}$
PISA [24]	No	Yes	Weak	Weak	Yes	Adj.	$\frac{1}{1+2\alpha}$
FPPW	Yes	No	Weak	Weak	Yes	≈ 1	$\approx \frac{1}{2}$

Cerberus [23] and PISA [24] elegantly provide fairness w.r.t. the hiring party. However, PISA fails to be deployed in cryptocurrencies with limited script languages such as Bitcoin and Cerberus sacrificing the channel privacy against watchtower. In particular, the Cerberus watchtower learns the distribution of funds in the channel. Thus, the main motivation of this chapter is designing a watchtower scheme for Bitcoin that achieves both: (1) fairness w.r.t. both the hired watchtower and her hiring party and (2) channel privacy against the watchtower.

The contributions of this chapter are as follows:

- In Section 5.3, we present a new privacy-preserving payment channel with a watchtower scheme for Bitcoin called FPPW, which is fair w.r.t. all channel participants and allows the channel parties to go offline for a long period of time. To be more precise, FPPW is an extension of a new variant of the Generalized channel, called NVG, which is introduced in Section 5.3.2.1. Furthermore, in Section 5.5, we prove that our design achieves fairness w.r.t. all channel participants and unlike Cerberus, it provides weak privacy against the watchtower. We also show that the coverage of FPPW is $\frac{1}{2}$ which is better than that of Cerberus and PISA. Table 5.1 summarises the mentioned properties for FPPW and compares it with PISA and Cerberus.
- In Section 5.6, we propose a fee handling mechanism that allows the channel participants to determine the fee for different transactions at the time when fraud occurs. Furthermore, a proof-of-concept implementation of FPPW channels on Bitcoin is provided in Section 5.9.

5.3 FPPW Overview

5.3.1 System Model

Cryptographic primitives that have been used in FPPW (i.e. digital signature, hard relation and adaptor signature as defined in Section 3.2) are cryptographically secure. There is an authenticated and secure end-to-end communication channel between channel parties. The watchtower and channel parties are rational and might deviate from the protocol if it increases their profit. Also, each pair of participants might collude with each other if it raises the total profit of colluding participants. The watchtower is an always online service provider, but channel parties can go offline for a long period (approximately T rounds). Furthermore, the underlying blockchain contains a distributed ledger that achieves security [56]. When a valid transaction is propagated in the blockchain network, it is definitely included in the blockchain ledger immediately (i.e. the confirmation delay Δ is 1).

Remark 5.1. *FPPW channels can work with any confirmation delay. However, we assume that the confirmation delay is 1 to simplify the protocol and its analysis.*

5.3.2 Overview

A payment channel (as defined in Definition 4.1) contains a sequence of state updates between two parties where only its first and last states are recorded on the blockchain. The two channel parties process all the intermediate state updates off-chain. This eliminates the need to confirm every state update, i.e. every transaction, on the blockchain. However, as one may submit an intermediate state (which is already revoked by a later state) to the blockchain, the channel parties will need to get online frequently to monitor and punish such misbehaviours. Such a requirement may be impractical for some users. Thus, the watchtower is introduced as a third party to act on behalf of the channel parties.

FPPW is a fair and privacy-preserving watchtower service for generalised channels [14]. The watchtower in an FPPW channel obtains no data on intermediate state updates. To provide fairness towards the watchtower, the FPPW service rewards the watchtower for channel creation and per channel update. Furthermore, to achieve fairness w.r.t. channel parties, the watchtower must lock some collateral, which can be redeemed by the watchtower if the watchtower is responsive upon fraudulent channel closures. If the watchtower is dishonest and the channel is closed at an old state, protocol guarantees that the cheated party can penalise the watchtower by taking its collateral. Watchtower can terminate its employment at any time. Then, the channel parties can update the

channel on-chain and hire a new watchtower or continue using the channel. In the latter case, channel parties must get online frequently.

In the rest of this section, we will provide an overview of FPPW. To do that, we will present a simple payment channel called NVG and then, we add a watchtower service to NVG which is fair with respect to the hiring party. Finally, we make our solution fair to the watchtower by allowing the watchtower to terminate its service.

5.3.2.1 NVG: A New Variant of the Generalized Channel

Fig. 5.1 depicts a New Variant of the Generalized channel [14], called NVG channel. An NVG channel is created once channel parties publish a funding transaction on the blockchain and hence fund a 2-of-2 multi-signature output¹ on the ledger. The i^{th} channel state includes a commit transaction $\text{TX}_{\text{CM},i}$ as well as a split transaction $\text{TX}_{\text{SP},i}$. The commit transaction sends the channel funds to a new joint account which is shared between the channel parties. The output of the commit transaction has two sub-conditions. The first sub-condition which is not time-locked, as we will explain later, is used for revocation purposes. The second sub-condition is relatively time-locked by t rounds with $t > \Delta$ and is met by the corresponding split transaction. Split transaction distributes the channel funds between the channel parties and hence represents the channel state.

The transaction $\text{TX}_{\text{CM},i}$ requires signatures of both parties A and B to be published. To generate $\sigma_{\text{TX}_{\text{CM},i}}^B$, party A generates a statement/witness pair $(Y_{A,i}, y_{A,i})$ and sends the statement $Y_{A,i}$ to B . Then, party B uses the pre-signing algorithm pSign of the adaptor signature and A 's statement $Y_{A,i}$ to generate a pre-signature $\tilde{\sigma}_{\text{TX}_{\text{CM},i}}^B$ on $[\text{TX}_{\text{CM},i}]$ and sends the pre-signature to A . Thus, whenever it is necessary, A is able to use the adaptation algorithm adapt of the adaptor signature to transform the pre-signature to the signature $\sigma_{\text{TX}_{\text{CM},i}}^B$ and publish $\text{TX}_{\text{CM},i}$ on-chain. This also enables B to apply the extraction algorithm Ext on the published signature and its corresponding pre-signature to extract the witness value $y_{A,i}$. The witness value, as will be seen, allows the honest party to punish the dishonest channel party by claiming all the channel funds, given that the published commit transaction is already revoked.

As one may submit an intermediate state (which is already replaced by a later state) to the blockchain, the channel parties will need to punish such misbehaviours. Thus, upon channel update from state i to $i + 1$, a revocation transaction $\text{TX}_{\text{RV},i}$ is created by parties. Unlike the split transaction, the revocation transaction can immediately spend the output of the corresponding commit transaction $\text{TX}_{\text{CM},i}$ using its first sub-condition which

¹An output that utilizes a script containing two public keys, requiring their corresponding signatures for spending the output.

does not contain any time-lock. Thus, if the revoked commit transaction $TX_{CM,i}$ is published by a channel party, let's say A , party B can immediately publish the revocation transaction $TX_{RV,i}$. Moreover, since commit transactions are signed using the adaptor signature, once $TX_{CM,i}$ is published by A , the witness $y_{A,i}$ is revealed to B . Thus, only B who knows both $y_{A,i}$ and $y_{B,i}$ can meet the condition $Y_{A,i} \wedge Y_{B,i}$ in the output of the revocation transaction and hence B will actually be the owner of all the channel funds. Broadcast of the latest commit transaction does not pose any risk to its broadcaster because the parties have not signed its corresponding revocation transaction yet.




Figure 5.1: NVG Channel Transactions Flow

5.3.2.2 Adding a Watchtower Service with Fairness w.r.t. the Hiring Party to NVG

The watchtower in the introduced scheme is given the revocation transactions for the revoked states and is supposed to publish the one whose corresponding commit transaction is published on the blockchain. However, such a watchtower service would be unfair with respect to the hiring party as the watchtower might become unresponsive upon fraud. To resolve this issue, the watchtower publishes a transaction called *collateral* transaction TX_{CL} and locks its collateral in a 3-of-3 multi-signature output shared between channel parties and the watchtower. The value of the collateral c equals the channel capacity $a + b$. If the watchtower does not appropriately act upon fraud, the cheated party has this chance to publish a transaction called *penalty* transaction 1 $TX_{PN_1,i}$ and take the watchtower's collateral. In other words, if the revoked commit transaction $TX_{CM,i}$ is published on the blockchain, either its corresponding revocation transaction $TX_{RV,i}$ is published by the watchtower or split transaction $TX_{SP,i}$ and then penalty transaction 1 $TX_{PN_1,i}$ are published by the cheating party and honest party, respectively.

As Fig. 5.2 shows, to deploy this, we add an auxiliary output with the least value supported by Bitcoin (denoted by ϵ) to the commit transaction. If the watchtower is responsive, once the revoked commit transaction $TX_{CM,i}$ is published, the watchtower instantly publishes the revocation transaction $TX_{RV,i}$ and invalidates both the split $TX_{SP,i}$ and penalty transaction 1 $TX_{PN_1,i}$. Otherwise, the dishonest party, let's say A , also publishes the split transaction and invalidates the revocation transaction. But B , who is

capable to extract A 's witness (i.e. $y_{A,i}$), adds the required signatures to the penalty transaction 1 $TX_{PN_1,i}$ and publishes it on the blockchain. Penalty transaction 1 also spends the collateral transaction's output. The only output of the penalty transaction 1 $TX_{PN_1,i}$ can be claimed by B who knows both $y_{A,i}$ and $y_{B,i}$. Therefore, since the watchtower's collateral equals the channel capacity, B will not lose any funds in the channel.




Figure 5.2: Adding a Fair Watchtower to NVG

5.3.2.3 Allowing Watchtower to Terminate its Service

Channel parties in the introduced scheme can go offline for any arbitrary period as the watchtower cannot redeem its collateral without authorisation from both channel parties. However, this would be unfair with respect to the watchtower because the watchtower's collateral might be locked forever. Therefore, there must be a way for the watchtower to terminate its service. So, we add a transaction called *reclaim* transaction $TX_{RC,i}$ to our solution. Once the watchtower decides to terminate its service, it publishes the reclaim transaction on-chain. Then, the watchtower can claim the reclaim transaction's output after T rounds which is significantly larger than t . We also add a penalty transaction 2 $TX_{PN_2,i}$ to each revoked state. The only difference between this transaction and penalty transaction 1 is that penalty transaction 2 spends the output of the reclaim transaction (rather than the collateral transaction's output) as well as the auxiliary output of the commit transaction. It means that while the output of the reclaim transaction is unspent, the watchtower might be penalised by channel parties. So channel parties can practically go offline for a long period (approximately T rounds). Fig. 5.3 depicts the transaction flows for the FPPW channel.




Figure 5.3: FPPW Channel Transaction flow

5.4 FPPW Protocol Description

The lifetime of an FPPW channel can be divided into 5 phases including *create*, *update*, *close*, *react* and *terminate*. We explain these phases in the following sections. The first three phases (create, update and close) correspond to different phases of a payment channel as defined in Definition 4.1. A watchtower scheme is also integrated into the FPPW design. The phases create, update and react in an FPPW channel respectively correspond to the phases start, update and react of a watchtower as defined in Definition 4.4. Also, both phases close and terminate in an FPPW channel lead to the watchtower termination as defined in Definition 4.4. Table 5.2 summarises the mentioned correspondences.

The cryptographic primitives, used in these phases, are as follows: A digital signature scheme $\Pi = (\text{Gen}, \text{Sign}, \text{Vrfy})$; a hard relation $\mathcal{R}$ with generating algorithm $\text{GenR} = \text{Gen}$; an adaptor signature scheme $\Xi_{\Pi, \mathcal{R}} = (\text{pSign}, \text{pVrfy}, \text{Adapt}, \text{Ext})$. The more detailed definitions of these cryptographic primitives can be found in Section 3.2. We assume that the watchtower is hired by both channel parties. However, FPPW can be simply extended to situations where only one party hires the watchtower. FPPW for such scenarios will be presented in Section 5.10.

Table 5.2: Different Phases of an FPPW Channel

FPPW Phases	Payment Channel Phases	Watchtower Phases
Create	Create	Start
Update	Update	Update
Close	Close	Terminate
React	-	React
Terminate	-	Terminate

5.4.1 Create

FPPW channel creation phase includes a funding transaction, a commit transaction and a split transaction. The funding transaction locks funds of the channel parties in a 2-of-2 multi-signature output and can be claimed only if both parties agree and cooperate with each other. The commit transaction is held by both channel parties and sends all the channel funds to a joint account that can be spent by the corresponding split transaction after t rounds. The split transaction actually represents the channel state and distributes the channel funds between the channel parties. The quantity t , which represents the dispute period, exists to ensure that there is enough time for punishing the dishonest channel party in the case of fraud (i.e. if the published commit transaction corresponds with a revoked state). Parties finally publish the funding transaction on the blockchain to create the channel. However, since its output can be spent if both parties cooperate, one party might lock the funds by being unresponsive. To avoid such situations, before signing and publishing the funding transaction, both channel parties must sign commit and split transactions.

Additionally, two other transactions are created in this phase: (1) the collateral transaction and (2) the reclaim transaction. These two are used for watchtower services. Using the collateral transaction, the watchtower locks its collateral in a 3-of-3 multi-signature output shared between channel parties and the watchtower. Collateral is awarded to the cheated channel party if the watchtower does not appropriately act upon fraud. The value of the collateral equals the channel capacity. Using the reclaim transaction, the watchtower can start the process of reclaiming its collateral. The watchtower can finally redeem its collateral by claiming the output of the reclaim transaction after a large relative time-lock of T rounds with $T \gg t$ which is called the penalty period. If channel parties get online at least once every $T - 1$ rounds, they will always have enough time to take the dishonest watchtower's collateral as compensation and prevent an unresponsive watchtower from redeeming its collateral. However, if the honest watchtower has published the reclaim transaction to withdraw its service, channel parties will have two options. They can either update the channel on-chain with a new watchtower or remain

almost always online to prevent fraudulent channel closures. The collateral transaction is finally recorded on-chain. However, to avoid any hostage situation, before publishing the collateral transaction, the watchtower must receive the channel parties' signatures on the reclaim transaction.

All the above-mentioned transactions are further explained hereinafter.

- **Funding transaction:** Using this transaction, channel parties A and B open an FPPW channel. If A (B , respectively) uses the x^{th} (y^{th} , respectively) output of a transaction with transaction identifier of $txid_A$ ($txid_B$, respectively) to fund the channel with a (b , respectively) coins, the funding transaction is as follows²:

$$\begin{aligned} \text{TX}_{\text{FU}}.\text{Input} &:= (txid_A \| x, txid_B \| y), \\ \text{TX}_{\text{FU}}.\text{Output} &:= \{(a + b + \epsilon, pk_A \wedge pk_B)\}, \\ \text{TX}_{\text{FU}}.\text{Witness} &:= ((1, \sigma_{\text{TX}_{\text{FU}}}^{A,1}), (1, \sigma_{\text{TX}_{\text{FU}}}^{B,2})). \end{aligned} \quad (5.1)$$

where ϵ is the minimum value supported by the Bitcoin blockchain and a and b are the initial balance of A and B in the channel (regardless of the negligible value $\epsilon/2$). In other words, A and B fund the channel with $a + \epsilon/2$ and $b + \epsilon/2$, respectively. Output of TX_{FU} is a 2-of-2 multi-signature output shared between A and B . The public keys pk_A and pk_B of A and B are generated using the key generation algorithm of the underlying digital signature Gen.

- **Commit transaction:** There exists one commit transaction per state but only the first one ($\text{TX}_{\text{CM},i}$ with $i = 0$) is created at the channel create phase. $\text{TX}_{\text{CM},i}$ is as follows:

$$\begin{aligned} \text{TX}_{\text{CM},i}.\text{Input} &:= \text{TX}_{\text{FU}}.\text{txid} \| 1, \\ \text{TX}_{\text{CM},i}.\text{Output} &:= ((a + b, \varphi_1 \vee \varphi_2), \\ &\quad (\epsilon, \varphi'_1 \vee \varphi'_2 \vee \varphi'_3)) \\ \text{TX}_{\text{CM},i}.\text{Witness} &:= \{(1, \{\sigma_{\text{TX}_{\text{CM},i}}^A, \sigma_{\text{TX}_{\text{CM},i}}^B\})\} \end{aligned} \quad (5.2)$$

with $\varphi_1 := pk_A \wedge pk_B \wedge t^+$, $\varphi_2 := pk_A \wedge pk_B \wedge pk_W$, $\varphi'_1 := pk_B \wedge Y_{A,i} \wedge t^+$, $\varphi'_2 := pk_A \wedge pk_B \wedge pk_W$ and $\varphi'_3 := pk_A \wedge Y_{B,i} \wedge t^+$ where $Y_{A,i}$ and $Y_{B,i}$ are statements of a hard relation $\mathcal{R}$ generated by A and B for the i^{th} state using the generating algorithm GenR and t^+ shows relative time-lock of t rounds. The first output with a value of $a + b$ is the main output. Normally, if parties act honestly and $\text{TX}_{\text{CM},i}$ is published

²We assume that funding sources of TX_{FU} are two typical UTXOs owned by A and B .

on-chain, the first sub-condition of its main output ($pk_A \wedge pk_B \wedge t^+$) is met by $TX_{SP,i}$ after t rounds. The second output of $TX_{CM,i}$ with a value of ϵ is the auxiliary output, which as will be explained in Section 5.4.2, is only used for watchtower purposes.

The transaction $TX_{CM,i}$ requires signatures of both parties A and B to be published. To sign the commit transaction $TX_{CM,i}$, each party uses their counterparty's statement to generate a pre-signature on the commit transaction for their counterparty. Then, each party can use the corresponding witness to adapt the pre-signature, received from their counterparty to a valid digital signature on the commit transaction. So, if the commit transaction $TX_{CM,i}$ is published by A , the witness $y_{A,i}$ is revealed to B and vice versa. The witness value, as will be seen in Section 5.4.2, might be used to penalise the dishonest channel party or the unresponsive watchtower.

Remark 5.2. A has two public keys in the first output of $TX_{CM,i}$, which for simplicity, we denote them both by pk_A . However, in practice, such public keys are selected dis-jointly. This is also extended to other participants and other outputs.

- **Split transaction:** $TX_{SP,i}$ actually represents the i^{th} channel state where only the first one ($TX_{SP,i}$ with $i = 0$) is created in the channel create phase. This transaction is as follows:

$$\begin{aligned} TX_{SP,i}.Input &:= (TX_{CM,i}.txid\|1), \\ TX_{SP,i}.Output &:= (\theta_1, \theta_2, \dots), \\ TX_{SP,i}.Witness &:= ((1, \{\sigma_{TX_{SP,i}}^A, \sigma_{TX_{SP,i}}^B\})) \end{aligned} \quad (5.3)$$

The $TX_{SP,i}$ spends the main output of $TX_{CM,i}$ by meeting the sub-condition $pk_A \wedge pk_B \wedge t^+$.

- **Collateral transaction:** TX_{CL} locks the collateral of the watchtower on-chain and its output can be spent if A , B and W cooperate. The collateral value c equals $a + b$. If W uses the z^{th} output of a transaction with transaction identifier of $txid_W$ to lock c coins, the collateral transaction TX_{CL} would be defined as follows:

$$\begin{aligned} TX_{CL}.Input &:= (txid_W\|z), \\ TX_{CL}.Output &:= \{(c, pk_A \wedge pk_B \wedge pk_W)\}, \\ TX_{CL}.Witness &:= \{(1, \sigma_{TX_{CL}}^{W,1})\}. \end{aligned} \quad (5.4)$$

- **Reclaim transaction:** This transaction spends the output of TX_{CL} and its output can be spent by A , B and W if they cooperate or by W after a long relative time-lock period. The TX_{RC} is defined as follows:

$$\begin{aligned}
 TX_{RC}.Input &:= TX_{CL}.txid\|1), \\
 TX_{RC}.Output &:= \{(c, (pk_A \wedge pk_B \wedge pk_W) \vee (pk_W \wedge T^+))\}, \\
 TX_{RC}.Witness &:= \{(1, (\sigma_{TX_{RC}}^{A,1}, \sigma_{TX_{RC}}^{B,1}, \sigma_{TX_{RC}}^{W,1}))\}.
 \end{aligned} \tag{5.5}$$

The second sub-condition in the output is used by the watchtower to redeem its collateral after T rounds and terminate its service. However, as will be mentioned in the following sections, the first sub-condition is used to penalise the unresponsive watchtower.

Fig. 5.4 summarises the channel create phase. Section 5.7 provides details of the corresponding protocol.

5.4.2 Update

Assume that an FPPW channel is in state i with $i \geq 0$ and channel parties decide to update it from state i to $i + 1$. This is performed in two sub-phases. In the first sub-phase, channel parties create a new commit transaction and a new split transaction for the new state. However, to avoid any hostage situation, they sign the split transaction before signing the commit transaction. In the second sub-phase, channel parties revoke the previous state by signing one revocation and two penalty transactions. At most one out of these three transactions might be published on-chain upon fraud (i.e. upon broadcast of the revoked commit transaction). While the revocation transaction might be used to penalise the cheating channel party, penalty transactions might be utilised for punishing the dishonest watchtower.

The revocation transaction is the only transaction that spends both outputs of the revoked commit transaction using their non-time-locked sub-conditions $pk_A \wedge pk_B \wedge pk_W$. Thus, once a dishonest channel party publishes the revoked commit transaction, the watchtower or the counterparty can immediately publish the revocation transaction. It invalidates both penalty transactions because they also spend the auxiliary output of the revoked commit transaction. The single output of the revocation transaction is spendable by someone who knows the witness value y of both channel parties (i.e. party A can claim it if party B has published the revoked commit transaction and vice versa).

Now assume that a dishonest channel party publishes the revoked commit transaction but the watchtower does not react in time. Then the dishonest channel party might also




Figure 5.4: A summary of FPPW Channel Create.

publish the corresponding split transaction after t rounds. This spends the main output of the revoked commit transaction and invalidates the revocation transaction. However, since the honest channel party go offline for at most $T - 1$ rounds, it gets online when the watchtower has not completed reclaiming its collateral yet (i.e. the watchtower has not broadcast the reclaim transaction or has not spent its output yet). Thus, the honest party can publish one of two penalty transactions. Both penalty transactions spend the auxiliary output of the revoked commit transaction as well as the output of collateral and reclaim transactions, respectively. Similar to the revocation transaction, only the honest cheated party can claim the output of the published penalty transaction.

The introduced transactions will be explained further below:

- **Revocation transaction:** When parties A and B want to revoke $\text{TX}_{\text{CM},i}$, each channel participant (A , B and W) generates all the required signatures for the revocation transaction $\text{TX}_{\text{RV},i}$ and sends the signatures to other two participants. $\text{TX}_{\text{RV},i}$ is as follows:

$$\begin{aligned} \text{TX}_{\text{RV},i}.\text{Input} &:= (\text{TX}_{\text{CM},i}.\text{txid}\|1, \text{TX}_{\text{CM},i}.\text{txid}\|2), \\ \text{TX}_{\text{RV},i}.\text{Output} &:= (a + b + \epsilon, Y_{A,i} \wedge Y_{B,i}), \\ \text{TX}_{\text{RV},i}.\text{Witness} &:= ((2, (\sigma_{\text{TX}_{\text{RV},i}}^{A,1}, \sigma_{\text{TX}_{\text{RV},i}}^{B,1}, \sigma_{\text{TX}_{\text{RV},i}}^{W,1})), (2, (\sigma_{\text{TX}_{\text{RV},i}}^{A,2}, \sigma_{\text{TX}_{\text{RV},i}}^{B,2}, \sigma_{\text{TX}_{\text{RV},i}}^{W,2}))) \end{aligned} \quad (5.6)$$

The $\text{TX}_{\text{RV},i}$ spends both outputs of $\text{TX}_{\text{CM},i}$ using the non-time-locked sub-condition $pk_A \wedge pk_B \wedge pk_W$ and sends all the channel funds to output with condition $Y_{A,i} \wedge Y_{B,i}$. When a dishonest party, let's say A , publishes the revoked $\text{TX}_{\text{CM},i}$, A must wait for t rounds before being able to publish $\text{TX}_{\text{SP},i}$. However, W or B can immediately publish $\text{TX}_{\text{RV},i}$. Since $\text{TX}_{\text{CM},i}$ has been published by A , party B can obtain $y_{A,i}$. Thus, only party B who knows both $y_{A,i}$ and $y_{B,i}$ will own all the channel funds.

- **Penalty transaction 1:** The penalty transaction 1 $\text{TX}_{\text{PN}_1,i}$ is used to penalise W , given that a dishonest party publishes $\text{TX}_{\text{CM},i}$ and spends its main output using $\text{TX}_{\text{SP},i}$. The $\text{TX}_{\text{PN}_1,i}$ is defined as follows:

$$\begin{aligned} \text{TX}_{\text{PN}_1,i}.\text{Input} &:= (\text{TX}_{\text{CM},i}.\text{txid}, \text{TX}_{\text{CL},i}.\text{txid}\|1), \\ \text{TX}_{\text{PN}_1,i}.\text{Output} &:= (c + \epsilon, Y_{A,i} \wedge Y_{B,i}), \\ \text{TX}_{\text{PN}_1,i}.\text{Witness} &:= (W_1, (1, (\sigma_{\text{TX}_{\text{PN}_1,i}}^{A,2}, \sigma_{\text{TX}_{\text{PN}_1,i}}^{B,2}, \sigma_{\text{TX}_{\text{PN}_1,i}}^{W,2}))) \end{aligned} \quad (5.7)$$

where $W_1 := (1, (\sigma_{\text{TX}_{\text{PN}_1,i}}^{B,1}, y_{A,i}))$ given that broadcaster of $\text{TX}_{\text{CM},i}$ is A or $W_1 := (3, (\sigma_{\text{TX}_{\text{PN}_1,i}}^{A,1}, y_{B,i}))$ otherwise. When parties want to revoke $\text{TX}_{\text{CM},i}$, A and W (B and W) compute the required signatures for the second input of $\text{TX}_{\text{PN}_1,i}$ and send the signatures to B (A). Now assume that one party, let's say A , publishes the revoked $\text{TX}_{\text{CM},i}$ and spends its main output after t rounds. Then, B obtains $y_{A,i}$ and hence can add the witness W_1 to $\text{TX}_{\text{PN}_1,i}$ and publish it, given that $\text{TX}_{\text{CL},i}.\text{Output}$ is still unspent. Then, the transaction $\text{TX}_{\text{PN}_1,i}$ spends the second output of $\text{TX}_{\text{CM},i}$ using the time-locked sub-condition $pk_B \wedge Y_{A,i} \wedge t^+$ as well as the output of the collateral transaction. Only B can claim output of $\text{TX}_{\text{PN}_1,i}$. A similar scenario occurs if B is the broadcaster of $\text{TX}_{\text{CM},i}$.

- **Penalty transaction 2:** There exists one penalty transaction 2 $\text{TX}_{\text{PN}_2,i}$ per state. It is exactly the same as $\text{TX}_{\text{PN}_1,i}$, with the only difference that it spends $\text{TX}_{\text{RC},i}.\text{Output}$

(rather than $\text{TX}_{\text{CL}}.\text{Output}$) using the sub-condition $pk_A \wedge pk_B \wedge pk_W$. Thus, it is useful for cases where the watchtower does not react upon fraud but by publishing TX_{RC} tries to reclaim its collateral. However, since the honest party goes offline for at most $T - 1$ rounds, it gets online when $\text{TX}_{\text{RC}}.\text{Output}$ is still unspent. Thus, the honest party can add the required witness to $[\text{TX}_{\text{PN}_2,i}]$ and publish it. The $\text{TX}_{\text{PN}_2,i}$ is defined as follows:

$$\begin{aligned} \text{TX}_{\text{PN}_2,i}.\text{Input} &:= (\text{TX}_{\text{CM},i}.\text{txid}, \text{TX}_{\text{RC}}.\text{txid} \| 1), \\ \text{TX}_{\text{PN}_2,i}.\text{Output} &:= (c + \epsilon, Y_{A,i} \wedge Y_{B,i}), \\ \text{TX}_{\text{PN}_2,i}.\text{Witness} &:= (W_1, (1, (\sigma_{\text{TX}_{\text{PN}_2,i}}^{A,2}, \sigma_{\text{TX}_{\text{PN}_2,i}}^{B,2}, \sigma_{\text{TX}_{\text{PN}_2,i}}^{W,2}))) \end{aligned} \quad (5.8)$$

where $W_1 := (1, (\sigma_{\text{TX}_{\text{PN}_2,i}}^{B,1}, y_{A,i}))$ given that broadcaster of $\text{TX}_{\text{CM},i}$ is A or $W_1 := (3, (\sigma_{\text{TX}_{\text{PN}_2,i}}^{A,1}, y_{B,i}))$ otherwise..

Fig. 5.5 summarises the channel update phase. Section 5.7 provides details of the corresponding protocol.

Remark 5.3. *Watchtower is actively involved in steps 6 and 7 of the channel update phase (See Fig. 5.5). Therefore, this phase fails to complete if the watchtower is unavailable. Section 5.8 introduces an update protocol for such scenarios.*

5.4.3 Close

Assume that the channel parties A and B have updated their channel n times and then A and/or B decide to close it. They can close the channel cooperatively. To do so, A and B create a new transaction, called modified split transaction $\text{TX}_{\overline{\text{SP}}}$, and publish it on-chain. The $\text{TX}_{\overline{\text{SP}}}$ is defined as follows:

$$\begin{aligned} \text{TX}_{\overline{\text{SP}}}.\text{Input} &:= (\text{TX}_{\text{FU},i}.\text{txid} \| 1), \\ \text{TX}_{\overline{\text{SP}}}.\text{Output} &:= (\theta_1, \theta_2, \dots), \\ \text{TX}_{\overline{\text{SP}}}.\text{Witness} &:= ((1, (\sigma_{\text{TX}_{\overline{\text{SP}}}}^A, \sigma_{\text{TX}_{\overline{\text{SP}}}}^B))) \end{aligned} \quad (5.9)$$

Outputs of this transaction might be similar to those for $\text{TX}_{\text{SP},n}$. Note that the value of auxiliary output of $\text{TX}_{\text{CM},n}$ (ϵ) can also be given to A and B ($\epsilon/2$ each) through outputs of $\text{TX}_{\overline{\text{SP}}}$. If one of the channel parties gets unresponsive, its counterparty can still close the channel non-collaboratively by publishing $\text{TX}_{\text{CM},n}$ and then $\text{TX}_{\text{SP},n}$ on-chain. The collaborative and non-collaborative channel closure protocols can be found in Section 5.7.




Figure 5.5: FPPW Channel Update.

5.4.4 React

It is always possible that a channel party publishes a revoked commit transaction $TX_{CM,i}$ on-chain. Then, the watchtower or the counterparty publishes the corresponding revocation transaction within $t - 1$ rounds. Only the honest counterparty can claim the output of the revocation transaction. If the watchtower is unresponsive and the honest party is offline, a malicious party can publish a revoked commit transaction $TX_{CM,i}$ with $i < n$ and its corresponding split transaction $TX_{SP,i}$ on-chain. Then the honest party, who gets online once every $T - 1$ rounds, can penalise the unresponsive watchtower by

publishing either $\text{TX}_{\text{PN}_1,i}$ or $\text{TX}_{\text{PN}_2,i}$. Protocols for these scenarios can be found in Section 5.7.

5.4.5 Watchtower Terminate

In this phase, W decides to terminate its employment. To do that, W publishes TX_{RC} and spends its output after T rounds. Since A and B do not go offline for more than $T - 1$ rounds, they get online during this T -round interval and observe that TX_{RC} is on the chain. Then parties can close the channel and open a new one with a new watchtower. Parties can also continue using this channel without any watchtower. To do so, channel parties must check the blockchain at least once every $t - 1$ rounds to prevent fraudulent channel closures. New channel updates can be performed according to the Generalized channels [14] or its new variant introduced in Section 5.3.2.1.

5.5 Security Analysis

In this section, we analyse privacy and coverage of FPPW protocol through Theorems 5.1 and 5.2, respectively. Then, we analyse the fairness of FPPW through Theorem 5.3.

Theorem 5.1. *FPPW provides weak privacy against watchtower as defined in Definition 4.6.*

Proof. Assume that the conditions mentioned in the two-stage privacy game (see Section 4.3.2) are satisfied. By observing different steps and transactions of the protocol, one can see that only split transactions contain information on $x_{A,i}$ and $x_{B,i}$ with $i \in [1, n - 1]$. However, these transactions are never published on-chain or sent to the watchtower or any external entity. Other transactions in the protocol contain no information regarding $x_{A,i}$ or $x_{B,i}$ with $i \in [1, n - 1]$. Note that the monetary value of outputs of $\text{TX}_{\text{CM},i}$, $\text{TX}_{\text{RV},i}$, $\text{TX}_{\text{PN}_1,i}$, $\text{TX}_{\text{PN}_2,i}$, TX_{CL} , and TX_{RC} of the first payment channels are the same as those for the second one. Furthermore, TX_{FU} , $\text{TX}_{\text{SP},n}$ or TX_{SP} contain no information regarding the i^{th} channel state with $i \in [1, n - 1]$. Thus, the view of any adversary $\mathcal{A}$ on $(\mathbf{x}_{A,[1,n-1]}, \mathbf{x}_{B,[1,n-1]})$ is indistinguishable from its view on $(\mathbf{x}_{A',[1,n-1]}, \mathbf{x}_{B',[1,n-1]})$. $\square$

Theorem 5.2. *FPPW provides β -coverage with $\beta = 1/2$ based on Definition 4.11.*

Proof. Assume that we have N payment channels, with channel capacities $X_i = a_i + b_i$, $i \in [1, N]$. Thus, the total capacity of the channels is $\mathcal{X} = \sum_{i=1}^N X_i$. Since the i^{th} channel collateral c_i equals $a_i + b_i$, the total watchtower collateral is $\mathcal{C} = \sum_{i=1}^N c_i = \sum_{i=1}^N X_i = \mathcal{X}$. Thus, we have $\beta = \frac{\mathcal{X}}{\mathcal{X} + \mathcal{C}} = 1/2$. $\square$

As mentioned earlier, Theorem 5.3 analyses the fairness of FPPW. Lemmas 5.1 and 5.2 are utilised to prove this theorem. They show how FPPW guarantees that funds of the honest channel party and the honest watchtower are safe in the channel. We first prove Lemmas 5.1 and 5.2 and then, considering these two lemmas, we prove Theorem 5.3.

Lemma 5.1. *Let Π be a EUF – CMA secure digital signature, $\mathcal{R}$ be a hard relation and Ξ be a secure adaptor digital signature. For an FPPW channel, assume that the honest channel party $P \in \{A, B\}$ checks the blockchain at the end of the channel creation phase and then gets online periodically with a period of at most $T - 1$ rounds. The probability that P loses any funds in the channel is negligible.*

Proof. Without loss of generality let $P = A$. Based on Lemma 5.4, A does not lose any funds with non-negligible probability unless a revoked commit transaction is published on-chain. However, at the end of the channel creation phase, no revoked commit transaction $\text{TX}_{\text{CM},i}$ exists to be published on-chain. Furthermore, the channel creation completes once TX_{CL} is recorded on-chain and since TX_{RC} spends the output of TX_{CL} , right after the end of the channel create phase, neither a revoked commit transaction nor the reclaim transaction is on-chain. Now according to assumptions, A checks the chain at the end of the channel create phase and goes offline for at most $T - 1$ rounds. The next time that A gets online, there will be 4 possibilities regarding the broadcast of a revoked $\text{TX}_{\text{CM},i}$ or TX_{RC} during the time interval when A has been offline:

1. Only TX_{RC} has been published on-chain,
2. Both a revoked $\text{TX}_{\text{CM},i}$ and TX_{RC} have been published,
3. Only a revoked $\text{TX}_{\text{CM},i}$ has been published, or
4. neither a revoked $\text{TX}_{\text{CM},i}$ nor TX_{RC} has been published.

For the first 3 possibilities, based on Lemmas 5.5, 5.6, and 5.7, the probability that A loses any funds is negligible. For the 4th possibility, the condition set is the same as the end of the channel create phase when neither a revoked commit transaction nor the reclaim transaction was on the blockchain. Therefore, since A will again go offline for at most $T - 1$ rounds, the whole process repeats. $\square$

Lemma 5.2. *Let Π be a EUF – CMA secure digital signature, $\mathcal{R}$ be a hard relation and Ξ be a secure adaptor digital signature. For an FPPW channel, assume that the honest watchtower W checks the blockchain at the end of the channel creation phase and then remains online. The probability that W loses any funds in the channel is negligible.*

Proof. Based on Lemma 5.8, W does not lose any funds with non-negligible probability unless a revoked commit transaction and then at least t rounds later its corresponding $TX_{PN_1,i}$ or $TX_{PN_2,i}$ are published. However, at the end of the channel create phase, no revoked commit transaction exists to be published yet. Also, following our assumptions, W remains online after the channel create phase. Now assume that a revoked $TX_{CM,i}$ is published through the block $\mathcal{B}_j$. The time-locked sub-conditions of outputs of $TX_{CM,i}$ cannot be met within $t - 1$ rounds. Also, meeting the non-time-locked sub-conditions of the first and second output of $TX_{CM,i}$ requires W 's signature and W does not grant such authorisations on any transaction other than $TX_{RV,i}$. Thus, once $TX_{CM,i}$ is published on the chain, due to our assumption regarding the security of the underlying digital signature, it is of negligible probability that any adversary $\mathcal{A}$ spends the first or second output of $TX_{CM,i}$ within $t - 1$ rounds using any transaction other than $TX_{RV,i}$. Furthermore, W has received $TX_{RV,i}$ (through step 6 of the channel update phase from state i to $i + 1$ in Fig. 5.5), before giving authorisation on the second input of $TX_{PN_1,i}$ or $TX_{PN_2,i}$ (through step 7 of the channel update phase from state i to $i + 1$ in Fig. 5.5). Thus, once $TX_{CM,i}$ is published, W can publish $TX_{RV,i}$ through one of the blocks $\mathcal{B}_{j+1} \dots \mathcal{B}_{j+t-1}$, meaning that the honest watchtower has at least $t - 1$ rounds time to publish $TX_{RV,i}$, which is enough based on our blockchain assumptions regarding the value of the confirmation delay. Following Lemma 5.9, it is of negligible probability that broadcast of $TX_{RV,i}$ causes the honest watchtower W to lose any funds in the channel. $\square$

Theorem 5.3. *Let Π be a EUF – CMA secure digital signature, $\mathcal{R}$ be a hard relation and Ξ be a secure adaptor digital signature. FPPW provides channel party α -fairness with $\alpha = 1$ and watchtower fairness as defined in Definition 4.9 and 4.10, respectively.*

Proof. The honest channel party always have at least one non-revoked commit transaction and its corresponding split transaction by broadcasting which she can close the channel. This proves that FPPW meets the first requirement of Definition 4.9. Furthermore, we know that based on FPPW protocol, the honest channel party checks the chain at least once every $T - 1$ round and according to Lemma 5.1, the probability that the honest channel party loses any funds in the channel is negligible. This proves that FPPW provides channel party α -fairness with $\alpha = 1$.

The watchtower in FPPW is paid for channel creation and each channel update and hence her reward amount is non-zero. Also, we know that based on FPPW protocol, the honest watchtower always remains online and according to Lemma 5.2, the probability that such an honest watchtower loses any funds in the channel is negligible. Additionally, the watchtower can publish TX_{RC} at any time and redeem her collateral after T rounds. Thus, FPPW meets both requirements of Definition 4.10. $\square$

Lemma 5.3. *Let Π be a EUF – CMA secure digital signature, $\mathcal{R}$ be a hard relation and Ξ be a secure adaptor digital signature. For an FPPW channel, it is of negligible probability that broadcast of $\text{TX}_{\text{RV},i}$, $\text{TX}_{\text{PN}_1,i}$ or $\text{TX}_{\text{PN}_2,i}$ causes the honest channel party $P \in \{A, B\}$ to lose any funds in the channel.*

Proof. Without loss of generality let $P = A$. The transaction $\text{TX}_{\text{RV},i}$ spends the main output of the revoked $\text{TX}_{\text{CM},i}$ and hence cannot be published unless $\text{TX}_{\text{CM},i}$ is on-chain. Based on the protocol, the honest party A never broadcasts the revoked $\text{TX}_{\text{CM},i}$ on-chain. The party A only creates the pre-signature $\tilde{\sigma}_{\text{CM},i}$ on the transaction $\text{TX}_{\text{CM},i}$. Thus, if $\text{TX}_{\text{CM},i}$ is published, the probability that A fails to obtain $y_{B,i}$ is negligible. Otherwise, a EUF – CMA security or witness extractability of the used adaptor signature is violated. Furthermore, $\text{TX}_{\text{RV},i}$ has only one output with the condition of $Y_{A,i} \wedge Y_{B,i}$ and the value of $a+b+\epsilon$. Since A privately preserves its witness value $y_{A,i}$, the probability that any PPT adversary claims $\text{TX}_{\text{RV},i}.\text{Output}$ is negligible. Otherwise, the utilised hard relation would break. Therefore, it is of negligible probability that A (who knows both $y_{A,i}$ and $y_{B,i}$) fails to claim $\text{TX}_{\text{RV},i}.\text{Output}$.

Also transactions $\text{TX}_{\text{PN}_1,i}$ and $\text{TX}_{\text{PN}_2,i}$ spend the auxiliary output of $\text{TX}_{\text{CM},i}$ as well as output of TX_{CL} and TX_{RC} , respectively. These transactions have one output with the condition of $Y_{A,i} \wedge Y_{B,i}$ and the value of $c + \epsilon = a + b + \epsilon$. Since output condition for $\text{TX}_{\text{PN}_1,i}$ and $\text{TX}_{\text{PN}_2,i}$ is the same as that of $\text{TX}_{\text{RV},i}$, the proof for $\text{TX}_{\text{PN}_1,i}$ and $\text{TX}_{\text{PN}_2,i}$ is also similar to that of $\text{TX}_{\text{RV},i}$. $\square$

Remark 5.4. *Before stating the next Lemma, it must be noted that while the channel update phase from state i to $i + 1$ is incomplete yet, it is possible that $\text{TX}_{\text{CM},i+1}$ is published on-chain and t rounds later $\text{TX}_{\text{SP},i+1}$ is also broadcast. In such situations, we assume that the channel party does not lose any funds in the channel even if her counterparty's balance in state $i + 1$ is larger than that of state i . In other words, it is assumed that during steps 6 and 7 of the channel update phase (see Fig. 5.5), there are two valid channel states where channel closure with each one does not cause the honest party to lose any funds in the channel. A similar assumption is also made for other payment channels of type replace by revocation [15, 14].*

Lemma 5.4. *Let Π be a EUF – CMA secure digital signature, $\mathcal{R}$ be a hard relation and Ξ be a secure adaptor digital signature. For an FPPW channel with n channel updates, it is of negligible probability that the honest party $P \in \{A, B\}$ loses any funds using any scenario other than the broadcast of $\text{TX}_{\text{CM},i}$ with $i < n$.*

Proof. Without loss of generality let $P = A$. Funds of A are locked in $\text{TX}_{\text{FU}}.\text{Output}$. It is of negligible probability that any PPT adversary $\mathcal{A}$ spends the output of TX_{FU} without the honest party A 's authorisation. Otherwise, the underlying digital signature would be forgeable. Furthermore, $\text{TX}_{\text{SP}}, \text{TX}_{\text{CM},i}$ with $i = [0, n-1]$, $\text{TX}_{\text{CM},n}$ and possibly $\text{TX}_{\text{CM},n+1}$ (given

that channel update phase from state n to $n + 1$ has started) are the only transactions in the protocol that spend the output of TX_{FU} and A grants authorisation for. Thus, these transactions will be discussed further to see how each one of them can cause the honest party A to be cheated out of its funds.

Since TX_{SP} represents the final agreed state of the channel, its broadcast cannot cause A to lose any funds in the channel. Additionally, since both sub-conditions in the main output of $\text{TX}_{\text{CM},n}$ include pk_A , due to our assumption regarding the security of the underlying digital signature, it is of negligible probability that the main output of $\text{TX}_{\text{CM},n}$ is spent without A 's authorisation. Since $\text{TX}_{\text{SP},n}$ is the only transaction in the protocol that spends the main output of $\text{TX}_{\text{CM},n}$ and A grants authorisation for, the probability of spending the main output of $\text{TX}_{\text{CM},n}$ using any transaction other than $\text{TX}_{\text{SP},n}$ is negligible. However, since $\text{TX}_{\text{SP},n}$ also represents the final state of the channel, its broadcast cannot cause A to lose any funds. According to Remark 5.4, similar statements can be stated for $\text{TX}_{\text{CM},n+1}$ and $\text{TX}_{\text{SP},n+1}$. Thus, cheating the honest party A using any scenario other than broadcast of $\text{TX}_{\text{CM},i}$ with $i < n$ is of negligible probability.

□

Case 1. Let there exist an FPPW channel with n channel updates where $n \geq 0$. Assume that the honest channel party $P \in \{A, B\}$ gets online when the last published block on the blockchain is $\mathcal{B}_j$. Party P observes that $\text{TX}_{\text{CM},i}$ is unpublished but TX_{RC} has been published through the block $\mathcal{B}_k$ with $k \leq j$.

Lemma 5.5. *Let Π be a EUF – CMA secure digital signature, $\mathcal{R}$ be a hard relation and Ξ be a secure adaptor digital signature. If conditions of Case 1 are satisfied, the probability that P loses any funds in the channel is negligible.*

Proof. Without loss of generality let $P = A$. If party A checks the blockchain and observes that $\text{TX}_{\text{CM},i}$ is unpublished but TX_{RC} has been published, A can get online periodically with a period of at most $t - 1$ rounds. Based on Lemma 5.4, if A is going to lose some funds in the channel with non-negligible probability, the PPT adversary $\mathcal{A}$ must publish $\text{TX}_{\text{CM},i}$ with $i < n$ through the block $\mathcal{B}_k$ with $k > j$. It must be noted that since the next time that A gets online again, $\mathcal{B}_{j+t-1}$ is the latest block on the blockchain, if we have $k > j + t - 1$, all the conditions mentioned for Case 1 repeat with j being replaced with $j + t - 1$. Thus, we assume that we have $j < k \leq j + t - 1$.

time-locked sub-conditions of outputs of $\text{TX}_{\text{CM},i}$ cannot be met within $t - 1$ rounds. Also, their non-time-locked sub-conditions include pk_A . Therefore, it is of negligible probability that any adversary spends the first or the second output of $\text{TX}_{\text{CM},i}$ through one of the blocks $\mathcal{B}_{k+1}, \dots, \mathcal{B}_{k+t-1}$ without the honest party A 's authorisation. Otherwise, the underlying digital signature would break. However, A never grants such authorisations

on a transaction other than $\text{TX}_{\text{RV},i}$. Furthermore, the honest party A has created $\text{TX}_{\text{RV},i}$ through step 6 of the channel update phase from state i to $i + 1$ (See Fig. 7). When A gets online the last block on the blockchain is $\mathcal{B}_{j+t-1}$. Thus, A is able to publish $\text{TX}_{\text{RV},i}$ through one of the blocks $\mathcal{B}_{j+t}, \dots, \mathcal{B}_{k+t-1}$. Since we have $k - j \geq 1$, A always have at least 1 block time to publish $\text{TX}_{\text{RV},i}$, which is enough based on our blockchain assumption regarding the value of the confirmation delay. Also, according to Lemma 5.3, it is of negligible probability that broadcast of $\text{TX}_{\text{RV},i}$ causes the honest party A to lose any funds in the channel. $\square$

Case 2. Let there exist an FPPW channel with n channel updates where $n \geq 0$. Assume that the honest channel party $P \in \{A, B\}$ gets online when the last published block on the blockchain is $\mathcal{B}_j$. Party P observes that $\text{TX}_{\text{CM},i}$ with $i < n$ and TX_{RC} have been published on-chain through the blocks $\mathcal{B}_k$ and $\mathcal{B}_l$ respectively with $k, l \leq j$ and $j + 1 < l + T$.

Lemma 5.6. *Let Π be a EUF – CMA secure digital signature, $\mathcal{R}$ be a hard relation and Ξ be a secure adaptor digital signature. If conditions of Case 2 are satisfied, the probability that P loses any funds in the channel is negligible.*

Proof. Without loss of generality let $P = A$. All sub-conditions of auxiliary output of $\text{TX}_{\text{CM},i}$ include either pk_A or $Y_{A,i}$. Thus, it is of negligible probability that any PPT adversary $\mathcal{A}$ spends this output of $\text{TX}_{\text{CM},i}$ without A 's authorisation. Otherwise, the underlying digital signature or the hard relation would break. The channel party A grants such an authorisation only on $\text{TX}_{\text{RV},i}$. However, based on Lemma 5.3, if $\text{TX}_{\text{RV},i}$ is published, the probability that A loses any funds is negligible. Thus, we assume that when A gets online, $\text{TX}_{\text{RV},i}$ is unpublished and hence the auxiliary output of $\text{TX}_{\text{CM},i}$ is unspent yet. According to values of j and k two categories of cases are possible. We firstly consider cases with $j + 1 \geq k + t$ and prove that in such cases A can publish $\text{TX}_{\text{PN}_2,i}$. Then, we show that if $j + 1 < k + t$, A can still publish $\text{TX}_{\text{RV},i}$. Thus, according to Lemma 5.3, for all cases, the probability that A loses any funds is negligible.

Consider cases with $j + 1 \geq k + t$. The non-time-locked sub-condition of TX_{RC} .Output includes pk_A and hence due to our assumption regarding the security of the underlying digital signature, it is of negligible probability that the output of TX_{RC} is spent without A 's authorisation through the block $\mathcal{B}_m$ with $m < l + T$. Also A grants this authorisation only on $\text{TX}_{\text{PN}_2,i}$. Actually, all channel participants grant such authorisation on $\text{TX}_{\text{PN}_2,i}$. Additionally, since $\text{TX}_{\text{CM},i}$ is on-chain, the probability that A fails to obtain $y_{B,i}$ and hence fails to generate the required signatures for the first input of $\text{TX}_{\text{PN}_2,i}$ is negligible. Otherwise, aEUF – CMA security or extractability of the used adaptor signature is violated. Thus, party A might publish $\text{TX}_{\text{PN}_2,i}$ on-chain through one of the blocks $\mathcal{B}_{j+1}, \dots, \mathcal{B}_{l+T-1}$ and since based on assumptions of Case 2 we have $j + 1 \leq l + T - 1$, the honest party A

will have at least $l + T - 1 - j \geq 1$ rounds time to publish $\text{TX}_{\text{PN}_2,i}$ which is enough based on our assumptions regarding the value of confirmation delay.

Now let $j + 1 < k + t$. We know that without having the honest party A 's authorisation, it is of negligible probability that any PPT adversary is able to spend the first or the second output of $\text{TX}_{\text{CM},i}$ through one of the blocks $\mathcal{B}_{k+1}, \dots, \mathcal{B}_{k+t-1}$. All the channel participants grant such authorisations only on $\text{TX}_{\text{RV},i}$ and A has created this transaction in step 6 of the channel update phase from state i to $i + 1$ (See Fig. 5.5). Thus, A can publish it through one of the blocks $\mathcal{B}_{j+1}, \dots, \mathcal{B}_{k+t-1}$ and since $j + 1 < k + t$, A will have at least $k + t - 1 - j \geq 1$ rounds time to publish $\text{TX}_{\text{RV},i}$.

Therefore, for all cases, A can publish either $\text{TX}_{\text{RV},i}$ or $\text{TX}_{\text{PN}_2,i}$ and hence according to Lemma 5.3, the probability that A loses any funds is negligible. □

Case 3. Let there exist an FPPW channel with n channel updates where $n \geq 0$. Assume that the honest channel party $P \in \{A, B\}$ gets online when the last published block on the blockchain is $\mathcal{B}_j$. Party P observes that $\text{TX}_{\text{CM},i}$ has been published on-chain through the block $\mathcal{B}_k$ with $k \leq j$ but TX_{RC} is unpublished.

Lemma 5.7. *Let Π be a EUF – CMA secure digital signature, $\mathcal{R}$ be a hard relation and Ξ be a secure adaptor digital signature. If conditions of Case 3 are satisfied, the probability that P loses any funds in the channel is negligible.*

Proof. Without loss of generality let $P = A$. Similar to the proof of Lemma 5.6, we assume that $\text{TX}_{\text{RV},i}$ is unpublished because otherwise, the probability that A loses some funds is negligible. According to proof of Lemma 5.6, if we have $j + 1 < k + t$, A will have at least $k + t - 1 - j \geq 1$ rounds time to publish $\text{TX}_{\text{RV},i}$.

Now let $j + 1 \geq k + t$. We know that due to the security of the underlying digital signature, it is of negligible probability that someone spends TX_{CL} . Output without A 's authorisation and A grants this authorisation only on TX_{RC} and $\text{TX}_{\text{PN}_1,i}$. Also, since $\text{TX}_{\text{CM},i}$ is on-chain, the probability that A fails to obtain $y_{B,i}$ is negligible. Thus, it is of negligible probability that party A cannot generate the required signatures for the first input of $\text{TX}_{\text{PN}_1,i}$. Then, party A can publish $\text{TX}_{\text{PN}_1,i}$ on-chain through one of the blocks $\mathcal{B}_{j+1}, \mathcal{B}_{j+2}, \dots$. Based on Lemma 5.3, it is of negligible probability that broadcast of $\text{TX}_{\text{PN}_1,i}$ can cause A to lose any funds.

If before the broadcast of $\text{TX}_{\text{PN}_1,i}$, TX_{RC} is published by the watchtower, conditions of Case 2 get satisfied and due to Lemma 5.6, we know that it is of negligible probability that A loses any funds in the channel. □

Lemma 5.8. *Let Π be a EUF – CMA secure digital signature, $\mathcal{R}$ be a hard relation and Ξ be a secure adaptor digital signature. For an FPPW channel with n channel updates, an honest watchtower W does not lose any funds with non-negligible probability unless first a revoked commit transaction $TX_{CM,i}$ with $i < n$ is broadcast on the chain and at least t rounds later, either $TX_{PN_1,i}$ or $TX_{PN_2,i}$ is also published on the chain.*

Proof. Assume that the channel update phase has completed n times with $n \geq 0$. We discuss different scenarios using which the honest watchtower can be cheated. Funds of the watchtower are locked in $TX_{CL}.$ Output and condition of $TX_{CL}.$ Output includes pk_W . Thus, due to security of the underlying digital signature, it is of negligible probability that any PPT adversary $\mathcal{A}$ spends $TX_{CL}.$ Output without W 's authorisation. TX_{RC} and $TX_{PN_1,i}$ with $i = [0, n-1]$ are the only transactions in the protocol that spend $TX_{CL}.$ Output and W grants authorisation for. Both sub-conditions of $TX_{RC}.$ Output include pk_W . Thus, due to the security of the underlying digital signature, it is of negligible probability that any PPT adversary $\mathcal{A}$ spends $TX_{RC}.$ Output without W 's authorisation and $TX_{PN_2,i}$ with $i = [0, n-1]$ are only transactions in the protocol that spend $TX_{RC}.$ Output and W grants authorisation for. Therefore, all scenarios with non-negligible probability lead to the broadcast of $TX_{PN_1,i}$ or $TX_{PN_2,i}$. Since both $TX_{PN_1,i}$ and $TX_{PN_2,i}$ take the auxiliary output of $TX_{CM,i}$ as their first inputs, those transactions can only be published if $TX_{CM,i}$ is on-chain. Meeting the non-time-locked sub-condition of the auxiliary output of $TX_{CM,i}$ requires W 's authorisation and W does not grant such an authorisation on $TX_{PN_1,i}$ or $TX_{PN_2,i}$. Therefore, it is of negligible probability that any PPT adversary publishes $TX_{PN_1,i}$ or $TX_{PN_2,i}$ before t rounds being elapsed since the broadcast of $TX_{CM,i}$. $\square$

Lemma 5.9. *Let Π be a EUF – CMA secure digital signature, $\mathcal{R}$ be a hard relation and Ξ be a secure adaptor digital signature. For an FPPW channel with n channel updates, if $TX_{CM,i}$ with $i \in [0, n-1]$ is published on-chain, it is of negligible probability that broadcast of $TX_{RV,i}$ causes the honest watchtower W to lose any funds in the channel.*

Proof. Based on Lemma 5.8, the probability of cheating the watchtower without publishing $TX_{PN_1,i}$ or $TX_{PN_2,i}$ is negligible. However, since $TX_{RV,i}$, $TX_{PN_1,i}$ and $TX_{PN_2,i}$ spend auxiliary output of $TX_{CM,i}$, if $TX_{RV,i}$ is published on-chain, $TX_{PN_1,i}$ and $TX_{PN_2,i}$ cannot be recorded on-chain anymore. $\square$

5.6 Fee Handling

Once a revoked commit transaction is recorded on the blockchain, the watchtower must record its corresponding revocation transaction within $t - 1$ rounds. Otherwise, the watchtower might be penalised. However, the time it takes for a transaction to be

recorded on the blockchain depends on its fee value and the network congestion. The body of a revocation transaction is created during the channel update phase but it might be broadcast in the blockchain network later upon fraud. Thus, the fee amount must be large enough to ensure the watchtower that the revocation transaction will be accepted by miners within the dispute period. In other words, when channel participants are creating a revocation transaction, they must assume that the blockchain network will be highly congested at the time when fraud will occur.

An alternative approach involves manipulating the SIGHASH type parameter to accommodate an increase in the fee amount during congestion periods. SIGHASH is a component of a transaction that specifies which parts of the transaction data is signed by the participants. By specifying a specific SIGHASH type, participants can determine which parts of the transaction are included in the signature, and which parts can be modified without invalidating the signature. There are different SIGHASH types, each denoted by a specific number, and they serve various purposes [57]. For our use case, we propose the usage of SIGHASH of type 0x81 (SIGHASH_ALL | SIGHASH_ANYONECANPAY) for channel parties' signatures for both inputs of revocation transactions. Thus, a signature for each input applies to that input and the output. Therefore, when due to network congestion the considered fee for the revocation transaction is low, the watchtower can add some inputs to the revocation transaction to increase the fee amount, sign all inputs using SIGHASH of type 0x01 (SIGHASH_ALL) and submit it to the network. If there exists enough time, the watchtower might even repeat this process several times and raise this extra fee each time until one of the revocation transactions is accepted by the miners. This method can be used if revocation transactions are only held by the watchtower (i.e. if channel parties do not receive signatures of the watchtower on revocation transactions during the channel update phase).

A similar approach can also be used for penalty transactions. Channel parties and the watchtower can use SIGHASH of type 0x02 (SIGHASH_NONE) for the second input of penalty transactions. Then, signatures apply only to all inputs of penalty transactions. In this way, the watchtower can be certain that a penalty transaction cannot be published unless its corresponding commit transaction is on-chain. However, if a revoked $TX_{CM,i}$ is published by a channel party, let's say A , and its main output is spent by $TX_{SP,i}$, party B has the opportunity to set the output value of the penalty transaction according to the network congestion and sign the corresponding penalty transaction (to meet the sub-condition $pk_B \wedge Y_{A,i} \wedge t^+$) using SIGHASH of type 0x01 (SIGHASH_ALL). In this way, B can reduce the output value if the current fee is low and this difference value is used as the extra fee amount. If there exists enough time, B can even repeat this process multiple times, each time with a higher fee until one penalty transaction is recorded on-chain.

5.7 FPPW Protocol

In this section, protocols for different phases of FPPW will be presented. In different steps of the protocol, channel participants generate (or verify) some signatures or pre-signatures on protocol transactions. When a signature or pre-signature is going to be generated (or verified) for j^{th} input of the transaction TX_i , the input message to the signing (or verification) algorithm is denoted by $f([\text{TX}_i], j)$ [58].

Remark 5.5. *The FPPW protocol which is presented in this Section is slightly different from what was explained in Section 5.4. Those differences are as follows. In the channel create phase, $[\text{TX}_{\text{RV},0}]$ is also created and W 's signatures on this transaction are given to both A and B . This transaction lacks the channel parties' signatures. Those signatures are computed in the channel update phase when $\text{TX}_{\text{CM},0}$ is going to be revoked. However, A and B require W 's signature in this phase to be able to continue using the channel given that W will be non-cooperative or temporarily unavailable in the channel update phase from state 0 to 1. Otherwise, A and B will have to update the channel on-chain because $\text{TX}_{\text{CM},0}$ will be irrevocable. Similarly, in the update channel phase from state i to $i+1$, W also generates its signatures for both inputs of $\text{TX}_{\text{RV},i+1}$ and sends them to A and B . This will allow A and B to update the channel from state $i+1$ to $i+2$ and revoke $\text{TX}_{\text{CM},i+1}$ even if W will be temporarily unavailable or uncooperative. Otherwise, $\text{TX}_{\text{CM},i+1}$ will be irrevocable.*

FPPW channel create protocol is as follows:

Preconditions: A , B and W own $a+\epsilon/2$, $b+\epsilon/2$ and $c = a+b$ coins on-chain in output of transactions with transaction identifiers txid_A , txid_B and txid_W respectively. A , B and W know each other's public keys, pk_A , pk_B and pk_W and values of ϵ , a , b and c that will be used in the channel.

1. Create $[\text{TX}_{\text{FU}}]$:

- (a) $P \in \{A, B\} \xrightarrow{\text{txid}_P} \bar{P}, W$
- (b) If P receives $\text{txid}_{\bar{P}}$, it creates $[\text{TX}_{\text{FU}}]$ according to 5.1. Else it stops.
- (c) If W receives txid_P with $P = \{A, B\}$, it creates $[\text{TX}_{\text{FU}}]$ according to 5.1. Else it stops.

2. Create $[\text{TX}_{\text{CM},0}]$:

- (a) $P \in \{A, B\}$ generates $(Y_{P,0}, y_{P,0}) \leftarrow \text{GenR}$.
- (b) $P \xrightarrow{Y_{P,0}} \bar{P}, W$
- (c) If P receives $Y_{\bar{P},0}$, it creates $[\text{TX}_{\text{CM},0}]$ according to 5.2. Else it stops.

- (d) If W receives $Y_{P,0}$ with $P = \{A, B\}$, it creates $[TX_{CM,0}]$ according to 5.2. Else it stops.
3. Create $[TX_{SP,0}]$: Party $P \in \{A, B\}$ creates $[TX_{SP,0}]$ according to 5.3.
4. Create $[TX_{RV,0}]$:
- (a) W creates $[TX_{RV,0}]$ according to 5.6.
- (b) W computes $\sigma_{TX_{RV,0}}^{W,j} = \text{Sign}_{sk_W}(f([TX_{RV,0}], j))$ with $j = \{1, 2\}$.
- (c) $W \xrightarrow{\sigma_{TX_{RV,0}}^{W,1}, \sigma_{TX_{RV,0}}^{W,2}} A, B$.
- (d) If party $P \in \{A, B\}$ receives $\sigma_{TX_{RV,0}}^{W,1}$ and $\sigma_{TX_{RV,0}}^{W,2}$ from W s.t. $\text{Vrfy}_{pk_W}(f([TX_{RV,0}], j); \sigma_{TX_{RV,0}}^{W,j}) = 1$ with $j = \{1, 2\}$, it continues. Else it stops.
5. Create $TX_{SP,0}$:
- (a) Party $P \in \{A, B\}$ computes $\sigma_{TX_{SP,0}}^P = \text{Sign}_{sk_P}(f([TX_{SP,0}], 1))$.
- (b) $P \xrightarrow{\sigma_{TX_{SP,0}}^P} \bar{P}$
- (c) If party P receives $\sigma_{TX_{SP,0}}^{\bar{P}}$ s.t. $\text{Vrfy}_{pk_{\bar{P}}}(f([TX_{SP,0}], 1); \sigma_{TX_{SP,0}}^{\bar{P}}) = 1$, it continues. Else it stops.
- (d) Party P creates $TX_{SP,0}$.
6. Create $TX_{CM,0}$:
- (a) Party $P \in \{A, B\}$ computes $\tilde{\sigma}_{TX_{CM,0}}^P = \text{pSign}_{sk_P}(f([TX_{CM,0}], 1), Y_{\bar{P},0})$.
- (b) $P \xrightarrow{\tilde{\sigma}_{TX_{CM,0}}^P} \bar{P}$
- (c) If party P receives $\tilde{\sigma}_{TX_{CM,0}}^{\bar{P}}$ s.t. $\text{pVrfy}_{pk_{\bar{P}}}(f([TX_{CM,0}], 1), Y_{P,0}; \tilde{\sigma}_{TX_{CM,0}}^{\bar{P}}) = 1$, it computes $\sigma_{TX_{CM,0}}^{\bar{P}} = \text{Adapt}(\tilde{\sigma}_{TX_{CM,0}}^{\bar{P}}, y_{P,0})$, computes $\sigma_{TX_{CM,0}}^P = \text{Sign}_{sk_P}(f([TX_{CM,0}], 1))$, creates $TX_{CM,0}$ and continues. Else it stops.
7. Create TX_{FU} :
- (a) Party $P \in \{A, B\}$ computes $\sigma_{TX_{FU}}^P = \text{Sign}_{sk_P}(f([TX_{FU}], j))$ where $j := 1$ if $P = A$ or $j := 2$ otherwise.
- (b) $P \xrightarrow{\sigma_{TX_{FU}}^P} \bar{P}$
- (c) If party P receives $\sigma_{TX_{FU}}^{\bar{P}}$ s.t. $\text{Vrfy}_{pk_{\bar{P}}}(f([TX_{FU}], j); \sigma_{TX_{FU}}^{\bar{P}}) = 1$ with $j := 1$ if $P = B$ or $j := 2$ otherwise, it continues. Else it stops.

- (d) Party P creates TX_{FU} .
- 8. Publish TX_{FU} : Party P publishes TX_{FU} on-chain.
- 9. Create $[\text{TX}_{\text{CL}}]$:
 - (a) W creates $[\text{TX}_{\text{CL}}]$ according to 5.4.
 - (b) $W \xrightarrow{\text{txid}_W} A, B$.
 - (c) If party P receives txid_W , it creates $[\text{TX}_{\text{CL}}]$. Else it stops.
- 10. Create $[\text{TX}_{\text{RC}}]$: A, B and W creates $[\text{TX}_{\text{RC}}]$ according to 5.5.
- 11. Create TX_{RC} :
 - (a) Party $P \in \{A, B\}$ computes $\sigma_{\text{TX}_{\text{RC}}}^P = \text{Sign}_{sk_P}(f([\text{TX}_{\text{RC}}], 1))$.
 - (b) $P \xrightarrow{\sigma_{\text{TX}_{\text{RC}}}^P} W$
 - (c) If W receives $\sigma_{\text{TX}_{\text{RC}}}^P$ with $P = \{A, B\}$ s.t. $\text{Vrfy}_{pk_P}(f([\text{TX}_{\text{RC}}], 1); \sigma_{\text{TX}_{\text{RC}}}^P) = 1$, it continues. Else it stops.
 - (d) W computes $\sigma_{\text{TX}_{\text{RC}}}^W = \text{Sign}_{sk_W}(f([\text{TX}_{\text{RC}}], 1))$ and creates TX_{RC} .
- 12. Create TX_{CL} :
 - (a) W computes $\sigma_{\text{TX}_{\text{RV}}}^W = \text{Sign}_{sk_W}(f([\text{TX}_{\text{CL}}], 1))$.
 - (b) W creates TX_{CL} .
- 13. Publish TX_{CL} : W publishes TX_{CL} on-chain.

FPPW channel update protocol is as following:

Preconditions: The channel create phase is complete and TX_{FU} and TX_{CL} are on-chain. The channel update phase has completed i times and hence the channel is at state i .

- 1. Create $[\text{TX}_{\text{CM}, i+1}]$:
 - (a) Party $P \in \{A, B\}$ generates $(Y_{P, i+1}, y_{P, i+1}) \leftarrow \text{GenR}$.
 - (b) $P \xrightarrow{Y_{P, i+1}} \bar{P}, W$
 - (c) If Party P receives $Y_{\bar{P}, i+1}$, it creates $[\text{TX}_{\text{CM}, i+1}]$ according to 5.2. Else it stops.

- (d) If W receives $Y_{P,i+1}$ with $P = \{A, B\}$, it creates $[TX_{CM,i+1}]$ according to 5.2. Else it stops.
2. Create $[TX_{SP,i+1}]$: Party $P \in \{A, B\}$ creates $[TX_{SP,i+1}]$ according to 5.3.
3. Create $[TX_{RV,i+1}]$:
 - (a) W creates $[TX_{RV,i+1}]$ according to 5.6.
 - (b) W computes $\sigma_{TX_{RV,i+1}}^{W,j} = \text{Sign}_{sk_W}(f([TX_{RV,i+1}], j))$ with $j = \{1, 2\}$.
 - (c) $W \xrightarrow{\sigma_{TX_{RV,i+1}}^{W,1}, \sigma_{TX_{RV,i+1}}^{W,2}} A, B$.
 - (d) If party $P \in \{A, B\}$ receives $\sigma_{TX_{RV,i+1}}^{W,1}$ and $\sigma_{TX_{RV,i+1}}^{W,2}$ from W s.t. $\text{Vrfy}_{pk_W}(f([TX_{RV,i+1}], j); \sigma_{TX_{RV,i+1}}^{W,j}) = 1$ with $j = \{1, 2\}$, it continues. Else it stops.
4. Create $TX_{SP,i+1}$:
 - (a) Party $P \in \{A, B\}$ computes $\sigma_{TX_{SP,i+1}}^P = \text{Sign}_{sk_P}(f([TX_{SP,i+1}], 1))$.
 - (b) $P \xrightarrow{\sigma_{TX_{SP,i+1}}^P} \bar{P}$
 - (c) If party P receives $\sigma_{TX_{SP,i+1}}^{\bar{P}}$ s.t. $\text{Vrfy}_{pk_{\bar{P}}}(f([TX_{SP,i+1}], 1); \sigma_{TX_{SP,i+1}}^{\bar{P}}) = 1$, it continues. Else it stops.
 - (d) Party P creates $TX_{SP,i+1}$.
5. Create $TX_{CM,i+1}$:
 - (a) Party $P \in \{A, B\}$ computes $\tilde{\sigma}_{TX_{CM,i+1}}^P = \text{pSign}_{sk_P}(f([TX_{CM,i+1}], 1), Y_{P,i+1})$.
 - (b) $P \xrightarrow{\tilde{\sigma}_{TX_{CM,i+1}}^P} \bar{P}$
 - (c) If party P receives $\tilde{\sigma}_{TX_{CM,i+1}}^{\bar{P}}$ s.t. $\text{pVrfy}_{pk_{\bar{P}}}(f([TX_{CM,i+1}], 1), Y_{P,i+1}; \tilde{\sigma}_{TX_{CM,i+1}}^{\bar{P}}) = 1$, it computes $\sigma_{TX_{CM,i+1}}^{\bar{P}} = \text{Adapt}(\tilde{\sigma}_{CM,i+1}^{\bar{P}}, y_{P,i+1})$, computes $\sigma_{TX_{CM,i+1}}^P = \text{Sign}_{sk_P}(f([TX_{CM,i+1}], 1))$, creates $TX_{CM,i+1}$ and continues. Else it execute the non-collaborative closure phase (from P 's point of view the channel is still at state i).
6. Create $TX_{RV,i}$:
 - (a) Party $P \in \{A, B\}$ creates $[TX_{RV,i}]$ according to 5.6.
 - (b) Party P computes $\sigma_{TX_{RV,i}}^{P,1} = \text{Sign}_{sk_P}(f([TX_{RV,i}], 1))$ and $\sigma_{TX_{RV,i}}^{P,2} = \text{Sign}_{sk_P}(f([TX_{RV,i}], 2))$.

- (c) $P \xrightarrow{\sigma_{TX_{RV,i}}^{P,1}, \sigma_{TX_{RV,i}}^{P,2}} \bar{P}$
- (d) If party P receives $\sigma_{TX_{RV,i}}^{\bar{P},1}$ and $\sigma_{TX_{RV,i}}^{\bar{P},2}$ from $\bar{P}$ s.t. $\text{Vrfy}_{pk_{\bar{P}}}(f([TX_{RV,i}], 1); \sigma_{TX_{RV,i}}^{\bar{P},1}) = 1$ and $\text{Vrfy}_{pk_{\bar{P}}}(f([TX_{RV,i}], 2); \sigma_{TX_{RV,i}}^{\bar{P},2}) = 1$, it continues. Else it executes the non-collaborative closure phase (The channel now is at state $i + 1$).
- (e) $P \xrightarrow{\sigma_{TX_{RV,i}}^{P,1}, \sigma_{TX_{RV,i}}^{P,2}} W$
- (f) If W receives $\sigma_{TX_{RV,i}}^{P,1}$ and $\sigma_{TX_{RV,i}}^{P,2}$ from $P = \{A, B\}$ s.t. $\text{Vrfy}_{pk_P}(f([TX_{RV,i}], 1); \sigma_{TX_{RV,i}}^{P,1}) = 1$ and $\text{Vrfy}_{pk_P}(f([TX_{RV,i}], 2); \sigma_{TX_{RV,i}}^{P,2}) = 1$, it continues. Else it stops.
- (g) A, B and W create $TX_{RV,i}$ according to 5.6.
7. Create $[TX_{PN_1,i}]$ and $[TX_{PN_2,i}]$
- (a) Party $P \in \{A, B\}$ creates $[TX_{PN_1,i}]$ and $[TX_{PN_2,i}]$ according to 5.7 and 5.8.
- (b) Party P computes $\sigma_{TX_{PN_1,i}}^{P,2} = \text{Sign}_{sk_P}(f([TX_{PN_1,i}], 2))$ and $\sigma_{TX_{PN_2,i}}^{P,2} = \text{Sign}_{sk_P}(f([TX_{PN_2,i}], 2))$.
- (c) $P \xrightarrow{\sigma_{TX_{PN_1,i}}^{P,2}, \sigma_{TX_{PN_2,i}}^{P,2}} \bar{P}$
- (d) If P receives signatures $\sigma_{TX_{PN_1,i}}^{\bar{P},2}$ and $\sigma_{TX_{PN_2,i}}^{\bar{P},2}$ s.t. $\text{Vrfy}_{pk_{\bar{P}}}(f([TX_{PN_1,i}], 2); \sigma_{TX_{PN_1,i}}^{\bar{P},2}) = 1$ and $\text{Vrfy}_{pk_{\bar{P}}}(f([TX_{PN_2,i}], 2); \sigma_{TX_{PN_2,i}}^{\bar{P},2}) = 1$, it continues. Else it executes the non-collaborative closure phase (The channel now is at state $i + 1$) or gets online at least once every $t - 1$ blocks.
- (e) W creates $[TX_{PN_1,i}]$ and $[TX_{PN_2,i}]$ according to 5.7 and 5.8
- (f) W computes $\sigma_{TX_{PN_1,i}}^{W,2} = \text{Sign}_{sk_W}(f([TX_{PN_1,i}], 2))$ and $\sigma_{TX_{PN_2,i}}^{W,2} = \text{Sign}_{sk_W}(f([TX_{PN_2,i}], 2))$.
- (g) $W \xrightarrow{\sigma_{TX_{PN_1,i}}^{W,2}, \sigma_{TX_{PN_2,i}}^{W,2}} P$ with $P = \{A, B\}$.
- (h) If $P \in \{A, B\}$ receives signatures $\sigma_{TX_{PN_1,i}}^{W,2}$ and $\sigma_{TX_{PN_2,i}}^{W,2}$ from W s.t. $\text{Vrfy}_{pk_W}(f([TX_{PN_1,i}], 2); \sigma_{TX_{PN_1,i}}^{W,2}) = 1$ and $\text{Vrfy}_{pk_W}(f([TX_{PN_2,i}], 2); \sigma_{TX_{PN_2,i}}^{W,2}) = 1$ it continues. Else it executes the non-collaborative closure phase (The channel now is at state $i + 1$) or gets online at least once every $t - 1$ blocks.

FPPW channel collaborative closure protocol is as following:

Preconditions: The channel create phase is complete and TX_{FU} and TX_{CL} are on-chain. The channel is at state n .

1. Create TX_{SP} :
 - (a) Party $P \in \{A, B\}$ creates $[\text{TX}_{\text{SP}}]$ according to 5.9.
 - (b) P computes $\sigma_{\text{TX}_{\text{SP}}}^P = \text{Sign}_{sk_P}(f([\text{TX}_{\text{SP}}], 1))$.
 - (c) $P \xrightarrow{\sigma_{\text{TX}_{\text{SP}}}^P} \bar{P}$
 - (d) If P receives $\sigma_{\text{TX}_{\text{SP}}}^{\bar{P}}$ from $\bar{P}$ s.t. $\text{Vrfy}_{pk_{\bar{P}}}(f([\text{TX}_{\text{SP}}], 1); \sigma_{\text{TX}_{\text{SP}}}^{\bar{P}}) = 1$, it continues. Else it executes the non-collaborative closure phase (from P 's point of view the channel is still at state n).
2. Publish TX_{SP} : Party $P \in \{A, B\}$ publishes TX_{SP} on-chain.

FPPW channel non-collaborative closure protocol is as following:

Preconditions: The channel create phase is complete and TX_{FU} and TX_{CL} are on-chain. The channel is at state n .

1. Party $P \in \{A, B\}$ publishes $\text{TX}_{\text{CM},n}$ on-chain.
2. Once $\text{TX}_{\text{CM},n}$ is recorded on-chain, P waits for t rounds and then publishes $\text{TX}_{\text{SP},n}$ on-chain.

The protocol for penalising the cheating party is as follows:

preconditions: The channel create phase is complete and TX_{FU} and TX_{CL} are on-chain. The channel is at state n . $\text{TX}_{\text{CM},i}$ with $i < n$ is recorded on-chain by a channel party. The watchtower is always online.

1. W observes that $\text{TX}_{\text{CM},i}$ is on-chain.
2. W publishes $\text{TX}_{\text{RV},i}$ on-chain before t rounds being elapsed since broadcast of $\text{TX}_{\text{CM},i}$.

The protocol for penalising the unresponsive watchtower is as follows:

preconditions: The channel create phase is complete and TX_{FU} and TX_{CL} are on-chain. The channel update phase has successfully completed n times. $\text{TX}_{\text{CM},i}$ with

$i < n$ is recorded on-chain. Parties check the blockchain at least once every $T - 1$ round.

1. Party P observes that $\mathcal{B}_j$ is the latest block on the blockchain and $\text{TX}_{\text{CM},i}$ has been published through the block $\mathcal{B}_k$ with $k \leq j$ but its first output has not been spent by $\text{TX}_{\text{RV},i}$.
2. if $j + 1 - k < t$:
 - (a) P publishes $\text{TX}_{\text{RV},i}$ on the blockchain.

Otherwise

- (a) P extract $\text{TX}_{\text{CM},i}.\text{Witness}[1]$ as $(1, (\sigma_{\text{TX}_{\text{CM},i}}^A, \sigma_{\text{TX}_{\text{CM},i}}^B))$.
- (b) P computes $y_{\bar{P},i} = \text{Ext}(\sigma_{\text{TX}_{\text{CM},i}}^P, \tilde{\sigma}_{\text{CM},i}^P, Y_{\bar{P},i})$.
- (c) If TX_{RC} is unpublished:
 - P computes $\sigma_{\text{TX}_{\text{PN}_1,i}}^{P,1} = \text{Sign}_{sk_P}(f([\text{TX}_{\text{PN}_1,i}], 1))$ and $\sigma_{\text{TX}_{\text{PN}_1,i}}^{P,1} = \text{Sign}_{y_{\bar{P},i}}(f([\text{TX}_{\text{PN}_1,i}], 1))$.
 - P creates $\text{TX}_{\text{PN}_1,i}$ using $[\text{TX}_{\text{PN}_1,i}], \text{TX}_{\text{PN}_1,i}.\text{Witness}[1] = (1, (\sigma_{\text{TX}_{\text{PN}_1,i}}^{P,1}, \sigma_{\text{TX}_{\text{PN}_1,i}}^{P,1}))$ with $j = 1$ if $P = B$ or $j = 3$ otherwise and $\text{TX}_{\text{PN}_1,i}.\text{Witness}[2] = (1, (\sigma_{\text{TX}_{\text{PN}_1,i}}^{A,2}, \sigma_{\text{TX}_{\text{PN}_1,i}}^{B,2}, \sigma_{\text{TX}_{\text{PN}_1,i}}^{W,2}))$.
 - P publishes $\text{TX}_{\text{PN}_1,i}$ on-chain.

Else:

- P computes $\sigma_{\text{TX}_{\text{PN}_2,i}}^{P,1} = \text{Sign}_{sk_P}(f([\text{TX}_{\text{PN}_2,i}], 1))$ and $\sigma_{\text{TX}_{\text{PN}_2,i}}^{P,1} = \text{Sign}_{y_{\bar{P},i}}(f([\text{TX}_{\text{PN}_2,i}], 1))$.
- P creates $\text{TX}_{\text{PN}_2,i}$ using $[\text{TX}_{\text{PN}_2,i}], \text{TX}_{\text{PN}_2,i}.\text{Witness}[1] = (j, (\sigma_{\text{TX}_{\text{PN}_2,i}}^{P,1}, \sigma_{\text{TX}_{\text{PN}_2,i}}^{P,1}))$ with $j = 1$ if $P = B$ or $j = 3$ otherwise and $\text{TX}_{\text{PN}_2,i}.\text{Witness}[2] = (1, (\sigma_{\text{TX}_{\text{PN}_2,i}}^{A,2}, \sigma_{\text{TX}_{\text{PN}_2,i}}^{B,2}, \sigma_{\text{TX}_{\text{PN}_2,i}}^{W,2}))$.
- P publishes $\text{TX}_{\text{PN}_2,i}$ on-chain.

5.8 Temporarily unavailable watchtower

FPPW can adapt to situations where channel parties want to update the channel state but the watchtower is temporarily unavailable. A way to deal with such occasions is that channel parties wait for the watchtower to get responsive. However, it disturbs the main functionality of the payment channel. Another solution is updating the channel by skipping those steps that require the watchtower to sign the new revocation transaction

(see step 3 of the channel update protocol in Section 5.7) and new penalty transactions (see step 7 of the channel update protocol in Section 5.7). This solution also has two problems. Firstly, without the watchtower cooperation, the new commit transaction gets irrevocable. Thus, channel parties would not be able to update the channel multiple times. Secondly, if the watchtower gets uncooperative, channel parties will be forced to update the channel on-chain even if they both agree to remain always online and continue using the channel.

To resolve this issue, channel parties can take steps in the channel update phase using the following commit transactions:

$$\begin{aligned}
\text{TX}_{\text{CM},i+1}.\text{Input} &:= \text{TX}_{\text{FU}}.\text{txid}\|1, \\
\text{TX}_{\text{CM},i+1}.\text{Output} &:= ((a + b, \varphi_1 \vee \varphi_2 \vee \varphi_3), (\epsilon, \varphi'_1 \vee \varphi'_2 \vee \varphi'_3)) \\
\text{TX}_{\text{CM},i+1}.\text{Witness} &:= \{(1, \{\sigma_{\text{TX}_{\text{CM},i}}^A, \sigma_{\text{TX}_{\text{CM},i}}^B\})\}
\end{aligned} \tag{5.10}$$

with $\varphi_1 := pk_A \wedge pk_B \wedge 2t^+$, $\varphi_2 := pk_A \wedge pk_B \wedge pk_W$, $\varphi_3 := pk_A \wedge pk_B \wedge t^+$, $\varphi'_1 := pk_B \wedge Y_{A,i+1} \wedge t^+$, $\varphi'_2 := pk_A \wedge pk_B \wedge pk_W$ and $\varphi'_3 := pk_A \wedge Y_{B,i+1} \wedge t^+$ where $Y_{A,i+1}$ and $Y_{B,i+1}$ are statements of a hard relation $\mathcal{R}$ generated by A and B for the $i + 1^{\text{th}}$ state using the generating algorithm GenR.

Also, a new type of revocation transaction is introduced as follows:

$$\begin{aligned}
\text{TX}_{\text{RV}',i+1}.\text{Input} &:= (\text{TX}_{\text{CM},i+1}.\text{txid}\|1), \\
\text{TX}_{\text{RV}',i+1}.\text{Output} &:= (a + b, Y_{A,i+1} \wedge Y_{B,i+1}), \\
\text{TX}_{\text{RV}',i+1}.\text{Witness} &:= \{(3, (\sigma_{\text{TX}_{\text{RV}',i+1}}^{A,1}, \sigma_{\text{TX}_{\text{RV}',i+1}}^{B,1}))\}
\end{aligned} \tag{5.11}$$

Then, the update protocol is as follows:

Preconditions: The channel create phase is complete and TX_{FU} and TX_{CL} are on-chain. The channel update phase has completed i times and hence the channel is at state i . The watchtower is unavailable.

1. Create $[\text{TX}_{\text{CM},i+1}]$:

(a) Party P generates $(Y_{P,i+1}, y_{P,i+1}) \leftarrow \text{GenR}$.

(b) $P \xrightarrow{Y_{P,i+1}} \tilde{P}$

- (c) If Party P receives $Y_{\bar{P},i+1}$, it creates $[TX_{CM,i+1}]$ according to 5.10. Else it stops.
2. Create $[TX_{SP,i+1}]$: Party P creates $[TX_{SP,i+1}]$ according to 5.3.
 3. Create $[TX_{RV',i+1}]$: Party P creates $[TX_{RV',i+1}]$ according to 5.11.
 4. Create $TX_{SP,i+1}$:
 - (a) Party P computes $\sigma_{TX_{SP,i+1}}^P = \text{Sign}_{sk_P}([f(TX_{SP,i+1}), 1])$.
 - (b) $P \xrightarrow{\sigma_{TX_{SP,i+1}}^P} \bar{P}$
 - (c) If party P receives $\sigma_{TX_{SP,i+1}}^{\bar{P}}$ s.t. $\text{Vrfy}_{pk_{\bar{P}}}(f([TX_{SP,i+1}], 1); \sigma_{TX_{SP,i+1}}^{\bar{P}}) = 1$, it continues. Else it stops.
 - (d) Party P creates $TX_{SP,i+1}$.
 5. Create $TX_{CM,i+1}$:
 - (a) Party P computes $\tilde{\sigma}_{CM,i+1}^P = \text{pSign}_{sk_P}(f([TX_{CM,i+1}], 1), Y_{\bar{P},i+1})$.
 - (b) $P \xrightarrow{\tilde{\sigma}_{CM,i+1}^P} \bar{P}$
 - (c) If party P receives $\tilde{\sigma}_{CM,i+1}^{\bar{P}}$ s.t. $\text{pVrfy}_{pk_{\bar{P}}}(f([TX_{CM,i+1}], 1), Y_{\bar{P},i+1}; \tilde{\sigma}_{CM,i+1}^{\bar{P}}) = 1$, it computes $\sigma_{TX_{CM,i+1}}^{\bar{P}} = \text{Adapt}(\tilde{\sigma}_{CM,i+1}^{\bar{P}}, y_{P,i+1})$, computes $\sigma_{TX_{CM,i+1}}^P = \text{Sign}_{sk_P}(f([TX_{CM,i+1}], 1))$, creates $TX_{CM,i+1}$ and continues. Else it execute the non-collaborative closure phase (from P 's point of view the channel is still at state i).
 6. Create $TX_{RV,i}$ or $TX_{RV',i}$:
 - (a) If $TX_{CM,i}$ is according to 5.2,
 - i. Party P computes $\sigma_{TX_{RV,i}}^{P,1} = \text{Sign}_{sk_P}(f([TX_{RV,i}], 1))$ and $\sigma_{TX_{RV,i}}^{P,2} = \text{Sign}_{sk_P}(f([TX_{RV,i}], 2))$.
 - ii. $P \xrightarrow{\sigma_{TX_{RV,i}}^{P,1}, \sigma_{TX_{RV,i}}^{P,2}} \bar{P}$
 - iii. If party P receives $\sigma_{TX_{RV,i}}^{\bar{P},1}$ and $\sigma_{TX_{RV,i}}^{\bar{P},2}$ from $\bar{P}$ s.t. $\text{Vrfy}_{pk_{\bar{P}}}(f([TX_{RV,i}], 1); \sigma_{TX_{RV,i}}^{\bar{P},1}) = 1$ and $\text{Vrfy}_{pk_{\bar{P}}}(f([TX_{RV,i}], 2); \sigma_{TX_{RV,i}}^{\bar{P},2}) = 1$, it continues. Else it executes the non-collaborative closure phase (The channel now is at state $i + 1$).
 - iv. Party P creates $TX_{RV,i}$.
 - (b) Otherwise:

- i. Party P computes $\sigma_{\text{TX}_{\text{RV},i}}^{P,1} = \text{Sign}_{sk_P}(f([\text{TX}_{\text{RV},i}], 1))$.
- ii. $P \xrightarrow{\sigma_{\text{TX}_{\text{RV},i}}^{P,1}} \bar{P}$.
- iii. If party P receives $\sigma_{\text{TX}_{\text{RV},i}}^{\bar{P},1}$ from $\bar{P}$ s.t. $\text{Vrfy}_{pk_{\bar{P}}}(f([\text{TX}_{\text{RV},i}], 1); \sigma_{\text{TX}_{\text{RV},i}}^{\bar{P},1}) = 1$, it continues. Else it executes the non-collaborative closure phase (The channel now is at state $i + 1$).
- iv. Party P creates $\text{TX}_{\text{RV},i}$.

Remark 5.6. The first (second) condition in step 6 corresponds with the case when watchtower was available (unavailable) during the channel update from state $i - 1$ to i .

Now assume that one of the revoked $\text{TX}_{\text{CM},i}$, for which $\text{TX}_{\text{RV},i}$ has been created, is published by a channel party. The cheating party must wait $2t$ rounds to be able to publish the split transaction $\text{TX}_{\text{SP},i}$. However, its online counterparty can wait for t rounds and then publish the $\text{TX}_{\text{RV},i}$ and take all the funds of the channel. The steps of this procedure are as follows:

Preconditions: The channel create phase is complete and TX_{FU} and TX_{CL} are on-chain. The channel update phase has successfully completed n times. $\text{TX}_{\text{CM},i}$ with $i < n$ is recorded on-chain. Parties have created $\text{TX}_{\text{RV},i}$. Channel parties are always online.

1. P observes that $\text{TX}_{\text{CM},i}$ is on-chain. P waits for t blocks.
2. P publishes $\text{TX}_{\text{RV},i}$ on-chain.

When the watchtower becomes available, $\text{TX}_{\text{RV},i}$, $\text{TX}_{\text{PN}_1,i}$ and $\text{TX}_{\text{PN}_2,i}$ (respectively according to 5.6, 5.7 and 5.8) for the new agreed states can be created by A , B and W . The steps of this procedure are as follows:

Preconditions: The channel create phase is complete and TX_{FU} and TX_{CL} are on-chain. The channel update phase has successfully completed m times with $m \geq 1$. The watchtower W was unavailable during the latest k channel updates with $0 < k \leq m$. The watchtower W is now available.

1. A , B and W repeats the following steps for $i = m - k + 1$ to $i = m - 1$:
 - (a) Create $[\text{TX}_{\text{RV},i}]$:
 - i. W creates $[\text{TX}_{\text{RV},i}]$ according to 5.6.

- ii. W computes $\sigma_{\text{TX}_{\text{RV},i}}^{W,j} = \text{Sign}_{sk_W}([f(\text{TX}_{\text{RV},i}], j))$ with $j = \{1, 2\}$.
 - iii. $W \xrightarrow{\sigma_{\text{TX}_{\text{RV},i}}^{W,1}, \sigma_{\text{TX}_{\text{RV},i}}^{W,2}} A, B$.
 - iv. If party $P \in \{A, B\}$ receives $\sigma_{\text{TX}_{\text{RV},i}}^{W,1}$ and $\sigma_{\text{TX}_{\text{RV},i}}^{W,2}$ from W s.t. $\text{Vrfy}_{pk_W}(f([\text{TX}_{\text{RV},i}], j); \sigma_{\text{TX}_{\text{RV},i}}^{W,j}) = 1$ with $j = \{1, 2\}$, it continues. Else it stops.
- (b) Create $\text{TX}_{\text{RV},i}$:
- i. Party $P \in \{A, B\}$ creates $[\text{TX}_{\text{RV},i}]$ according to 5.6.
 - ii. Party P computes $\sigma_{\text{TX}_{\text{RV},i}}^{P,1} = \text{Sign}_{sk_P}(f([\text{TX}_{\text{RV},i}], 1))$ and $\sigma_{\text{TX}_{\text{RV},i}}^{P,2} = \text{Sign}_{sk_P}(f([\text{TX}_{\text{RV},i}], 2))$.
 - iii. $P \xrightarrow{\sigma_{\text{TX}_{\text{RV},i}}^{P,1}, \sigma_{\text{TX}_{\text{RV},i}}^{P,2}} \bar{P}$
 - iv. If party P receives $\sigma_{\text{TX}_{\text{RV},i}}^{\bar{P},1}$ and $\sigma_{\text{TX}_{\text{RV},i}}^{\bar{P},2}$ from $\bar{P}$ s.t. $\text{Vrfy}_{pk_{\bar{P}}}(f([\text{TX}_{\text{RV},i}], 1); \sigma_{\text{TX}_{\text{RV},i}}^{\bar{P},1}) = 1$ and $\text{Vrfy}_{pk_{\bar{P}}}(f([\text{TX}_{\text{RV},i}], 2); \sigma_{\text{TX}_{\text{RV},i}}^{\bar{P},2}) = 1$, it continues. Else it executes the non-collaborative closure phase (The channel now is at state m).
 - v. $P \xrightarrow{\sigma_{\text{TX}_{\text{RV},i}}^{P,1}, \sigma_{\text{TX}_{\text{RV},i}}^{P,2}} W$.
 - vi. If W receives $\sigma_{\text{TX}_{\text{RV},i}}^{P,1}$ and $\sigma_{\text{TX}_{\text{RV},i}}^{P,2}$ from $P = \{A, B\}$ s.t. $\text{Vrfy}_{pk_P}(f([\text{TX}_{\text{RV},i}], 1); \sigma_{\text{TX}_{\text{RV},i}}^{P,1}) = 1$ and $\text{Vrfy}_{pk_P}(f([\text{TX}_{\text{RV},i}], 2); \sigma_{\text{TX}_{\text{RV},i}}^{P,2}) = 1$, it continues. Else it stops.
 - vii. A, B and W create $\text{TX}_{\text{RV},i}$ using $[\text{TX}_{\text{RV},i}]$ and $\text{TX}_{\text{RV},i}.\text{Witness}[j] = (2, (\sigma_{\text{TX}_{\text{RV},i}}^{A,j}, \sigma_{\text{TX}_{\text{RV},i}}^{B,j}, \sigma_{\text{TX}_{\text{RV},i}}^{W,j}))$ with $j = \{1, 2\}$.
- (c) Create $[\text{TX}_{\text{PN}_1,i}]$ and $[\text{TX}_{\text{PN}_2,i}]$
- i. Party $P \in \{A, B\}$ creates $[\text{TX}_{\text{PN}_1,i}]$ and $[\text{TX}_{\text{PN}_2,i}]$ according to 5.7 and 5.8.
 - ii. Party P computes $\sigma_{\text{TX}_{\text{PN}_1,i}}^{P,2} = \text{Sign}_{sk_P}(f([\text{TX}_{\text{PN}_1,i}], 2))$ and $\sigma_{\text{TX}_{\text{PN}_2,i}}^{P,2} = \text{Sign}_{sk_P}(f([\text{TX}_{\text{PN}_2,i}], 2))$, respectively.
 - iii. $P \xrightarrow{\sigma_{\text{TX}_{\text{PN}_1,i}}^{P,2}, \sigma_{\text{TX}_{\text{PN}_2,i}}^{P,2}} \bar{P}$
 - iv. If P receives signatures $\sigma_{\text{TX}_{\text{PN}_1,i}}^{\bar{P},2}$ and $\sigma_{\text{TX}_{\text{PN}_2,i}}^{\bar{P},2}$ s.t. $\text{Vrfy}_{pk_{\bar{P}}}(f([\text{TX}_{\text{PN}_1,i}], 2); \sigma_{\text{TX}_{\text{PN}_1,i}}^{\bar{P},2}) = 1$ and $\text{Vrfy}_{pk_{\bar{P}}}(f([\text{TX}_{\text{PN}_2,i}], 2); \sigma_{\text{TX}_{\text{PN}_2,i}}^{\bar{P},2}) = 1$, it continues. Else it executes the non-collaborative closure phase (The channel now is at state m) or gets online at least once every $t - 1$ blocks.
 - v. W creates $[\text{TX}_{\text{PN}_1,i}]$ and $[\text{TX}_{\text{PN}_2,i}]$ according to 5.7 and 5.8

- vi. W computes $\sigma_{\text{TX}_{\text{PN}_1,i}}^{W,2} = \text{Sign}_{sk_W}(f([\text{TX}_{\text{PN}_1,i}], 2))$ and $\sigma_{\text{TX}_{\text{PN}_2,i}}^{W,2} = \text{Sign}_{sk_W}(f([\text{TX}_{\text{PN}_2,i}], 2))$.
 - vii. $W \xrightarrow{\sigma_{\text{TX}_{\text{PN}_1,i}}^{W,2}, \sigma_{\text{TX}_{\text{PN}_2,i}}^{W,2}} P$ with $P = \{A, B\}$.
 - viii. If $P \in \{A, B\}$ receives signatures $\sigma_{\text{TX}_{\text{PN}_1,i}}^{W,2}$ and $\sigma_{\text{TX}_{\text{PN}_2,i}}^{W,2}$ from W s.t. $\text{Vrfy}_{pk_W}(f([\text{TX}_{\text{PN}_1,i}], 2); \sigma_{\text{TX}_{\text{PN}_1,i}}^{W,2}) = 1$ and $\text{Vrfy}_{pk_W}(f([\text{TX}_{\text{PN}_2,i}], 2); \sigma_{\text{TX}_{\text{PN}_2,i}}^{W,2}) = 1$ it continues. Else it executes the non-collaborative closure phase (The channel now is at state m) or gets online at least once every $t - 1$ blocks.
2. Create $[\text{TX}_{\text{RV},m}]$:
- (a) W creates $[\text{TX}_{\text{RV},m}]$ according to 5.6.
 - (b) W computes $\sigma_{\text{TX}_{\text{RV},m}}^{W,j} = \text{Sign}_{sk_W}(f([\text{TX}_{\text{RV},m}], j))$ with $j = \{1, 2\}$.
 - (c) $W \xrightarrow{\sigma_{\text{TX}_{\text{RV},m}}^{W,1}, \sigma_{\text{TX}_{\text{RV},m}}^{W,2}} A, B$.
 - (d) If party $P \in \{A, B\}$ receives $\sigma_{\text{TX}_{\text{RV},m}}^{W,1}$ and $\sigma_{\text{TX}_{\text{RV},m}}^{W,2}$ from W s.t. $\text{Vrfy}_{pk_W}(f([\text{TX}_{\text{RV},m}], j); \sigma_{\text{TX}_{\text{RV},m}}^{W,j}) = 1$ with $j = \{1, 2\}$, it continues. Else it stops.

The commit and revocation transactions introduced in 5.10 and 5.11 have an interesting property that allows the honest channel party to penalise both the cheating party and the unresponsive watchtower. Assume that there is a revoked commit transaction $\text{TX}_{\text{CM},i}$ according to 5.10 for which the watchtower has also signed penalty transactions. If this transaction is published by a cheating party, let's say A , the watchtower can immediately publish $\text{TX}_{\text{RV},i}$ to invalidate the penalty transactions. However, if the watchtower is unresponsive, after t rounds, B can firstly publish $\text{TX}_{\text{RV}',i}$ to penalise the cheating party B and then publish $\text{TX}_{\text{PN}_1,i}$ or $\text{TX}_{\text{PN}_2,i}$ to penalise the unresponsive watchtower. Split transaction $\text{TX}_{\text{SP},i}$ cannot be published within $2t - 1$ rounds since broadcast of $\text{TX}_{\text{CM},i}$.

5.9 FPPW Transactions Scripts

Bitcoin scripting is a fundamental component of the Bitcoin protocol that allows users to create and enforce conditions for spending bitcoins. It involves using a simple and stack-based scripting language to define the conditions that must be met to unlock the funds in a Bitcoin transaction output (UTXO). When bitcoins are sent to an address, the recipient specifies conditions in the form of a ScriptPubKey or simply the script. This script defines the spending conditions required to unlock and spend the funds. To spend

the bitcoins locked in a UTXO, the sender must provide a witness in the spending transaction. The witness is combined with the ScriptPubKey of the UTXO, and the resulting script is executed. The script execution is done by pushing data onto a stack and applying various script operations. If the final result of the script evaluation is true, the UTXO can be spent. The script operations include basic arithmetic, cryptographic operations, conditional statements (e.g. OP_IF and OP_ELSE), and signature checks (e.g. OP_CHECKMULTISIG) [59].

Funding transaction has one output with the following script where pubkeyA, pubkeyB are public keys of A and B, respectively:

```
2 <pubkeyA> <pubkeyB> 2 OP_CHECKMULTISIG
```

Commit transaction has one input that takes the output of the funding transaction with witness script 0 <pubkeyA_sig> <pubkeyB_sig>. It also has two outputs where the script of its first output (main output) is as follows:

```
OP_IF
```

```
  # Revocation
```

```
  3 <Rev_pubkeyA> <Rev_pubkeyB> <Rev_pubkeyW> 3 OP_CHECKMULTISIG
```

```
OP_ELSE
```

```
  # Split
```

```
  <delay t> OP_CHECKSEQUENCEVERIFY OP_DROP
```

```
  2 <Spl_pubkeyA> <Spl_pubkeyB> 2
```

```
OP_ENDIF
```

where <Rev_pubkeyA> and <Spl_pubkeyA> are public keys of A, <Rev_pubkeyB> and <Spl_pubkeyB> are public keys of B and <Rev_pubkeyW> is public key of W.

The script for the second output (auxiliary output) of the commit transaction is as follows:

```
OP_IF
```

```
  # Revocation
```

```
  3 <Rev_pubkeyA> <Rev_pubkeyB> <Rev_pubkeyW> 3 OP_CHECKMULTISIG
```

```
OP_ELSE
```

```
  <delay t> OP_CHECKSEQUENCEVERIFY OP_DROP
```

```
OP_IF
```

```
  # Penalty1 or Penalty2 by party B
```

```

    2 <Pen_pubkeyB> <YA> 2 OP_CHECKMULTISIG
OP_ELSE
    # Penalty1 or Penalty2 by party A
    2 <Pen_pubkeyA> <YB> 2 OP_CHECKMULTISIG
OP_ENDIF
OP_ENDIF

```

where $\langle \text{Rev_pubkeyA} \rangle$ and $\langle \text{Pen_pubkeyA} \rangle$ are public keys of A , $\langle \text{Rev_pubkeyB} \rangle$ and $\langle \text{Pen_pubkeyB} \rangle$ are public keys of B and $\langle \text{Rev_pubkeyW} \rangle$ is public key of W . Also, YA and YB are statement of A and B , respectively. The witness script for input of split transaction is $0 \langle \text{Spl_pubkeyA_Sig} \rangle \langle \text{Spl_pubkeyB_Sig} \rangle 0$

The revocation transaction has two inputs where its first and second inputs take the first and second outputs of the corresponding commit transaction, respectively. The witness script for both inputs is $0 \langle \text{Rev_pubkeyA_sig} \rangle \langle \text{Rev_pubkeyB_sig} \rangle \langle \text{Rev_pubkeyW_sig} \rangle$. It also has one output with the following script:

```
2 <YA> <YB> 2 OP_CHECKMULTISIG
```

Collateral transaction has one output with script $3 \langle \text{pubkeyA} \rangle \langle \text{pubkeyB} \rangle \langle \text{pubkeyW} \rangle 3 \text{OP_CHECKMULTISIG}$, where pubkeyA , pubkeyB and pubkeyW are the public keys of A , B and W , respectively.

The reclaim transaction has one input taking the collateral transaction output with witness script:

```
0 <pubkeyA_sig> <pubkeyB_sig> <pubkeyW_sig>
```

It also has a single output, with the following script:

```

OP_IF
    # Penalty2
    3 <Pen_pubkeyA> <Pen_pubkeyB> <Pen_pubkeyW> 3 OP_CHECKMULTISIG
OP_ELSE
    # normal
    <delay T> OP_CHECKSEQUENCEVERIFY OP_DROP
    <pubkeyW> OP_CHECKSIG
OP_ENDIF

```

where Pen_pubkeyA , Pen_pubkeyB and Pen_pubkeyW are public keys of A , B and W , respectively. pubkeyW is also the public key of W .

Penalty transaction 1 has two inputs. The first one takes the second output of the corresponding commit transaction with the following witness script if B is broadcasting the penalty transaction:

$0 \langle \text{Pen_pubkeyB_Sig} \rangle \langle \text{YA_Sig} \rangle 1 \ 0$

or with the following witness script, if A is broadcasting it:

$0 \langle \text{Pen_pubkeyA_Sig} \rangle \langle \text{YB_Sig} \rangle 0 \ 0$

The second input takes the output of the collateral transaction with the witness script $0 \langle \text{pubkeyA_Sig} \rangle \langle \text{pubkeyB_Sig} \rangle \langle \text{pubkeyW_Sig} \rangle$. Also, the script for its output is similar to that of the revocation transaction.

Penalty transaction 2 is similar to penalty transaction 1. The only difference is that its second input spends the output of the reclaim transaction. The witness for the second input is $0 \langle \text{Pen_pubkeyA_Sig} \rangle \langle \text{Pen_pubkeyB_Sig} \rangle \langle \text{Pen_pubkeyW_Sig} \rangle$

5.10 FPPW with One Hiring Party

If only one of the channel parties is willing to hire the watchtower, some changes must be applied to FPPW. The transaction flows for FPPW in such scenarios are depicted in Fig. 5.6. In this scenario, without loss of generality, we assume that party A is the hiring party and hence only A funds the extra ϵ . Also, the auxiliary output of the commit transaction has only two sub-conditions where the first one is used by A for penalty purposes and the second one can be used by both parties for revocation purposes. Output condition for TX_{CL} as well as the first sub-condition for output of TX_{RC} is $pk_A \wedge pk_W$ and public key of party B is not involved in these transactions anymore. Furthermore, output of $\text{TX}_{\text{PN}_1,i}$ and $\text{TX}_{\text{PN}_2,i}$ can be only claimed by party A . Therefore, If party B publishes a revoked commit transaction, W can publish its corresponding revocation transaction and then only A can claim its output. Otherwise, A can penalise the watchtower by publishing either $\text{TX}_{\text{PN}_1,i}$ or $\text{TX}_{\text{PN}_2,i}$. Similarly, if A publishes a revoked commit transaction, B can broadcast its corresponding revocation transaction and claim all the channel funds. The watchtower is only paid by A .

5.11 Conclusion

In this chapter, we presented a new payment channel with a watchtower for Bitcoin, called FPPW, which achieves fairness with respect to both the channel party and the watchtower as well as weak privacy against the watchtower. FPPW is the same as Cerberus with respect to the watchtower privacy against third parties meaning that it can

Chapter 6

Garrison: a storage efficient Bitcoin watchtower

We construct this chapter based on our published paper, “Garrison: a novel watchtower scheme for bitcoin” [60].

6.1 Introduction

In the previous chapter, we mostly focused on the privacy and fairness of a watchtower scheme and presented a fair and privacy-preserving watchtower scheme, called FPPW. However, in this chapter, our attention shifts towards evaluating the storage costs incurred by both the watchtower and channel parties. These costs play a crucial role in determining the operational expenses associated with running services such as hubs in PCHs or watchtower services. The storage size of the watchtower in Monitor [15], DCWC* [32] as well as Generalized [14], Cerberus [23] and FPPW [54] channels increases linearly with each channel update and hence the watchtower’s storage costs would be $\mathcal{O}(n)$ with n being the number of channel updates.

Outpost [22] is a novel payment channel with a watchtower scheme that reduces the watchtower’s storage requirements per channel from $\mathcal{O}(n)$ to $\mathcal{O}(\log n)$. This consequently reduces the operational costs of maintaining watchtowers. Although elegantly designed, Outpost suffers from the following shortcomings,

- The storage cost of each channel party is still $\mathcal{O}(n)$.
- Each party has his own version of the channel state where this state duplication causes the number of transactions to exponentially increase with the number of applications on top of each other [14]. In other words, to add an application (e.g.

Table 6.1: Comparison of different dispute period-based payment channels with n channel updates, m HTLC outputs on average per state and k channel splits on top of each other.

Scheme	Party's St. Cost	Watch. St. Cost	on-chain TX. ^a	off-chain TX. ^b
Lightning [21]	$\mathcal{O}(\log n)$	$\mathcal{O}(N)$	$\mathcal{O}(m)$	$\mathcal{O}(2^k)$
Generalized [32]	$\mathcal{O}(\log n)$	$\mathcal{O}(n)$	$\mathcal{O}(1)$	$\mathcal{O}(1)$
Outpost [22]	$\mathcal{O}(n)$	$\mathcal{O}(\log n)$	$\mathcal{O}(m)$	$\mathcal{O}(2^k)$
FPPW [54]	$\mathcal{O}(n)$	$\mathcal{O}(n)$	$\mathcal{O}(1)$	$\mathcal{O}(1)$
Cerberus [23]	$\mathcal{O}(n)$	$\mathcal{O}(n)$	$\mathcal{O}(m)$	$\mathcal{O}(2^k)$
Garrison	$\mathcal{O}(\log n)$	$\mathcal{O}(\log n)$	$\mathcal{O}(1)$	$\mathcal{O}(1)$

^aNumber of on-chain transactions upon dispute.

^bNumber of off-chain transactions per state.

Virtual channel [27]) on top of the channel, parties must split their channel into sub-channels. If parties recursively split their channel k times, then to update their last layer sub-channel, they must create $\mathcal{O}(2^k)$ different versions of the channel state.

- Outpost works based on “punish-per-output” pattern, meaning that if there are m outputs in the published old state, the cheated party must claim each output separately [14]. Then, the required on-chain transactions upon dispute would be $\mathcal{O}(m)$ with m being the number of outputs in the published old state.

Therefore, the main motivation of this chapter is designing a Bitcoin payment channel with a watchtower scheme which is storage-efficient for channel parties and the watchtower and also avoids state duplication and punish-per-output pattern.

The contribution of this chapter is to present a new payment channel with a watchtower for Bitcoin, called Garrison, for which the storage cost of channel parties and the watchtower would be logarithmic in the maximum number of channel updates. Furthermore, both channel parties store the same version of transactions. Additionally, regardless of the number of outputs in each channel state, there exists a single revocation transaction per state. We also prove the security of the Garrison channel based on the security of its underlying cryptographic primitives. Table 6.1 presents a comparison between Garrison and other Bitcoin payment channels that work based on dispute period.

6.2 Notations

In this section, we add a new notation to the existing notations introduced in Section 3.3. We already know that a Bitcoin output θ is a tuple of two attributes, $\theta = (\text{cash}, \varphi)$, where

$\theta.\text{cash}$ denotes the number of coins held in this output. However, the OP_RETURN output is a special output that does not hold any coins and is used to add some arbitrary data to the blockchain. Such an output is denoted by $\theta = (0, \text{data})$, where *data* is its arbitrary data. OP_RETURN outputs are illustrated by blocked lines (instead of directional arrows) in charts. As an example, the second output in the transaction TX in Fig. 6.1 is an OP_RETURN output.




Figure 6.1: A sample transaction flow.

6.3 Garrison Overview

6.3.1 System Model

Channel parties exchange data using an authenticated and secure communication channel. Channel participants might deviate from the protocol if it increases their profit. Furthermore, the underlying blockchain contains a distributed ledger that achieves security [56]. If a valid transaction is propagated in the blockchain network, it is included in the blockchain ledger within Δ rounds (i.e. the confirmation delay is Δ).

6.3.2 Garrison Overview

This section overviews the Garrison channel between *A* (Alice) and *B* (Bob). To introduce Garrison, we use the payment channel NVG, introduced in Section 5.3.2.1, as the starting point and modify it step by step to mitigate its limitations.

6.3.2.1 Reducing the Storage Requirements of the Watchtower

One of the limitations of NVG channel is that all revocation transactions must be stored by channel parties or their watchtowers to be published upon fraud. To reduce the storage requirements of the watchtower, similar to Outpost [22], our main idea is to storing the revocation transaction $\text{TX}_{\text{RV},i}$ inside the commit transaction $\text{TX}_{\text{CM},i}$. Then, once $\text{TX}_{\text{CM},i}$ is published, the watchtower extracts $\text{TX}_{\text{RV},i}$ and records it on the blockchain. However, we have $\text{TX}_{\text{RV},i}.\text{Input} = \text{TX}_{\text{CM},i}.\text{txid}\|1$. Thus, if $\text{TX}_{\text{RV},i}$ is created, signed and finally stored

inside $\text{TX}_{\text{CM},i}$, then $[\text{TX}_{\text{CM},i}]$ and hence $\text{TX}_{\text{CM},i}.\text{txid}$ and $\text{TX}_{\text{RV},i}$ change. Thus, there is a self-loop situation [22]. To solve this issue, we add an auxiliary output with the value of ϵ to commit transactions where ϵ is the minimum value supported by the Bitcoin blockchain. We also add an auxiliary transaction between each commit transaction and its corresponding split transaction. This new transaction $\text{TX}_{\text{AU},i}$ spends the auxiliary output of the commit transaction. The signatures of party A and party B on $[\text{TX}_{\text{RV},i}]$ are stored in an `OP_RETURN` output of the auxiliary transaction $\text{TX}_{\text{AU},i}$. The split transaction $\text{TX}_{\text{SP},i}$ spends the main output of $\text{TX}_{\text{CM},i}$ as well as the main output of the auxiliary transaction $\text{TX}_{\text{AU},i}$. Based on this design, parties can be sure that once the revoked commit transaction $\text{TX}_{\text{CM},i}$ is published on the blockchain, its split transaction $\text{TX}_{\text{SP},i}$ cannot be published unless $\text{TX}_{\text{AU},i}$ is also on the blockchain. Furthermore, due to the time-lock in the main output of $\text{TX}_{\text{AU},i}$, once this transaction is published on-chain, $\text{TX}_{\text{SP},i}$ cannot be published within $t - 1$ rounds. However, the honest party or the watchtower can extract the signatures on $[\text{TX}_{\text{RV},i}]$ from $\text{TX}_{\text{AU},i}$ and publish $\text{TX}_{\text{RV},i}$ immediately. Fig. 6.2 depicts the transaction flows.




Figure 6.2: Reducing the Storage Requirements of the Watchtower

However, this scheme has the following issues:

- To create and publish the revocation transaction, the watchtower must also know the value of $Y_{A,i}$ and $Y_{B,i}$.
- Typically, revocation transaction of state i must be created once parties update the channel state from state i to $i + 1$. However, in the proposed scheme signatures for $\text{TX}_{\text{RV},i}$ is stored in $\text{TX}_{\text{AU},i}$ and hence $\text{TX}_{\text{RV},i}$ must actually be created once parties update the channel state from state $i - 1$ to i . It means if an honest party records the latest commit and auxiliary transactions on the blockchain, the counterparty might publish the revocation transaction and take all the channel funds.

To solve the first mentioned issue, $Y_{A,i}$ and $Y_{B,i}$ are stored in an `OP_RETURN` output that is added to the commit transaction $\text{TX}_{\text{CM},i}$. To solve the second mentioned issue, we add

two statements from the hard relation $\mathcal{R}$, $R_{A,i}$ and $R_{B,i}$, to the first sub-condition of the main output of $\text{TX}_{\text{CM},i}$, where $R_{A,i}$ ($R_{B,i}$) is generated by A (B) for the state i . Then, once the latest commit and auxiliary transactions are published by A , party B cannot record the revocation transaction as he does not know his counterparty's witness $r_{A,i}$. The witnesses $r_{A,i}$ and $r_{B,i}$ are exchanged between the parties and are given to the watchtower once parties have created $\text{TX}_{\text{CM},i+1}$, $\text{TX}_{\text{AU},i+1}$ and $\text{TX}_{\text{SP},i+1}$. Thus, $\text{TX}_{\text{FU}}.\text{txid}$, public keys pk_A and pk_B as well as r values of both parties are all data needed by the watchtower to watch the channel for both parties. Fig. 6.3 depicts the mentioned modifications.

The security requirement for r values is that B (or the watchtower) must not be able to compute $r_{A,j}$ given that he knows $r_{A,i}$ with $i < j$. Otherwise, when A submits the latest commit transaction $\text{TX}_{\text{CM},j}$, party B uses $r_{A,i}$ to compute $r_{A,j}$. Then, B publishes the revocation transaction $\text{TX}_{\text{RV},j}$ and claims its output. If r values are randomly generated, the mentioned security requirement is met but the storage cost of channel parties and the watchtower would be $\mathcal{O}(n)$. To reduce the storage and meet the stated security requirement, parties generate their r values in a binary Merkle tree and use them from the deepest leaf nodes in the tree to the root [61]. In more detail, in a binary Merkle tree, each node has two child nodes where having the value of a node, the value of each of its child nodes can simply be computed using a one-way function. But deriving the value of a node from its child nodes' values is computationally infeasible. Thus, since r values are used from the deepest leaf nodes in the tree, the stated security requirement is achieved. Moreover, the storage needed by each channel party (or the watchtower) to store r values, received from her counterparty, will be $\mathcal{O}(\log n)$ because upon receipt of a node value, its child nodes' values can be removed from the storage.




Figure 6.3: Adding Y and R Values to Commit Transactions

6.3.2.2 Reducing the Storage Requirements of channel parties

Although the storage of the watchtower is $\mathcal{O}(\log n)$, channel parties still have to store all the signatures of their counterparties on the revocation transactions. Otherwise, the dishonest channel party publishes a revoked commit transaction $\text{TX}_{\text{CM},i}$ without publishing its auxiliary transaction $\text{TX}_{\text{AU},i}$. Then, the channel funds could be locked forever. This raises a hostage situation. The scheme Outpost suffers from this problem which is why the storage requirement of channel parties is $\mathcal{O}(n)$. To solve this problem, we add one sub-condition, $Y_{A,i} \wedge Y_{B,i} \wedge 3t^+$, to the main output of the commit transaction $\text{TX}_{\text{CM},i}$. This sub-condition allows the honest channel party to claim all the channel funds in such hostage situations. In other words, if party A publishes the revoked commit transaction $\text{TX}_{\text{CM},i}$, she has $3T$ rounds time to publish $\text{TX}_{\text{AU},i}$ and $\text{TX}_{\text{SP},i}$ before B can claim the channel funds by meeting the sub-condition $Y_{A,i} \wedge Y_{B,i} \wedge 3t^+$. If during this interval, $\text{TX}_{\text{AU},i}$ is published, party B instantly establishes and publishes $\text{TX}_{\text{RV},i}$ and claims its output. To do so, each party must have r values of both parties stored. Since these keys are generated in a Merkle tree, the storage requirements of each channel party for storing these values would be $\mathcal{O}(\log n)$ (See Fig. 6.4).

Once party A publishes $\text{TX}_{\text{CM},i}$, party B must be able to use Ext algorithm to extract the value of $y_{A,i}$. To do so, he must know the corresponding pre-signature $\tilde{\sigma}_{\text{TX}_{\text{CM},i}}^B$. If parties store all their own pre-signatures, their storage cost would be $\mathcal{O}(n)$. To acquire lower storage costs, parties must be able to regenerate the required pre-signature, once a commit transaction is published. To achieve this goal, random values which are required to generate pre-signatures must be generated in a Merkle tree and be used from the root to the deepest leaf node in the tree. In this way, once the commit transaction $\text{TX}_{\text{CM},i}$ is published by A , party B can regenerate the required random value, recompute the corresponding pre-signature $\tilde{\sigma}_{\text{TX}_{\text{CM},i}}^B$ and finally extract the value of $y_{A,i}$. Thus, the storage requirements would be still $\mathcal{O}(\log n)$.

Additionally, party B must know the value of $y_{B,i}$ to meet $Y_{A,i} \wedge Y_{B,i}$. The security requirement for y values is that A must not be able to compute $y_{B,i}$ given that he knows $y_{B,j}$ with $j > i$. Otherwise, once B submits the latest commit transaction $\text{TX}_{\text{CM},j}$, A computes $y_{B,j}$ and hence derives $y_{B,i}$ with $i < j$ and then try to publish $\text{TX}_{\text{CM},i}$ before $\text{TX}_{\text{CM},j}$ being published on the ledger. Then, A might be able to claim all the channel funds by meeting the third sub-condition of the main output of $\text{TX}_{\text{CM},i}$ or by publishing the revocation transaction $\text{TX}_{\text{RV},i}$ and claiming its output. If y values are randomly generated, the mentioned security requirement is met but the parties' storage cost would be $\mathcal{O}(n)$. To reduce the storage and simultaneously meet the stated security requirement, parties generate their y values in a Merkle tree and give the corresponding Y values to their counterparties from the root to the deepest leaf nodes in the tree.




Figure 6.4: Reducing Storage Requirements of Channel Parties

6.4 Garrison Protocol Description

The lifetime of a Garrison channel can be divided into 4 phases including *create*, *update*, *close*, and *punish*. These phases will be explained in the following.

6.4.1 Create

The channel creation phase completes once the funding transaction TX_{FU} , the commit transactions $TX_{CM,0}$, the split transaction $TX_{SP,0}$, the auxiliary transaction $TX_{AU,0}$ and body of the revocation transaction $[TX_{RV,0}]$ are created, and TX_{FU} is published on the blockchain. In this phase, parties do not have access to $TX_{RV,0}$ as they have not exchanged $r_{A,0}$ and $r_{B,0}$ yet. At the end of the channel creation phase, the channel would be at state 0. Since the output of the funding transaction can only be spent if both parties agree, one party might become unresponsive to raise a hostage situation. To avoid this, parties must sign the commit, revocation, auxiliary and split transactions before signing and publishing the funding transaction. Fig. 6.5 summarises the channel creation phase.

The introduced transactions will be explained further below:

- **Funding transaction** Given that A (B , respectively) uses the x^{th} (y^{th} , respectively) output of a transaction with transaction identifier of $txid_A$ ($txid_B$, respectively) to fund the channel with a (b , respectively) coins, the funding transaction




Figure 6.5: Summary of Garrison channel creation phase.

is as follows¹:

$$\begin{aligned} TX_{FU}.Input &:= (txid_A \| x, txid_B \| y), \\ TX_{FU}.Output &:= \{(a + b, pk_A \wedge pk_B)\}, \\ TX_{FU}.Witness &:= ((1, \sigma_{TX_{FU}}^{A,1}), (1, \sigma_{TX_{FU}}^{B,2})). \end{aligned} \quad (6.1)$$

Commit transaction There exists one commit transaction per state but only the first one ($TX_{CM,i}$ with $i = 0$) is created at the channel creation phase.

$$\begin{aligned} TX_{CM,i}.Input &:= TX_{FU}.txid \| 1, \\ TX_{CM,i}.Output &:= ((a + b, \varphi_1 \vee \varphi_2 \vee \varphi_3), \\ &\quad (\epsilon, pk_A \wedge pk_B), \\ &\quad (0, (Y_{A,i}, Y_{B,i}))) \\ TX_{CM,i}.Witness &:= \{(1, \{\sigma_{TX_{CM,i}}^A, \sigma_{TX_{CM,i}}^B\})\} \end{aligned} \quad (6.2)$$

with $\varphi_1 := (pk_A \wedge pk_B \wedge R_{A,i} \wedge R_{B,i})$, $\varphi_2 := (Y_{A,i} \wedge Y_{B,i} \wedge 3t^+)$, and $\varphi_3 := (pk_A \wedge pk_B \wedge t^+)$ where $Y_{A,i}$ and $R_{A,i}$ ($Y_{B,i}$ and $R_{B,i}$) are statements of a hard relation $\mathcal{R}$ generated by A (B) for the i^{th} state and t is any number such that $t > \Delta$. The first and second outputs of the transaction are the main and auxiliary outputs. Normally, if $TX_{CM,i}$ is the last commit transaction and is published on-chain, first its auxiliary output

¹We assume that funding sources of TX_{FU} are two typical UTXOs owned by A and B .

and then its main output is spent by the auxiliary and split transactions, respectively. The third output of $TX_{CM,i}$ is an `OP_RETURN` output containing values of $Y_{A,i}$ and $Y_{B,i}$. Parties A and B use their counterparties' statements $Y_{B,i}$ and $Y_{A,i}$ and the underlying adaptor signature to generate a pre-signature on the commit transaction for their counterparties. Thus, once A publishes the commit transaction $TX_{CM,i}$, she also reveals her witness $y_{B,i}$.

Remark 6.1. *Each Bitcoin transaction can have at most one `OP_RETURN` output with the size constraint of 80 bytes. To store $Y_{A,i}$ and $Y_{B,i}$ inside an `OP_RETURN` output, their compressed version, each with a 33-byte length, are stored.*

- **Revocation transaction** The revocation transaction for state i is denoted by $TX_{RV,i}$ and is as follows, where at the channel creation phase only $TX_{RV,i}$ with $i = 0$ is created:

$$\begin{aligned} TX_{RV,i}.Input &:= TX_{CM,i}.txid\|1, \\ TX_{RV,i}.Output &:= \{(a + b, Y_{A,i} \wedge Y_{B,i})\}, \\ TX_{RV,i}.Witness &:= \{(1, \{\sigma_{TX_{RV,i}}^A, \sigma_{TX_{RV,i}}^B, r_{A,i}, r_{B,i}\})\} \end{aligned} \quad (6.3)$$

The $TX_{RV,i}$ spends the main output of $TX_{CM,i}$ using its non-time-locked sub-condition $pk_A \wedge pk_B \wedge R_{A,i} \wedge R_{B,i}$ and sends all the channel funds to an output with the condition $Y_{A,i} \wedge Y_{B,i}$. When a dishonest party, let's say A , publishes the revoked $TX_{CM,i}$, she must publish $TX_{AU,i}$ and then wait for t rounds before being able to publish $TX_{SP,i}$. However, given that the state i is revoked, B knows the value of $r_{A,i}$ and hence creates the revocation transaction $TX_{RV,i}$ and instantly publishes it on the blockchain. The output of $TX_{RV,i}$ can only be claimed by B because no one else knows the witness $y_{B,i}$.

- **Auxiliary transaction** Auxiliary transaction for state i is as follows, where at the channel creation phase only $TX_{AU,i}$ with $i = 0$ is created:

$$\begin{aligned} TX_{AU,i}.Input &:= TX_{CM,i}.txid\|2, \\ TX_{AU,i}.Output &:= ((\epsilon, pk_A \wedge pk_B \wedge t^+), \\ &\quad (0, (\sigma_{TX_{RV,i}}^A, \sigma_{TX_{RV,i}}^B))) \\ TX_{AU,i}.Witness &:= \{(1, \{\sigma_{TX_{AU,i}}^A, \sigma_{TX_{AU,i}}^B\})\} \end{aligned} \quad (6.4)$$

This transaction spends the auxiliary output of the commit transaction and its output is spent by the split transaction. In other words, the split transaction cannot

be published unless the auxiliary transaction is on the blockchain. The second output of $\text{TX}_{\text{AU},i}$ is an `OP_RETURN` output containing signatures of both parties on the corresponding revocation transaction.

Remark 6.2. *Each encoded Bitcoin signature can be up to 73 bytes long. Thus, due to the size constraint of the `OP_RETURN` output, two separate signatures do not fit into the auxiliary transaction. To solve this issue, A and B can aggregate their public keys pk_A and pk_B to form an aggregated public key pk [62] and change φ_1 in $\text{TX}_{\text{CM},i}$ to $(pk \wedge R_{A,i} \wedge R_{B,i})$. Then, rather than two separate signatures on the revocation transaction, they generate a multisignature (with up to 73-byte size) and store it inside the `OP_RETURN` output of $\text{TX}_{\text{AU},i}$.*

- **Split transaction**

$\text{TX}_{\text{SP},i}$ actually represents the i^{th} channel state and is as follows, where at the channel creation phase only the first one, $\text{TX}_{\text{SP},i}$ with $i = 0$, is created:

$$\begin{aligned} \text{TX}_{\text{SP},i}.\text{Input} &:= (\text{TX}_{\text{CM},i}.\text{txid}\|1, \text{TX}_{\text{AU},i}.\text{txid}\|1), \\ \text{TX}_{\text{SP},i}.\text{Output} &:= (\theta_1, \theta_2, \dots), \\ \text{TX}_{\text{SP},i}.\text{Witness} &:= ((3, \{\sigma_{\text{TX}_{\text{SP},i}}^A, \sigma_{\text{TX}_{\text{SP},i}}^B\}), (1, \{\sigma_{\text{TX}_{\text{SP},i}}^A, \sigma_{\text{TX}_{\text{SP},i}}^B\})) \end{aligned} \quad (6.5)$$


The split transaction spends the main output of the commit transaction and the first output of the auxiliary transaction.

6.4.2 Update

Let the channel be in state $i \geq 0$ and channel parties decide to update it to state $i + 1$. The update process is performed in two sub-phases. In the first sub-phase, channel parties create $\text{TX}_{\text{CM},i+1}$, $\text{TX}_{\text{SP},i+1}$, $\text{TX}_{\text{AU},i+1}$, and $[\text{TX}_{\text{RV},i+1}]$ for the new state. In the second sub-phase, channel parties revoke the state i by exchanging $r_{A,i}$ and $r_{B,i}$ and giving these values to the watchtower. We assume that the watchtower is also paid after each channel update. Fig. 6.6 summarises the channel update phase.

6.4.3 Close

Assume that the channel parties A and B have updated their channel n times and then A and/or B decide to close it. They can close the channel cooperatively. To do so, A and

Figure 6.6: Summary of Garrison channel update phase from state i to $i + 1$.

B create the below transaction, called modified split transaction $TX_{\overline{SP}}$, and publish it on the blockchain:

$$\begin{aligned} TX_{\overline{SP}}.Input &:= TX_{FU}.txid \| 1, \\ TX_{\overline{SP}}.Output &:= TX_{SP,n}.Output, \\ TX_{\overline{SP}}.Witness &:= \{(1, \{\sigma_{TX_{\overline{SP}}}^A, \sigma_{TX_{\overline{SP}}}^B\})\}. \end{aligned} \quad (6.6)$$

If one of the channel parties, e.g. party B , becomes unresponsive, A can still non-collaboratively close the channel. To do so, she publishes $TX_{CM,n}$ and $TX_{AU,n}$ on the ledger. Then, she waits for t rounds, and finally publishes $TX_{SP,n}$.

6.4.4 Punish

Let the channel be at state n . If a channel party, e.g. party A , publishes $TX_{CM,i}$ and then $TX_{AU,i}$ with $i < n$ on the blockchain, party B or his watchtower can create the transaction $TX_{RV,i}$ and publish it within t rounds. If only $TX_{CM,i}$ is published, party B claims its first output by meeting its second sub-condition $Y_{A,i} \wedge Y_{A,i} \wedge 3t^+$.

Remark 6.3. *If the watchtower is non-responsive, the channel might be closed in an old state. The paper [33] proposes a reputation system, called HashCached, which forces watchtowers to be responsive without requiring them to lock any funds as collateral. Garrison might be used with the HashCached system.*

6.5 Security Analysis

In this section, we prove that for the Garrison channel, it is of negligible probability that an honest party loses any funds.

Lemma 6.1. *Let Π be a EUF – CMA secure digital signature, $\mathcal{R}$ be a hard relation and Ξ be a secure adaptor digital signature. Then, for a Garrison channel with n channel updates, the broadcast of $\text{TX}_{\text{RV},i}$ with $i < n$ causes the honest channel party $P \in \{A, B\}$ to lose any funds in the channel with negligible probability.*

Proof. Without loss of generality let $P = A$. The transaction $\text{TX}_{\text{RV},i}$ with $i < n$ spends the main output of the revoked $\text{TX}_{\text{CM},i}$ and hence cannot be published unless $\text{TX}_{\text{CM},i}$ is on-chain. The transaction $\text{TX}_{\text{CM},i}$ spends the output of TX_{FU} . Since the condition in TX_{FU} .Output contains pk_A , this output cannot be spent without A 's authorisation. Otherwise, the security of the underlying digital signature would be violated. Based on the protocol, the honest party A never broadcasts the revoked $\text{TX}_{\text{CM},i}$ on-chain and her pre-signature $\tilde{\sigma}_{\text{TX}_{\text{CM},i}}$ on the transaction $\text{TX}_{\text{CM},i}$ is the only authorisation he grants for spending TX_{FU} .Output using $\text{TX}_{\text{CM},i}$. Thus, if $\text{TX}_{\text{CM},i}$ is published, the probability that A fails to obtain $y_{B,i}$ is negligible. Otherwise, a EUF – CMA security or witness extractability of the used adaptor signature is violated. Furthermore, $\text{TX}_{\text{RV},i}$ has only one output with the condition of $Y_{A,i} \wedge Y_{B,i}$ and the value of $a + b$. Since A privately preserves its witness value $y_{A,i}$, the probability that any PPT adversary claims $\text{TX}_{\text{RV},i}$.Output is negligible. Otherwise, the utilised hard relation would break. Therefore, it is of negligible probability that A (who knows both $y_{A,i}$ and $y_{B,i}$) fails to claim $\text{TX}_{\text{RV},i}$.Output and obtain all the channel funds. $\square$

Lemma 6.2. *Let Π be a EUF – CMA secure digital signature, $\mathcal{R}$ be a hard relation and Ξ be a secure adaptor digital signature. Then, for a Garrison channel between A and B with $P \in \{A, B\}$ being the honest party, if P 's counterparty publishes $\text{TX}_{\text{CM},i}$, it is with negligible probability that*

- P fails to obtain the data required to meet the second sub-condition of $\text{TX}_{\text{CM},i}$.Output[1]. φ .
- any PPT adversary can meet the second sub-condition of $\text{TX}_{\text{CM},i}$.Output[1]. φ .

Proof. Without loss of generality let $P = A$. Similar to the proof of Lemma 6.1, if B publishes $\text{TX}_{\text{CM},i}$, the probability that A fails to obtain $y_{B,i}$ is negligible. Otherwise, aEUf – CMA security or witness extractability of the used adaptor signature is violated. The witness $y_{A,i}$ has also been created by A and hence he has the whole data required to meet $Y_{A,i} \wedge Y_{B,i} \wedge 3t^+$. Furthermore, given that A privately preserves its witness value $y_{A,i}$, the probability that any PPT adversary meets this sub-condition is negligible. Otherwise, the utilised hard relation would break. $\square$

Lemma 6.3. *Let Π be a EUF – CMA secure digital signature, $\mathcal{R}$ be a hard relation and Ξ be a secure adaptor digital signature. Then, for a Garrison channel with n channel updates, if the honest party $P \in \{A, B\}$ publishes $\text{TX}_{\text{CM},n}$, P loses funds in the channel with negligible probability.*

Proof. Without loss of generality let $P = A$. We assume that A publishes $\text{TX}_{\text{CM},n}$ in the block $\mathcal{B}_j$ of the blockchain and prove that it is of negligible probability that A fails to publish $\text{TX}_{\text{SP},n}$. Then, since $\text{TX}_{\text{SP},n}$ corresponds with the latest channel state, its broadcast cannot cause A to lose any funds.

The condition $\text{TX}_{\text{CM},n}.\text{Output}[2].\varphi$ contains pk_A and hence it is of negligible probability that this output is spent without A 's authorisation. Otherwise, the security of the underlying digital signature is violated. The honest party A grants such an authorisation only on the transaction $\text{TX}_{\text{AU},n}$ which is held by both A and B . Based on the protocol, once $\text{TX}_{\text{CM},n}$ is published on the blockchain by A , he also instantly submits $\text{TX}_{\text{AU},n}$ to the blockchain. According to our assumptions regarding the blockchain, $\text{TX}_{\text{AU},n}$ is published on the blockchain in the block $\mathcal{B}_{j+k}$ with $0 < k \leq \Delta < t$. Similarly, the first output of $\text{TX}_{\text{AU},n}$ can only be spent by $\text{TX}_{\text{SP},n}$. According to the protocol, A holds $\text{TX}_{\text{SP},n}$ and submits it to the blockchain t rounds after $\text{TX}_{\text{AU},n}$ is published on-chain. Thus, given that the first input of $\text{TX}_{\text{SP},n}$ (or equivalently the first output of $\text{TX}_{\text{CM},n}$) is still unspent, based on our assumptions regarding the blockchain, $\text{TX}_{\text{SP},n}$ is published on the blockchain in the block $\mathcal{B}_{j+k+l+t}$ with $0 < l \leq \Delta < t$. Now, we prove that, when $\mathcal{B}_{j+k+l+t}$ with $0 < l, k < t$ is added to the blockchain, the first output of $\text{TX}_{\text{CM},n}$, $\text{TX}_{\text{CM},n}.\text{Output}[1]$, is still unspent.

The first and third sub-conditions of $\text{TX}_{\text{CM},n}.\text{Output}[1]$ contains $R_{A,n}$ and pk_A , respectively and hence it is of negligible probability that these two sub-conditions are met without A 's authorisation. Otherwise, the underlying hard relation or digital signature would break. Party A grants such an authorisation only on $\text{TX}_{\text{SP},n}$. Moreover, the second sub-condition $Y_{A,i} \wedge Y_{B,i} \wedge 3t^+$ cannot be met in block $\mathcal{B}_{j+k+l+t}$ with $0 < l, k < t$ because $j + k + l + t < j + 3t$. $\square$

Lemma 6.4. *Let Π be a EUF – CMA secure digital signature, $\mathcal{R}$ be a hard relation and Ξ be a secure adaptor digital signature. Then, for a Garrison channel with n channel updates*

and with $P \in \{A, B\}$ being the honest party, if P 's counterparty publishes $\text{TX}_{\text{CM},n}$, it is of negligible probability that P loses any funds in the channel.

Proof. Without loss of generality let $P = A$. The proof is similar to the proof of Lemma 6.3. The only difference is that following Lemma 6.2, it is of negligible probability that A fails to meet the second sub-condition of $\text{TX}_{\text{CM},n}.\text{Output}[1]$. Therefore, A can either publishes both $\text{TX}_{\text{AU},n}$ and $\text{TX}_{\text{SP},n}$ or claim $\text{TX}_{\text{CM},n}.\text{Output}[1]$ by meeting its second sub-condition. None of these two cases can cause the honest party A to lose any funds in the channel. $\square$

Lemma 6.5. *Let Π be a EUF – CMA secure digital signature, $\mathcal{R}$ be a hard relation and Ξ be a secure adaptor digital signature. Then, for a Garrison channel with n channel updates and with $P \in \{A, B\}$ being the honest party, if any adversary publishes $\text{TX}_{\text{CM},i}$ with $i < n$, it is of negligible probability that P loses any funds in the channel.*

Proof. Without loss of generality let $P = A$. The output $\text{TX}_{\text{CM},i}.\text{Output}[1]$ includes 3 sub-conditions, one of which must be met to cheat A out of its funds. The first sub-condition contains pk_A and hence it is of negligible probability that this output is spent without A 's authorisation. Otherwise, the security of the used digital signature is violated. The honest party A grants such an authorisation only on the transaction $\text{TX}_{\text{RV},i}$. However, according to Lemma 6.1, it is of negligible probability that broadcast of $\text{TX}_{\text{RV},i}$ causes A to lose any funds. Moreover, according to Lemma 6.2, it is of negligible probability that any PPT adversary can meet the second sub-condition. Now, we prove that if the third sub-condition is used to cheat A out of her funds, it leads to a contradiction.

Assume that the third sub-condition of $\text{TX}_{\text{CM},i}.\text{Output}[1]$ is used to cheat A out of her funds. This sub-condition contains pk_A and hence it is of negligible probability that this condition is met without A 's authorisation. Otherwise, the security of the underlying digital signature is violated. The honest party A grants such an authorisation only on the transaction $\text{TX}_{\text{SP},i}$. Assume that $\text{TX}_{\text{SP},i}$ is included in the block $\mathcal{B}_k$ of the blockchain. The transaction $\text{TX}_{\text{SP},i}$ cannot be added to the blockchain unless its inputs are some unspent outputs on the blockchain. It means that $\text{TX}_{\text{AU},i}$ is also on the blockchain and following the condition in $\text{TX}_{\text{AU},i}.\text{Output}[1]$, the transaction $\text{TX}_{\text{AU},i}$ must have been published in the block $\mathcal{B}_j$ with $j \leq k - t$. However, based on the protocol, once A or her watchtower observes $\text{TX}_{\text{AU},i}$ on the blockchain, they create the corresponding revocation transaction $\text{TX}_{\text{RV},i}$ and submit it to the blockchain. According to our blockchain assumptions, this transaction is published on the blockchain in block $\mathcal{B}_l$ with $j < l \leq j + \Delta < j + t \leq k$. However, once $\text{TX}_{\text{RV},i}$ is published in the block $\mathcal{B}_l$ of the blockchain, the transaction $\text{TX}_{\text{SP},i}$ becomes invalid and cannot be published in block $\mathcal{B}_k$ of the blockchain which leads to a contradiction. $\square$

Theorem 6.1. *Let Π be a EUF – CMA secure digital signature, $\mathcal{R}$ be a hard relation and Ξ be a secure adaptor digital signature. Then, for a Garrison channel, an honest party $P \in \{A, B\}$ loses any funds in the channel with negligible probability.*

Proof. Without loss of generality let $P = A$. Funds of A are locked in TX_{FU} .Output. It is of negligible probability that any PPT adversary $\mathcal{A}$ spends the output of TX_{FU} without the honest party A 's authorisation. Otherwise, the underlying digital signature would be forgeable. Furthermore, $\text{TX}_{\overline{\text{SP}}}$, $\text{TX}_{\text{CM},i}$ with $i = [0, n-1]$, $\text{TX}_{\text{CM},n}$ are the only transactions in the protocol that spend the output of TX_{FU} and A grants authorisation for. Thus, these transactions will be discussed further. Since $\text{TX}_{\overline{\text{SP}}}$ represents the final agreed state of the channel, its broadcast cannot cause A to lose any funds. Moreover, according to Lemmas 6.3 and 6.4, it is of negligible probability that broadcast of $\text{TX}_{\text{CM},n}$ causes A to be cheated out of her funds. Also, based on the protocol, A never publishes $\text{TX}_{\text{CM},i}$ with $i = [0, n-1]$ and according to Lemma 6.5, if one of these transactions is published by the adversary, it causes A to lose any funds with negligible probability. This concludes the proof. $\square$

6.6 Garrison Transactions Scripts

The funding transaction is similar to the corresponding transaction in the FPPW channel. Commit transaction has one input that takes the output of the funding transaction with witness script of 0 $\langle \text{pubkeyA_sig} \rangle \langle \text{pubkeyB_sig} \rangle$. It also has three outputs where the script of its first output (main output) is as follows:

OP_IF

Revocation

4 $\langle \text{Rev_pubkeyA} \rangle \langle \text{Rev_pubkeyB} \rangle \langle \text{R_pubkeyA} \rangle \langle \text{R_pubkeyB} \rangle$ 4 OP_CHECKMULTISIG

OP_ELSE

OP_IF

$\langle \text{delay } 3t \rangle$ OP_CHECKSEQUENCEVERIFY OP_DROP

2 $\langle \text{Y_pubkeyA} \rangle \langle \text{Y_pubkeyB} \rangle$ 2 OP_CHECKMULTISIG

OP_ELSE

Split

$\langle \text{delay } t \rangle$ OP_CHECKSEQUENCEVERIFY OP_DROP

2 $\langle \text{Spl_pubkeyA} \rangle \langle \text{Spl_pubkeyB} \rangle$ 2 OP_CHECKMULTISIG

OP_ENDIF

OP_ENDIF

where $\langle \text{Rev_pubkeyA} \rangle$ and $\langle \text{Spl_pubkeyA} \rangle$ are public keys of A and $\langle \text{R_pubkeyA} \rangle$ and $\langle \text{Y_pubkeyA} \rangle$ are revocation key and publishing key of A , respectively. Also, $\langle \text{Rev_pubkeyB} \rangle$, $\langle \text{Spl_pubkeyB} \rangle$, $\langle \text{R_pubkeyB} \rangle$ and $\langle \text{Y_pubkeyB} \rangle$ are the corresponding parameters for B .

The script for the second output (auxiliary output) of the commit transaction is as following:

$2 \langle \text{Aux_pubkeyA} \rangle \langle \text{Aux_pubkeyB} \rangle 2 \text{ OP_CHECKMULTISIG}$

where $\langle \text{Aux_pubkeyA} \rangle$ and $\langle \text{Aux_pubkeyB} \rangle$ are public keys of A and B respectively.

The script for the third output of the commit transaction is as following:

$\text{OP_RETURN } \langle \text{Y_pubkeyA} \rangle \langle \text{Y_pubkeyB} \rangle$

where $\langle \text{Y_pubkeyA} \rangle$ and $\langle \text{Y_pubkeyB} \rangle$ are the same as the corresponding values used in the script in the first output of the commit transaction.

The revocation transaction spends the main output of a revoked commit transaction with the witness script of $0 \langle \text{Rev_pubkeyA_sig} \rangle \langle \text{Rev_pubkeyB_sig} \rangle \langle \text{R_pubkeyA_sig} \rangle \langle \text{R_pubkeyB_sig} \rangle 1$. It also has one output with the script of $2 \langle \text{Y_pubkeyA} \rangle \langle \text{Y_pubkeyB} \rangle 2 \text{ OP_CHECKMULTISIG}$.

The auxiliary transaction spends the auxiliary output of the commit transaction using the witness script of $0 \langle \text{Aux_pubkeyA_sig} \rangle \langle \text{Aux_pubkeyB_sig} \rangle$. It also has two outputs. Its first output contains the following script:

$\langle \text{delay } t \rangle \text{ OP_CHECKSEQUENCEVERIFY OP_DROP}$

$2 \langle \text{Spl_pubkeyA} \rangle \langle \text{Spl_pubkeyB} \rangle 2$

The second output of the auxiliary transaction is as follows:

$\text{OP_RETURN } \langle \text{Rev_pubkeyA_sig} \rangle \langle \text{Rev_pubkeyB_sig} \rangle$ where $\langle \text{Rev_pubkeyA_sig} \rangle$ and $\langle \text{Rev_pubkeyB_sig} \rangle$ are the corresponding signatures used in the witness script of the corresponding revocation transaction.

The split transaction has two inputs. The first one takes the main output of the corresponding commit transaction with the witness script of $0 \langle \text{Spl_pubkeyA_Sig} \rangle \langle \text{Spl_pubkeyB_Sig} \rangle 0 0$. The second input takes the first output of the corresponding auxiliary transaction with the witness script of $0 \langle \text{Spl_pubkeyA_Sig} \rangle \langle \text{Spl_pubkeyB_Sig} \rangle$.

6.7 Garrison Protocol

In this section, protocols for different phases of Garrison will be presented. In different steps of the protocol, channel participants generate (or verify) some signatures or pre-signatures on protocol transactions. When a signature or pre-signature is going to be generated (or verified) for j^{th} input of the transaction TX_i , the input message to the signing (or verification) algorithm is denoted by $f([\text{TX}_i], j)$ [58].

Garrison channel creation protocol is as follows:

Preconditions: A and B own $a + \epsilon/2$ and $b + \epsilon/2$ coins on-chain in output of transactions with transaction identifiers txid_A and txid_B respectively. A and B know each other's public keys and values of ϵ , a and b that are going to be used in the channel.

1. Create $[\text{TX}_{\text{FU}}]$:

$$(a) P \in \{A, B\} \xrightarrow{\text{txid}_P} \bar{P}$$

(b) If P receives $\text{txid}_{\bar{P}}$, it creates $[\text{TX}_{\text{FU}}]$ according to 6.1. Else it stops.

2. Create $[\text{TX}_{\text{CM},0}]$:

(a) $P \in \{A, B\}$ generates $(Y_{P,0}, y_{P,0}) \leftarrow \text{GenR}$ and $(R_{P,0}, r_{P,0}) \leftarrow \text{GenR}$.

$$(b) P \xrightarrow{Y_{P,0}, R_{P,0}} \bar{P}$$

(c) If P receives $Y_{\bar{P},0}$, it creates $[\text{TX}_{\text{CM},0}]$ according to 6.2. Else it stops.

3. Create $[\text{TX}_{\text{RV},0}]$:

(a) Party $P \in \{A, B\}$ creates $[\text{TX}_{\text{RV},0}]$ according to 6.3.

(b) Party P computes $\sigma_{\text{TX}_{\text{RV},0}}^P = \text{Sign}_{sk_P}(f([\text{TX}_{\text{RV},0}]))$.

$$(c) P \xrightarrow{\sigma_{\text{TX}_{\text{RV},0}}^P} \bar{P}.$$

(d) If party P receives $\sigma_{\text{TX}_{\text{RV},0}}^{\bar{P}}$ from $\bar{P}$ s.t. $\text{Vrfy}_{pk_{\bar{P}}}(f([\text{TX}_{\text{RV},0}], 1); \sigma_{\text{TX}_{\text{RV},0}}^{\bar{P}}) = 1$, it continues. Else it stops.

4. Create $[\text{TX}_{\text{AU},0}]$: Party $P \in \{A, B\}$ creates $[\text{TX}_{\text{AU},0}]$ according to 6.4.

5. Create $[\text{TX}_{\text{SP},0}]$: Party $P \in \{A, B\}$ creates $[\text{TX}_{\text{SP},0}]$ according to 6.5.

6. Create $\text{TX}_{\text{SP},0}$:

(a) Party $P \in \{A, B\}$ computes $\sigma_{\text{TX}_{\text{SP},0}}^{P,j} = \text{Sign}_{sk_P}(f([\text{TX}_{\text{SP},0}], j))$ with $j := \{1, 2\}$.

- (b) $P \xrightarrow{\sigma_{\text{TX}_{\text{SP},0}}^{P,1}, \sigma_{\text{TX}_{\text{SP},0}}^{P,2}} \bar{P}$
- (c) If party P receives $\sigma_{\text{TX}_{\text{SP},0}}^{\bar{P},j}$ with $j := \{1, 2\}$ s.t. $\text{Vrfy}_{pk_{\bar{P}}}(f([\text{TX}_{\text{SP},0}], j); \sigma_{\text{TX}_{\text{SP},0}}^{\bar{P},j}) = 1$ for $j := \{1, 2\}$, it continues. Else it stops.
- (d) Party P creates $\text{TX}_{\text{SP},0}$ according to 6.5.
7. Create $\text{TX}_{\text{AU},0}$:
- (a) Party $P \in \{A, B\}$ computes $\sigma_{\text{TX}_{\text{AU},0}}^P = \text{Sign}_{sk_P}(f([\text{TX}_{\text{AU},0}], 1))$.
- (b) $P \xrightarrow{\sigma_{\text{TX}_{\text{AU},0}}^P} \bar{P}$
- (c) If party P receives $\sigma_{\text{TX}_{\text{AU},0}}^{\bar{P}}$ s.t. $\text{Vrfy}_{pk_{\bar{P}}}(f([\text{TX}_{\text{AU},0}], 1); \sigma_{\text{TX}_{\text{AU},0}}^{\bar{P}}) = 1$, it continues. Else it stops.
- (d) Party P creates $\text{TX}_{\text{AU},0}$ according to 6.4.
8. Create $\text{TX}_{\text{CM},0}$:
- (a) Party $P \in \{A, B\}$ computes $\tilde{\sigma}_{\text{TX}_{\text{CM},0}}^P = \text{pSign}_{sk_P}(f([\text{TX}_{\text{CM},0}], 1), Y_{\bar{P},0})$.
- (b) $P \xrightarrow{\tilde{\sigma}_{\text{TX}_{\text{CM},0}}^P} \bar{P}$
- (c) If party P receives $\tilde{\sigma}_{\text{TX}_{\text{CM},0}}^{\bar{P}}$ s.t. $\text{pVrfy}_{pk_{\bar{P}}}(f([\text{TX}_{\text{CM},0}], 1), Y_{P,0}; \tilde{\sigma}_{\text{TX}_{\text{CM},0}}^{\bar{P}}) = 1$, it computes $\sigma_{\text{TX}_{\text{CM},0}}^{\bar{P}} = \text{Adapt}(\tilde{\sigma}_{\text{TX}_{\text{CM},0}}^{\bar{P}}, y_{P,0})$, computes $\sigma_{\text{TX}_{\text{CM},0}}^P = \text{Sign}_{sk_P}(f([\text{TX}_{\text{CM},0}], 1))$, creates $\text{TX}_{\text{CM},0}$ according to 6.2 and continues. Else it stops.
9. Create TX_{FU} :
- (a) Party $P \in \{A, B\}$ computes $\sigma_{\text{TX}_{\text{FU}}}^P = \text{Sign}_{sk_P}(f([\text{TX}_{\text{FU}}], j))$ where $j := 1$ if $P = A$ or $j := 2$ otherwise.
- (b) $P \xrightarrow{\sigma_{\text{TX}_{\text{FU}}}^P} \bar{P}$
- (c) If party P receives $\sigma_{\text{TX}_{\text{FU}}}^{\bar{P}}$ s.t. $\text{Vrfy}_{pk_{\bar{P}}}(f([\text{TX}_{\text{FU}}], j); \sigma_{\text{TX}_{\text{FU}}}^{\bar{P}}) = 1$ with $j := 1$ if $P = B$ or $j := 2$ otherwise, it continues. Else it stops.
- (d) Party P creates TX_{FU} .
10. Publish TX_{FU} : Party $P \in \{A, B\}$ publishes TX_{FU} on-chain.

The Garrison channel update protocol is as follows:

Preconditions: The channel create phase is complete and TX_{FU} is on-chain. The channel update phase has been completed i times and hence the channel is at state i .

1. Create $[\text{TX}_{\text{CM},i+1}]$:

(a) $P \in \{A, B\}$ generates $(Y_{P,i+1}, y_{P,i+1}) \leftarrow \text{GenR}$ and $(R_{P,i+1}, r_{P,i+1}) \leftarrow \text{GenR}$.

(b) $P \xrightarrow{Y_{P,i+1}, R_{P,i+1}} \bar{P}$

(c) If P receives $Y_{\bar{P},i+1}$, it creates $[\text{TX}_{\text{CM},i+1}]$ according to 6.2. Else it stops.

2. Create $[\text{TX}_{\text{RV},i+1}]$:

(a) Party $P \in \{A, B\}$ creates $[\text{TX}_{\text{RV},i+1}]$ according to 6.3.

(b) Party P computes $\sigma_{\text{TX}_{\text{RV},i+1}}^P = \text{Sign}_{sk_P}(f([\text{TX}_{\text{RV},i+1}]))$.

(c) $P \xrightarrow{\sigma_{\text{TX}_{\text{RV},i+1}}^P} \bar{P}$.

(d) If party P receives $\sigma_{\text{TX}_{\text{RV},i+1}}^{\bar{P}}$ from $\bar{P}$ s.t. $\text{Vrfy}_{pk_{\bar{P}}}(f([\text{TX}_{\text{RV},i+1}], 1); \sigma_{\text{TX}_{\text{RV},i+1}}^{\bar{P}}) = 1$, it continues. Else it stops.

3. Create $[\text{TX}_{\text{AU},i+1}]$: Party $P \in \{A, B\}$ creates $[\text{TX}_{\text{AU},i+1}]$ according to 6.4.

4. Create $[\text{TX}_{\text{SP},i+1}]$: Party $P \in \{A, B\}$ creates $[\text{TX}_{\text{SP},i+1}]$ according to 6.5.

5. Create $\text{TX}_{\text{SP},i+1}$:

(a) Party $P \in \{A, B\}$ computes $\sigma_{\text{TX}_{\text{SP},i+1}}^{P,j} = \text{Sign}_{sk_P}(f([\text{TX}_{\text{SP},i+1}], j))$ with $j := \{1, 2\}$.

(b) $P \xrightarrow{\sigma_{\text{TX}_{\text{SP},i+1}}^{P,1}, \sigma_{\text{TX}_{\text{SP},i+1}}^{P,2}} \bar{P}$

(c) If party P receives $\sigma_{\text{TX}_{\text{SP},i+1}}^{\bar{P},j}$ with $j := \{1, 2\}$ s.t. $\text{Vrfy}_{pk_{\bar{P}}}(f([\text{TX}_{\text{SP},i+1}], j); \sigma_{\text{TX}_{\text{SP},i+1}}^{\bar{P},j}) = 1$ for $j := \{1, 2\}$, it continues. Else it stops.

(d) Party P creates $\text{TX}_{\text{SP},i+1}$ according to 6.5.

6. Create $\text{TX}_{\text{AU},i+1}$:

(a) Party $P \in \{A, B\}$ computes $\sigma_{\text{TX}_{\text{AU},i+1}}^P = \text{Sign}_{sk_P}(f([\text{TX}_{\text{AU},i+1}], 1))$.

(b) $P \xrightarrow{\sigma_{\text{TX}_{\text{AU},i+1}}^P} \bar{P}$

- (c) If party P receives $\sigma_{\text{TX}_{\text{AU},i+1}}^{\bar{P}}$ s.t. $\text{Vrfy}_{pk_{\bar{P}}}(f([\text{TX}_{\text{AU},i+1}], 1); \sigma_{\text{TX}_{\text{AU},i+1}}^{\bar{P}}) = 1$, it continues. Else it stops.
 - (d) Party P creates $\text{TX}_{\text{AU},i+1}$ according to 6.4.
7. Create $\text{TX}_{\text{CM},i+1}$:
- (a) Party $P \in \{A, B\}$ computes $\tilde{\sigma}_{\text{TX}_{\text{CM},i+1}}^P = \text{pSign}_{sk_P}(f([\text{TX}_{\text{CM},i+1}], 1), Y_{\bar{P},i+1})$.
 - (b) $P \xrightarrow{\tilde{\sigma}_{\text{TX}_{\text{CM},i+1}}^P} \bar{P}$
 - (c) If party P receives $\tilde{\sigma}_{\text{TX}_{\text{CM},i+1}}^{\bar{P}}$ s.t. $\text{pVrfy}_{pk_{\bar{P}}}(f([\text{TX}_{\text{CM},i+1}], 1), Y_{P,i+1}; \tilde{\sigma}_{\text{TX}_{\text{CM},i+1}}^{\bar{P}}) = 1$, it computes $\sigma_{\text{TX}_{\text{CM},i+1}}^{\bar{P}} = \text{Adapt}(\tilde{\sigma}_{\text{TX}_{\text{CM},i+1}}^{\bar{P}}, y_{P,i+1})$, computes $\sigma_{\text{TX}_{\text{CM},i+1}}^P = \text{Sign}_{sk_P}(f([\text{TX}_{\text{CM},i+1}], 1))$, creates $\text{TX}_{\text{CM},i+1}$ according to 6.2 and continues. Else it executes the non-collaborative closure phase (from P 's point of view the channel is still at state i).
8. Revoke $\text{TX}_{\text{CM},i}$:
- (a) $P \in \{A, B\} \xrightarrow{r_{P,i}} \bar{P}$.
 - (b) if P receives $r_{\bar{P},i}$ s.t. $(R_{\bar{P},i}, r_{\bar{P},i}) \in \mathcal{R}$, then continues. Else, it executes the non-collaborative closure phase (from P 's point of view the channel is at state $i + 1$).

Garrison channel collaborative closure protocol is as follows:

Preconditions: The channel create phase is complete and TX_{FU} is on-chain. The channel is at state n .

1. Create TX_{SP} :
 - (a) Party $P \in \{A, B\}$ creates $[\text{TX}_{\text{SP}}]$ according to 6.6.
 - (b) P computes $\sigma_{\text{TX}_{\text{SP}}}^P = \text{Sign}_{sk_P}(f([\text{TX}_{\text{SP}}], 1))$.
 - (c) $P \xrightarrow{\sigma_{\text{TX}_{\text{SP}}}^P} \bar{P}$
 - (d) If P receives $\sigma_{\text{TX}_{\text{SP}}}^{\bar{P}}$ from $\bar{P}$ s.t. $\text{Vrfy}_{pk_{\bar{P}}}(f([\text{TX}_{\text{SP}}], 1); \sigma_{\text{TX}_{\text{SP}}}^{\bar{P}}) = 1$, it continues. Else it executes the non-collaborative closure phase (from P 's point of view the channel is still at state n).
2. Publish TX_{SP} : Party $P \in \{A, B\}$ publishes TX_{SP} on-chain.

Garrison channel non-collaborative closure protocol is as follows:

Preconditions: The channel create phase is complete and TX_{FU} is on-chain. The channel is at state n .

1. Party $P \in \{A, B\}$ publishes $\text{TX}_{\text{CM},n}$ on-chain.
2. Once $\text{TX}_{\text{CM},n}$ is recorded on-chain, P waits for t rounds and then publishes $\text{TX}_{\text{SP},n}$ on-chain.

The protocol for penalizing the cheating party is as following:

preconditions: The channel create phase is complete and TX_{FU} is on-chain. The channel is at state $n > 0$. $\text{TX}_{\text{CM},i}$ with $i < n$ is recorded on-chain by a dishonest party. The honest party (or its watchtower) is always online.

1. The honest party P (or its watchtower) observes that $\text{TX}_{\text{CM},i}$ is on-chain. Party P (or its watchtower) extracts $y_{\bar{P},i} = \text{Ext}(\sigma_{\text{TX}_{\text{CM},i}}^P, \tilde{\sigma}_{\text{CM},i}^P, Y_{\bar{P},i})$.
2. If $\text{TX}_{\text{AU},i}$ is published within $3t$ rounds, Party P (or its watchtower) extracts $\sigma_{\text{TX}_{\text{RV},i}}^P$ and $\sigma_{\text{TX}_{\text{RV},i}}^{\bar{P}}$ from $\text{TX}_{\text{AU},i}.\text{Output}[2]$, computes $\sigma_{\text{TX}_{\text{RV},i}}^{\prime P} = \text{Sign}_{r_{P,i}}(f([\text{TX}_{\text{RV},i}], 1))$ and $\sigma_{\text{TX}_{\text{RV},i}}^{\prime \bar{P}} = \text{Sign}_{r_{\bar{P},i}}(f([\text{TX}_{\text{RV},i}], 1))$, creates $\text{TX}_{\text{RV},i}$ according to 6.3, publishes it on the blockchain, claims its output and stops. Else, party P claims the output $\text{TX}_{\text{CM},i}.\text{Output}[2]$ and stops.

6.8 Conclusion

In this chapter, we presented a payment channel with a watchtower, called Garrison, whose storage costs for both channel parties and the watchtower are logarithmic in the maximum number of channel updates. Garrison avoids state duplication. So, the number of transactions does not increase exponentially with the number of payment channels or applications built on top of each other. Furthermore, the number of transactions that are published on the blockchain upon fraud does not increase with the number of outputs in the published revoked state.

Regarding the properties defined in Chapter 4, Garrison provides agility as the watchtower contract might be made with any watchtower and it can initiate and terminate at any time. Like Monitor, DCWC, DCWC* and Outpost, Garrison provides strong privacy against the watchtower. Also, it potentially achieves weak watchtower privacy against the third party (Refer to Section 4.3.2 for more details.). The watchtower in Garrison is rewarded upon each channel update. So, Garrison provides watchtower fairness but since the watchtower does not lock any collateral per channel, the channel party might

Table 6.2: Different Properties of the Garrison Scheme.

	Bitcoin Support	Agility	Priv. ag. Watch.	Priv. ag. 3 rd Party	Watch. Fairness	α	β
Garrison	Yes	Yes	Strong	Weak	Yes	0	1

lose her funds and hence Garrison is unfair to the channel party. Correspondingly, it provides β -coverage with $\beta = 1$. Adding channel party fairness to Garrison has been left to future works. Table 6.2 summarises the mentioned properties for Garrison.

Chapter 7

Daric: a storage efficient channel with penalisation

We construct this chapter based on the full version of our published paper, “Daric: A Storage Efficient Payment Channel With Penalisation Mechanism” [63] (The full version paper is available at <https://eprint.iacr.org/2022/1295.pdf>).

7.1 Introduction

While the Lightning Network exhibits elegant design, it does have certain limitations. One such limitation is the linear increase in storage requirements for channel parties, as they must store all revocation secrets received from their counterparts as the number of channel updates grows. This leads to an escalation in the storage expenses, particularly for channel parties managing a large number of channels and frequent channel updates (e.g., hubs in the PCH use case). Additionally, to ensure the identification and penalization of misbehaving parties, the channel state is duplicated, resulting in each party having its own copy of the state. Then, if parties split their channel into sub-channels (for example in order to add an application like a *Virtual channel* [27] on top of the channel), the state of each sub-channel is duplicated and it must propagate on both duplicates of the parent channel. Thus, state duplication causes the number of transactions to exponentially rise with the number of applications k built on top of each other [14].

Towards a different direction, the payment channel eltoo [1] introduces ANYPREVOUT [31] (also known as NOINPUT) as a new Bitcoin signature type to deploy the concept of versioning. This allows channel parties to override the current channel state by creating a state with a higher version number, which can be published upon fraud. So, channel

parties in eltoo do not store any revocation secrets from old channel states. This simplifies the key management and offers more affordable watchtowers as the transaction with the highest version invalidates all previous states. Furthermore, if an honest party forgets about an update and publishes an outdated state, it does not result in the loss of funds.

However, eltoo is incentive incompatible because its lack of punishment might encourage a dishonest party to publish an old state; either the other side corrects it or the dishonest party wins [64]. The only discouraging factor—the fee for publishing the old state—is also determined by the dishonest party. Thus, she can set it to the minimum possible value, i.e. few cents for some blockchains such as Bitcoin hard forks (e.g. Litecoin and Bitcoin Cash) and less than 1 USD for Bitcoin. Moreover, the transaction fee is independent of the channel *capacity* (i.e. the total funds in the channel). Therefore, even for payment channels with a huge capacity of several BTCs (e.g., channels listed in [65]), the dishonest party’s cost will be still below 1 USD (See Section 7.6.2 for detailed analysis). Additionally, enforcing a large transaction fee or restricting the channel capacity (proposed in [66]) might be unfavourable to the honest party.

Furthermore, a dishonest party in eltoo might publish multiple outdated states to delay the channel closure process [67]. Thus, eltoo fails to achieve *bounded closure*, i.e. honest party is not guaranteed that the channel closure completes within a bounded time. This compromises the security of time-based payments, e.g. *Hashed Time-Lock Contract* (HTLC) (See Section 7.6.1 for further analysis).

Therefore, the main motivation of this chapter is designing a Bitcoin payment channel that (1) provides constant storage, (2) achieves bounded closure, (3) provides incentive compatibility, and (4) avoids state duplication.

The contributions of this chapter are as follows:

- We present a new Bitcoin payment channel, called Daric, which (1) is provably secure in the Universal Composability (UC) framework, (2) achieves constant size storage for both channel parties and the watchtower, (3) provides bounded closure, (4) provides punishment mechanism and hence achieves incentive compatibility, (5) avoids state duplication without needing any particular property (e.g. adaptor signature properties) for the underlying digital signature, and (6) attains unlimited lifetime, given that channel parties on average pay each other at most once per second. Table 7.1 compares Daric with other Bitcoin payment channels.
- We compare Daric and eltoo and show Daric is robust against an attack [67] to eltoo, which we also formalise in this chapter. We further perform a cost-benefit

Table 7.1: Comparison of different payment channels with n channel updates and k recursive channel splitting.

Scheme	Party's St. Req.	Watch. St. Req.	Unl. Life.	Incent. Comp.	# of TxS	Ada. Sig. Avoid.	Bnd. Cls.
Lightning [†] [15]	$\mathcal{O}(n)$	$\mathcal{O}(n)$	Y	Y	$\mathcal{O}(2^k)$	Y	Y
Generalized [†] [14]	$\mathcal{O}(n)$	$\mathcal{O}(n)$	Y	Y	$\mathcal{O}(1)$	N	Y
Outpost [22]	$\mathcal{O}(n)$	$\mathcal{O}(\log n)$	N	Y	$\mathcal{O}(2^k)$	Y	Y
FPPW [54]	$\mathcal{O}(n)$	$\mathcal{O}(n)$	Y	Y	$\mathcal{O}(1)$	N	Y
Cerberus [23]	$\mathcal{O}(n)$	$\mathcal{O}(n)$	Y	Y	$\mathcal{O}(2^k)$	Y	Y
Sleepy [†] [28]	$\mathcal{O}(n)$	N/A	N	Y	$\mathcal{O}(2^k)$	Y	Y
eltoo [1]	$\mathcal{O}(1)$	$\mathcal{O}(1)$	Y [§]	N	$\mathcal{O}(1)$	Y	N
Daric (this work)	$\mathcal{O}(1)$	$\mathcal{O}(1)$	Y [§]	Y	$\mathcal{O}(1)$	Y	Y

[†]: If parties pre-generate n keys in a Merkle tree, their storage requirements decrease to $\mathcal{O}(\log n)$ but the channel lifetime becomes limited to n channel updates.

[§]: Given that the channel update rate is at most one update per second.

analysis to assess the attacker's revenue in practice. We also show (1) Daric provides a higher deterrent effect against profit-driven attackers than eltoo and (2) unlike eltoo, Daric's deterrent effect is flexible.

- We compare Daric, eltoo, Lightning, Generalized, Sleepy, Cerberus, FPPW and Outpost channels with respect to the amount of data that is published on the blockchain in different channel closure scenarios (See Table 7.2). We show that Daric in the dishonest closure scenario outperforms Lightning with at least 1 HTLC output as well as all other schemes. In the non-collaborative closure scenario, Daric outperforms Lightning with at least 7 HTLC outputs as well as Generalized, eltoo and FPPW. Moreover, we compute the number of operations required for each channel update and show that (1) Unlike Lightning, Daric values are independent of the number of HTLC outputs m and (2) Daric is comparable with other schemes (see Table 7.2).

7.2 Notations and Background

7.2.1 Notations

In this section, we add some new notations to the ones introduced in Section 3.3. We use $\overline{[TX]}$ and $\overline{TX}$ to denote $(TX.nLT, TX.Output)$ and $(TX.nLT, TX.Output, TX.Witness)$, respectively. The absolute time-lock of i in an output condition is shown by $i^{\geq}$ and means the output cannot be spent unless the $nLockTime$ parameter in the spending transaction

is equal to or greater than i . Since a transaction may only be recorded on the blockchain if its $nLockTime$ is in the past, $i^{\geq}$ in an output condition ensures the output cannot be spent unless i is expired (i.e. i is in the past).

7.2.2 Background

7.2.2.1 Floating Transactions

Each signature in a Bitcoin transaction contains a flag, called SIGHASH, which specifies which part of the transaction has been signed. Typically, signatures are of type SIGHASH_ALL, meaning the signature authorises all inputs (i.e. references to previous outputs) and outputs. The SIGHASH of type ANYPREVOUT indicates that the signature does not authorise the inputs. This allows the signer to refer to any arbitrary UTXO whose condition is met by the transaction witness data. Such a transaction is called a *floating* transaction. The dotted arrow to TX'' in Fig. 7.1 shows that TX'' is a floating transaction whose signature matches the public key pk_C . This transaction is denoted by $\overline{TX''}$ to emphasise that since it is a floating transaction, its input is unspecified and can be any output with matching condition. The signature with the ANYPREVOUT flag is denoted by $\hat{\sigma}$.

7.2.2.2 eltoo [1]

An eltoo channel is created like a Lightning channel, but each state is represented by two transactions: (1) the *update* transaction and (2) the *settlement* transaction, where both parties have the same version of these two transactions. Each update transaction is a floating transaction that transfers all the channel funds to a new joint address. The update transaction's output can be spent by its corresponding settlement transaction, which splits the channel funds among parties. If A submits an old update transaction, she has to wait for a relative time-lock of T rounds before she can publish the corresponding settlement transaction. It gives some time to B to publish the latest update transaction (which is a floating transaction) and override the already published update transaction.

7.3 Daric Overview

To provide a high level overview of our solution, we start by reviewing the limitations of the Lightning channel and then gradually present our work.




Figure 7.1: A sample transaction flow.

7.3.1 Revocation Per State

Parties' and their watchtower's storage in a Lightning channel increases over time as they should store some revocation-related data for each revoked state. Our main idea to reduce their storage is transforming the revocation transactions into floating transactions. Thereby, participants only need to store the latest revocation transaction with the largest version number and use it upon fraud. However, for a Lightning channel, (1) the monetary value of each revocation transaction typically differs from one state to another, and (2) each commit transaction might have multiple HTLC outputs and hence the number of revocation transactions might also differ from one state to another. So, since revocation transactions of different states differ in value and number, it is infeasible to replace them all with the latest revocation transactions.

Therefore, our first modification is following the *punish-then-split* mechanism, introduced in [14]. According to this mechanism, the commit transaction sends the channel funds to a new joint output, which is controlled by both parties. The output of this commit transaction can be spent by its corresponding split transaction after t rounds where outputs of the split transaction split the channel funds between A and B . If A publishes a revoked commit transaction, B must spend its output within t rounds with the corresponding revocation transaction. This revocation transaction gives all the channel funds to B . Fig. 7.2 depicts the transaction flows for this channel where each party stores a single revocation transaction with fixed monetary value (i.e. $a+b$ coins) per state. In this figure, TX_{FU} denotes the funding transaction and $\text{TX}_{\text{CM},i}^A$, $\text{TX}_{\text{SP},i}^A$ and $\text{TX}_{\text{RV},i}^A$ (or respectively $\text{TX}_{\text{CM},i}^B$, $\text{TX}_{\text{SP},i}^B$ and $\text{TX}_{\text{RV},i}^B$) denote the commit, split and revocation transactions held by A (or respectively held by B) for state i .

7.3.2 Revocation Per Channel

In the scheme, depicted in Fig. 7.2, channel parties need to store a revocation transaction for each revoked state. Therefore, storage requirements of channel parties (or their




Figure 7.2: Transaction flows for a Lightning channel with punish-then-split mechanism.

watchtower) increase with each channel update. To solve this issue, we transform revocation transactions into floating transactions, i.e. the signatures in a revocation transaction, held by A , are of type ANYPREVOUT and meet the output condition of all commit transactions, held by B , and vice versa. It allows parties to only store the last revocation transaction.

7.3.3 Avoiding State Duplication

Since each state in the introduced scheme contains two split transactions (one for each party), the scheme suffers from state duplication. To avoid this, we transform split transactions into floating transactions. Then, each state contains one split transaction (held by both parties), which spends any of two commit transactions of that state.

7.3.4 State Ordering

Since split and revocation transactions are floating, it must be guaranteed that the latest commit transaction cannot be spent using any split or revocation transaction from previous states. Otherwise, the honest party, who has published the latest commit transaction, might lose some funds in the channel. To achieve this requirement, we repurpose [1] the $nLockTime$ parameter of split and revocation transactions to store the *state number*: the number of times the channel has been updated to date. Furthermore, we add the state number to the output condition of each commit transaction as an absolute time-lock. Then, since the absolute time-lock in the output condition of the last commit transaction would be larger than the $nLockTime$ parameter in any split or revocation transaction from previous states, the mentioned requirement is met.

7.3.5 Putting Pieces Together

The transaction flow for state i of Daric is depicted in Fig. 7.3. Let the channel be in state n . To close the channel, each party (e.g. A) can publish the latest commit transaction (e.g. $\text{TX}_{\text{CM},n}^A$), wait for T rounds and finally publish the latest split transaction $\text{TX}_{\text{SP},n}$. There is no revocation transaction for the latest state. If party B publishes a revoked commit transaction (i.e. $\text{TX}_{\text{CM},i}^B$ with $i < n$), then party A instantly publishes the latest revocation transaction $\text{TX}_{\text{RV},n-1}^A$ to take all the channel funds.




Figure 7.3: Transaction flows for state i of a Daric channel.

7.4 Daric Protocol Description

This section presents our protocol using the transaction flows depicted in Fig. 7.3. The lifetime of a Daric channel can be divided into 4 phases including create, update, close, and punish. We introduce these phases through sections 7.4.1 to 7.4.4. Section 7.11 provides the formal description of the protocol.

7.4.1 Create

To create the channel, A and B sign and publish the funding transaction TX_{FU} on the blockchain. By publishing this transaction, A and B fund the channel with a and b coins, respectively, but since the output of the funding transaction can only be spent if both parties agree, one party might become unresponsive to raise a hostage situation. To avoid this, before signing the funding transaction, parties commit to the initial channel state, i.e. state 0, by exchanging signatures for the corresponding commit and split

transactions. Let us explain the different steps of the channel creation phase in more detail.

- **Step 1:** At the first step, A and B send their funding sources (i.e. $txid_A$ and $txid_B$) to each other. This enables them to create the body of the funding transaction $[TX_{FU}]$.
- **Step 2:** Having the transaction identifier of TX_{FU} , parties create the body of the commit transactions, i.e. $[TX_{CM,0}^A]$ and $[TX_{CM,0}^B]$.
- **Steps 3:** Parties exchange the required signatures (with SIGHASH of type ANYPREVOUT) to create the floating transaction $\overline{TX_{SP,0}}$. This floating transaction could take output of $TX_{CM,0}^A$ or $TX_{CM,0}^B$ as its input.
- **Step 4:** Parties exchange the required signatures to create the commit transactions $TX_{CM,0}^A$ and $TX_{CM,0}^B$.
- **Step 5:** Parties exchange the required signatures to create the funding transactions TX_{FU} .
- **Step 6:** Parties publish the funding transaction on the blockchain.

The absolute time-lock in the output script of commit transactions and correspondingly the $nLockTime$ parameter in the split transaction must be in the past. Otherwise, parties have to wait to publish such transactions. As explained in Section 7.3.4, the time-lock is set to the state number and hence its value increases with each channel update. Absolute time-locks lower than 500,000,000 specify the block number after which the transaction can be included in a block. According to the value of the current block height, if we set the initial time-lock to the first state number, i.e. 0, the channel can be updated around 700,000 times. However, absolute time-locks equal to or larger than 500,000,000 specify the UNIX timestamp after which the transaction will be valid.

According to the value of the current timestamp, if we set the initial time-lock (and correspondingly $nLockTime$ parameter) to 500,000,000, the channel can be updated around 1 billion times [1]. Moreover, the current timestamp increases one unit per second, meaning if the average rate of the channel update is up to once per second, the channel can be updated an infinite number of times.

The above-mentioned transactions are further explained below.

- **Funding transaction:** Using this transaction, channel parties A and B open a Daric channel by funding a and b , coins into the channel, respectively. The funding

transaction is as follows¹:

$$\begin{aligned}
 \text{TX}_{\text{FU}}.\text{nLT} &:= 0, \\
 \text{TX}_{\text{FU}}.\text{Input} &:= (\text{txid}_A \| x, \text{txid}_B \| y), \\
 \text{TX}_{\text{FU}}.\text{Output} &:= \{(a + b, pk_A \wedge pk_B)\}, \\
 \text{TX}_{\text{FU}}.\text{Witness} &:= ((1, \sigma_{\text{TX}_{\text{FU}}}^{A,1}), (1, \sigma_{\text{TX}_{\text{FU}}}^{B,2})).
 \end{aligned}$$

- **Commit transaction:** There exist two versions of commit transaction per state, each held by one of the parties, but only the first ones ($\text{TX}_{\text{CM},i}^A$ and $\text{TX}_{\text{CM},i}^B$ with $i = 0$) are created at this phase. Commit transactions send the channel funds to a joint account. Commit transactions for state i are as follows:

$$\begin{aligned}
 \text{TX}_{\text{CM},i}^A.\text{nLT} &:= 0, \\
 \text{TX}_{\text{CM},i}^A.\text{Input} &:= \text{TX}_{\text{FU}}.\text{txid} \| 1, \\
 \text{TX}_{\text{CM},i}^A.\text{Output} &:= \{(a + b, \varphi_1 \vee \varphi_2)\}, \\
 \text{TX}_{\text{CM},i}^A.\text{Witness} &:= \{(1, \{\sigma_{\text{TX}_{\text{CM},i}^A}^A, \sigma_{\text{TX}_{\text{CM},i}^A}^B\})\}
 \end{aligned}$$

and

$$\begin{aligned}
 \text{TX}_{\text{CM},i}^B.\text{nLT} &:= 0, \\
 \text{TX}_{\text{CM},i}^B.\text{Input} &:= \text{TX}_{\text{FU}}.\text{txid} \| 1, \\
 \text{TX}_{\text{CM},i}^B.\text{Output} &:= \{(a + b, \varphi_1 \vee \varphi'_2)\}, \\
 \text{TX}_{\text{CM},i}^B.\text{Witness} &:= \{(1, \{\sigma_{\text{TX}_{\text{CM},i}^B}^A, \sigma_{\text{TX}_{\text{CM},i}^B}^B\})\}
 \end{aligned}$$

with $\varphi_1 := (pk_{\text{SP}}^A \wedge pk_{\text{SP}}^B \wedge t^+ \wedge i^{\geq})$, $\varphi_2 = (pk_{\text{RV}}^A \wedge pk_{\text{RV}}^B \wedge i^{\geq})$ and $\varphi'_2 = (pk_{\text{RV}}^A \wedge pk_{\text{RV}}^B \wedge i^{\geq})$. Meeting each sub-condition requires both parties' authorisation. The parameters $i^{\geq}$ and t^+ show absolute time-lock of i and relative time-lock of t rounds, respectively. The parameter t can be set to any value larger than the blockchain delay Δ .

- **Split transaction** There is one split transaction per state which is held by both channel parties, but only the first one ($\text{TX}_{\text{SP},i}$ with $i = 0$) is created at this phase. The split transaction for state i determines the balance of each channel party in

¹We assume that funding sources of TX_{FU} are two typical UTXOs owned by A and B .

the i^{th} channel state. For this transaction, we have:

$$\begin{aligned} \text{TX}_{\text{SP},i}.\text{nLT} &:= S_0 + i, \\ \text{TX}_{\text{SP},i}.\text{Input} &:= \text{TX}_{\text{CM},i}^P.\text{txid}\|1, \text{ with } P \in \{A, B\}. \\ \text{TX}_{\text{SP},i}.\text{Output} &:= (\theta_1, \theta_2, \dots), \\ \text{TX}_{\text{SP},i}.\text{Witness} &:= \{(1, \{\hat{\sigma}_{\text{TX}_{\text{SP},i}}^A, \hat{\sigma}_{\text{TX}_{\text{SP},i}}^B\})\} \end{aligned}$$

The transaction $\text{TX}_{\text{SP},i}$ is a floating transaction whose witness satisfies the first sub-condition of the output of $\text{TX}_{\text{CM},i}^A$ or $\text{TX}_{\text{CM},i}^B$ ($pk_{\text{SP}}^A \wedge pk_{\text{SP}}^B \wedge t^+ \wedge i \geq$). Normally, given that parties are honest and the channel is in state n , one of two commit transactions of the latest state, i.e. either $\text{TX}_{\text{CM},n}^A$ or $\text{TX}_{\text{CM},n}^B$, is published on-chain. Then, the first sub-condition of its output with the ($pk_{\text{SP}}^A \wedge pk_{\text{SP}}^B \wedge t^+ \wedge i \geq$) is satisfied by the corresponding split transaction $\text{TX}_{\text{SP},n}$ after t rounds.

To reduce the required communication between channel parties for each channel update, pk_{SP}^A and pk_{SP}^B do not change from one state to the next. However, since split transactions are floating, it must be guaranteed that the split transaction of state i cannot take the output of one of the commit transactions of the next states as its input because otherwise the output of the latest commit transaction, let's say $\text{TX}_{\text{CM},n}^A$, could be spent using an old split transaction $\text{TX}_{\text{SP},i}$ with $i < n$, which is undesirable. To meet this requirement, the $n\text{LockTime}$ parameter of $\text{TX}_{\text{SP},i}$ is set to i . Then since the first sub-condition of commit transactions of state j with $j > i$ are time-locked using $j \geq$ but $\text{TX}_{\text{SP},i}.\text{nLT} \geq j$ does not hold, $\text{TX}_{\text{SP},i}$ cannot spend output of $\text{TX}_{\text{CM},j}^A$ or $\text{TX}_{\text{CM},j}^B$.

7.4.2 Update

Let the channel be in state $i \geq 0$ and channel parties decide to update it to state $i + 1$. The update process is performed in two sub-phases. The first sub-phase is similar to steps 2 to 4 of the channel creation phase where channel parties create two new commit transactions $\text{TX}_{\text{CM},i+1}^A$ and $\text{TX}_{\text{CM},i+1}^B$ as well as a new split transaction $\overline{\text{TX}_{\text{SP},i+1}}$ for the new state. In the second sub-phase, channel parties revoke the state i by signing two revocation transactions $\overline{\text{TX}_{\text{RV},i}^A}$ and $\overline{\text{TX}_{\text{RV},i}^B}$. The revocation transaction $\overline{\text{TX}_{\text{RV},i}^A}$ (or respectively $\overline{\text{TX}_{\text{RV},i}^B}$) contains no input yet and can spend output of any commit transaction $\overline{\text{TX}_{\text{CM},j}^B}$ (or respectively $\overline{\text{TX}_{\text{CM},j}^A}$) with $j \leq i$. With each channel update, the state number and hence the time-lock value in the output condition of each commit transaction and $n\text{LockTime}$ in split and revocation transactions increase by one unit. Let us explain the different steps of the channel update phase in more detail.

- **Step 1:** Parties create the body of the commit transactions, i.e. $[TX_{CM,i+1}^A]$ and $[TX_{CM,i+1}^B]$.
- **Steps 2:** Parties exchange the required signatures (with SIGHASH of type ANYPREVOUT) to create the floating transaction $\overline{TX_{SP,i+1}}$. This floating transaction takes output of $TX_{CM,i+1}^A$ or $TX_{CM,i+1}^B$ as its input.
- **Step 3:** Parties exchange the required signatures to create the commit transactions $TX_{CM,i+1}^A$ and $TX_{CM,i+1}^B$.
- **Step 4:** Parties exchange the required signatures (with SIGHASH of type ANYPREVOUT) to create the floating transactions $\overline{TX_{RV,i}^A}$ and $\overline{TX_{RV,i}^B}$.

One of the parties might receive the signature on the split or commit transactions in steps 2 or 3 (or respectively receive the signature on the revocation transaction in step 4) but avoid signing the corresponding transaction for the other party. In such situations, the honest party non-collaboratively closes the channel with the latest valid channel state, i.e. state i (or respectively state $i + 1$). More technical details can be found in Section 7.11.

Revocation transaction will be introduced further below:

- **Revocation transaction** Once the channel is updated from state i to $i + 1$, two versions of revocation transaction are created for state i , one version for each channel party. The revocation transaction held by party A and party B for state i are denoted by $TX_{RV,i}^A$ and $TX_{RV,i}^B$, respectively, where

$$\begin{aligned}
 TX_{RV,i}^A.nLT &:= i, \\
 TX_{RV,i}^A.Input &:= TX_{CM,j,B}.txid\|1, \text{ with } j \leq i, \\
 TX_{RV,i}^A.Output &:= \{(a + b, pk_A)\}, \\
 TX_{RV,i}^A.Witness &:= \{(2, \{\hat{\sigma}_{TX_{RV,i}^A}^A, \hat{\sigma}_{TX_{RV,i}^A}^B\})\}
 \end{aligned}$$

and

$$\begin{aligned}
 TX_{RV,i}^B.nLT &:= i, \\
 TX_{RV,i}^B.Input &:= TX_{CM,j,A}.txid\|1, \text{ with } j \leq i, \\
 TX_{RV,i}^B.Output &:= \{(a + b, pk_B)\}, \\
 TX_{RV,i}^B.Witness &:= \{(2, \{\hat{\sigma}_{TX_{RV,i}^B}^A, \hat{\sigma}_{TX_{RV,i}^B}^B\})\}
 \end{aligned}$$

The transaction $TX_{RV,i}^A$ (or $TX_{RV,i}^B$) is a floating transaction whose witness satisfies the second sub-condition of the output of $TX_{CM,j,B}$ (or $TX_{CM,j,A}$) with $j \leq i$. Given that

the channel is in state n , if a dishonest channel party, let's say party B , publishes the old commit transaction $\text{TX}_{\text{CM},j,B}$ with $j < n$, he must wait for T rounds before being able to publish its corresponding split transaction $\text{TX}_{\text{SP},j}$. However, A can instantly publish the latest revocation transaction $\text{TX}_{\text{RV},n-1,A}$ and claim its output.

Since revocation transactions are floating and meet the second sub-condition of commit transactions, the public keys that are used in such sub-conditions must not change from one state to the next. However, it must be guaranteed that the revocation transaction of state i cannot take the output of commit transactions of the next states as input. Thus, similar to what we did for split transactions, the $n\text{LockTime}$ parameter of revocation transactions of state i is set to $S_0 + i$. Then since the second sub-condition of each commit transaction of state i is also time-locked using CLTV_{S_0+i} , the desired requirement is met.

A dishonest channel party, e.g. party A , must not be able to publish both a revoked commit transaction $\text{TX}_{\text{CM},i}^A$ with $i < n$ as well as a revocation transaction $\text{TX}_{\text{RV},j,A}$ with $j \geq i$ on the ledger. Otherwise, he can take all the channel funds. To prevent A from doing so, the revocation public keys that are used in commit transactions held by A (pk_{RV}^A and pk_{RV}^B) are different from those in commit transactions held by B (pk_{RV}^A and pk_{RV}^B). Therefore, the output of a revoked commit transaction held by A can only be spent by a revocation transaction held by B and vice versa.

7.4.3 Close

Assume while the channel between A and B is in state n , they decide to collaboratively close it. To do so, A and B exchange signatures for a new transaction, called modified split transaction TX_{SP} , and publish it on the blockchain. This transaction takes the funding transaction's output as its input and splits the channel funds among channel parties. The transaction TX_{SP} is as following:

$$\begin{aligned} \text{TX}_{\text{SP}}.\text{Input} &:= \text{TX}_{\text{FU}}.\text{txid}\|1, \\ \text{TX}_{\text{SP}}.\text{Output} &:= \text{TX}_{\text{SP},n}.\text{Output}, \\ \text{TX}_{\text{SP}}.\text{Witness} &:= \{(1, \{\sigma_{\text{TX}_{\text{SP}}}^A, \sigma_{\text{TX}_{\text{SP}}}^B\})\}. \end{aligned}$$

If one of the channel parties, e.g. party B , becomes unresponsive, its counterparty A can still non-collaboratively close the channel by publishing $\text{TX}_{\text{CM},n}^A$, adding the output of $\text{TX}_{\text{CM},n}^A$ as an input to $\overline{\text{TX}_{\text{SP},n}}$ to transform it into $\text{TX}_{\text{SP},n}$, and finally publishing $\text{TX}_{\text{SP},n}$ after t rounds.

7.4.4 Punish

Let the channel be in state n . If a dishonest channel party, let's say A , publishes an old commit transaction $\text{TX}_{\text{CM},i}^A$ with $i < n$ on the ledger, party B adds the output of $\text{TX}_{\text{CM},i}^A$ as an input to $\text{TX}_{\text{RV},n-1}^B$ in order to transform it into $\text{TX}_{\text{RV},n-1}^B$ and instantly publishes $\text{TX}_{\text{RV},n-1}^B$ on the blockchain.

7.5 Security Analysis Overview

In this section, we first provide some payment channel notations as well as our security model, which follows previous works on *layer-2* solutions [14, 68, 69, 29]. Then, we present desired properties of a payment channel and an ideal functionality $\mathcal{F}$ that attains those properties. Finally, we show that the Daric protocol is a realisation of the ideal functionality $\mathcal{F}$ and hence achieves its desired properties.

7.5.1 Notation and Security Model

We use an extended version of the *universal composability* framework [55] to formally model the security of our construction. This extended version [70], called the Global Universal Composability framework (GUC), supports a global setup. To simplify our model, we assume that the communication network is synchronous, meaning that the protocol is executed through multiple rounds and parties in the protocol are connected to each other via an authenticated communication channel which guarantees 1-round delivery. Transactions are recorded by a global ledger $\mathcal{L}(\Delta, \Sigma)$, where Σ is a signature scheme used by the blockchain and Δ is an upper bound on the blockchain delay: the number of rounds it takes a transaction to be accepted by the ledger. Section 7.9 provides more details on our security model.

We abbreviate a daric payment channel γ as an attribute tuple $\gamma := (\text{id}, \text{users}, \text{cash}, \text{st}, \text{sn}, \text{flag}, \text{st}')$, where $\gamma.\text{id} \in \{0, 1\}^*$ defines the channel identifier, $\gamma.\text{users}$ represents the identities of the channel users, $\gamma.\text{cash} \in \mathbb{R}^{\geq 0}$ is the channel capacity, $\gamma.\text{st} := (\theta_1, \dots, \theta_l)$ is a list of l outputs defining the channel state after the last complete channel update and $\gamma.\text{sn}$ is the state number. The flag $\gamma.\text{flag} \in \{1, 2\}$ and the state $\gamma.\text{st}'$ will be explained below.

The initial value of $\gamma.\text{flag}$ and $\gamma.\text{st}'$ are 1 and $\perp$, respectively. Assume that the channel has been updated $n \geq 0$ times and the channel state after the n^{th} update is st and hence we have $\gamma.\text{st} = st$. Now, assume that parties start the update process to update the state of the channel from state st to st' . From a particular point in the channel update process

onward, at least one of the parties has sufficient data to enforce the new state st' on the blockchain when parties have not completely revoked the state $\gamma.st$ yet. The flag $\gamma.flag$ is set to 2 to identify such occasions and $\gamma.st'$ is set to st' to maintain the new state. Thus, when $\gamma.flag = 2$, the channel might be finalised with either $\gamma.st$ or $\gamma.st'$. At the end of the channel update process, once the state st was revoked by both parties, $\gamma.st$ and $\gamma.st'$ are set to st' and $\perp$, respectively, and $\gamma.flag$ is set to 1.

7.5.2 Ideal Functionality Properties

This section closely follows [14] to introduce desired security and efficiency properties of a payment channel as follows:

- **Consensus on creation:** A channel γ is created only if both channel parties in the set $\gamma.users$ agree to create it.
- **Consensus on update:** A channel γ is updated only if both channel parties in the set $\gamma.users$ agree to update it. Also, parties reach an agreement on update acceptance or rejection within a bounded number of rounds (the bound might depend on the ledger delay Δ).
- **Bounded closure with punish:** An honest user $P \in \gamma.users$ has the assurance that within a bounded number of rounds (the bound might depend on the ledger delay Δ), she can finalise the channel state on the ledger either by enforcing a state that gives her $\gamma.cash$ coins, or by enforcing $\gamma.st$ if $\gamma.flag = 1$ or by enforcing either $\gamma.st$ $\gamma.st'$ otherwise.
- **Optimistic update:** If both parties in $\gamma.users$ are honest, the channel update completes with no ledger interaction.

Section 7.10 introduces an ideal functionality $\mathcal{F}$ that achieves these properties. Theorem 7.1 shows Daric protocol, denoted by π , is a realisation of $\mathcal{F}$ and hence achieves its desired properties. It follows from 14 Lemmas. Section 7.12 presents the full security proof.

Theorem 7.1. *Let Σ be an EUF – CMA secure signature scheme. Then, for any ledger delay $\Delta \in \mathbb{N}$, the protocol π UC-realises the ideal functionality $\mathcal{F}(T)$ with any $t > \Delta$.*

Proof. Theorem follows directly from Lemma 7.1, Lemma 7.4, Lemma 7.5 and Lemma 7.8, presented in Section 7.12. □

To enhance structure and prevent an overly lengthy section, we partition the necessary contents for our security proof into two separate chapters. We formally define π in Section 7.11 and then provide a simulator $\mathcal{S}$ in Section 7.12 where $\mathcal{S}$ has interaction with the ideal functionality $\mathcal{F}$ and $\mathcal{L}$. The simulator simulates the content and timing of all messages of the honest party to the adversary and also translates any message from the adversary into a message to the ideal functionality, such that an indistinguishable execution of the protocol in the ideal world is emulated. Also, in Section 7.12, we prove for any action that causes the ideal functionality to output Error with non-negligible probability, the simulator constructs a reduction against the existential unforgeability of the underlying signature scheme Σ with non-negligible success probability, which contradicts with our assumption regarding the security of Σ . This proves our protocol would be as secure as the ideal functionality $\mathcal{F}$ and provides its desirable properties.

7.6 Daric Versus Eltoo

In section 7.6.1, we present an attack to eltoo whose main purpose is to postpone the channel closure. We show this attack is practically profitable when applied to eltoo but it cannot be applied to Daric. In section 7.6.2, we analyse Daric and eltoo to compare their robustness against profit-driven attackers. We use the statistical data derived from the Lightning Network to enable such an analysis.

7.6.1 HTLC Security

This section presents an attack against HTLC security in eltoo (previously informally discussed in [67]) and analyses the attacker's revenue. Let the adversary represent two nodes on the PCN: node M_1 and node M_2 . Assume that the adversary has established N channels from M_1 to victim nodes $V_1, \dots, V_N$ and N channels from victim nodes to M_2 . The channel between M_1 and V_i is denoted with γ_i . The adversary performs N simultaneous HTLC payments from M_1 to M_2 through $V_1, \dots, V_N$. Let the payment value for all HTLCs be A coins and the time-lock for all these payments for M_1 's channels be T . Assume that M_2 accepts the payments and provides the required secrets for all HTLC payments and hence M_2 is paid $N \cdot A$ coins in total. Then, victims provide the secrets to the node M_1 . However, M_1 does not update her channels with victims. Therefore, victims attempt to claim all HTLCs on-chain. To prevent victims from closing their channels in time, M_1 takes the following steps:

1. Submit a valid *Delay* transaction TX_{De} with $N + 1$ inputs and $N + 1$ outputs where the i^{th} input-output pair corresponds with an outdated state of the channel γ_i and

the last input-output pair adds further funds to be used as the transaction fee, which is set to any value larger than A .

2. If TX_{De} is published and the time-lock t is still unexpired, go to step 1.
3. Once the time-lock t is expired, submit the latest channel state for all channels and claim their HTLC outputs.

In the above attack, to replace the already submitted transaction TX_{De} with the latest state of the channel γ_i , V_i has to set a transaction fee that is larger than the total absolute transaction fee of TX_{De} [71]. But since the transaction fee for TX_{De} is larger than A , V_i will be unwilling to pay such a transaction fee.

Once the HTLC time-lock is expired and the latest channel state is added to the ledger, there will be a race between M_1 and each victim to claim the HTLC output. The adversary will have a better chance to win the race if she has a better network connection with a higher number of nodes.

Now we perform a cost-benefit analysis to determine if the attack is profitable to the attacker. For a fixed value of A , with setting N to the largest possible value, the adversary (1) reduces the fee per channel for each delay transaction and (2) reduces the pace at which outdated states are added to the blockchain. A Bitcoin transaction can contain up to 100,000 VBytes (where each VByte equals four weight units) and each input-output pair contains 222 bytes of witness data and 84 bytes of non-witness data (See Appendix H.4 in the full version of the Daric paper [63] for more details). Therefore, TX_{De} can cover up to around $\frac{100,000}{0.25 \times 222 + 84} \approx 715$ eltoo channels. The minimum possible fee rate is 1 Satoshi per VByte. Thus, if A is set to 100,000 Satoshi, the total fee for each delay transaction would be 100,000 Satoshi.

At the time of writing this chapter (in January 2023), the average transaction fee is quite low and hence transactions with the minimum fee rate are added to the blockchain in 30 minutes. It means if HTLC time-locks are set to 3 days, 144 delay transactions are published before time-locks getting expired. In other words, the adversary pays $144A$ as a transaction fee to earn up to $715A$. In more congested times, it might take several hours for a transaction with a minimum fee rate to be added to the blockchain. Thus, the attack could be even more profitable to the attacker. This attack is inapplicable to Daric because once the attacker publishes an old commit transaction, the only valid transactions are the revocation transactions held by her counterparty.

7.6.2 Punishment Mechanism

Prior to providing a formal analysis, we provide intuitions as follows. The only cost for a dishonest party in eltoo is the fee for publishing the old state, which could be (1) less than

1 USD for Bitcoin and (2) independent of the channel capacity. However, given that the balance of each party in a Daric channel cannot be less than 1% of the channel capacity (which is currently deployed in the Lightning Network), the minimum amount that a dishonest party might lose would have the following properties: (1) It is proportional to the channel capacity, (2) Its value (around 20 USD on average in the Lightning Network in April 2022) is typically significantly larger than the transaction fee and (3) It is easily raised by increasing the minimum possible balance of each channel party from 1% of the channel capacity to a higher proportion. Therefore, Daric's deterrent effect against profit-driven attackers is higher and more flexible than that of eltoo.

Now, we present a more formal comparison between eltoo and Daric. We assume the channel party either stays online or employs a watchtower that is fair w.r.t the hiring party [54] (i.e. the watchtower guarantees its client's funds in the channel). For the former case, let p denote the probability that the honest channel party successfully reacts upon fraud, i.e. $1 - p$ is the probability that the honest party, due to crash failures or DoS attacks, fails to react. We show that (1) to discourage attacks by profit-driven parties, p for eltoo must be more significant than that of Daric, and (2) unlike Daric, an increase in the channel capacity in eltoo channels raises the minimum value of p that is required to prevent fraud. However, achieving large values of p (e.g. 0.9999) could be difficult for ordinary users. This indicates that eltoo needs a way to punish profit-driven attackers.

To monitor a channel, the watchtower's collateral equals the channel capacity [54, 23]. Let C denote the total capacity of the Bitcoin payment channel network and C_W denote the total capital that fair watchtowers have spent to watch their clients' channels. Then, the probability that a randomly selected payment channel is monitored by a fair watchtower is roughly computed as $\frac{C_W}{C}$.

Assume that a dishonest party $\mathcal{A}$ creates an eltoo channel with a channel capacity of $C_{\mathcal{A}}$ coins, where the initial balance of $\mathcal{A}$ and her counterparty are $C_{\mathcal{A}}$ and 0, respectively. For now, we assume that parties know if their counterparties are using a fair watchtower. We will relax this assumption later. If the channel is being monitored by a fair watchtower, $\mathcal{A}$ continues using the channel in an honest way. Otherwise, she sends all her balance to her counterparty in exchange for some products or services and then submits the initial channel state to the blockchain. In such a case, with a probability of $1 - p$ and p , $\mathcal{A}$'s revenue and her loss would be $C_{\mathcal{A}} - f$ and f , respectively, where f denotes the transaction fee. Thus, $\mathcal{A}$ is discouraged to attack iff:

$$(C_{\mathcal{A}} - f)(1 - p) - f \cdot p < 0 \Leftrightarrow p > 1 - \frac{f}{C_{\mathcal{A}}}.$$

For a Daric channel, $\mathcal{A}$ is discouraged to attack iff:

$$0.99 \cdot C_{\mathcal{A}} \cdot (1 - p) - 0.01 \cdot C_{\mathcal{A}} \cdot p < 0 \Leftrightarrow p > 0.99.$$

The threshold value for eltoo is typically more significant than that of Daric. At the time of writing this chapter, the average values of f for a transaction and $C_{\mathcal{A}}$ for a Lightning channel are around 0.000042 BTC and 0.069 BTC, respectively, leading to $1 - \frac{f}{C_{\mathcal{A}}} \approx 0.999$. But the adversary can practically set f to the lowest possible value (i.e. 1 Satoshi per VByte) leading to $f \approx 0.0000021^2$ BTC and $1 - \frac{f}{C_{\mathcal{A}}} \approx 0.9999$ for eltoo. Therefore, (1) to discourage attacks, the honest party would require to meet a higher p in eltoo than in Daric, (2) the threshold for eltoo depends on the channel capacity, and (3) the threshold for Daric can simply decrease from 0.99 to lower values.

In the above analysis, we assumed that $\mathcal{A}$ knows whether her counterparty is hiring any fair watchtower. Considering the opposite case, the probability that the channel is not being monitored by any fair watchtower and the honest party fails to react upon fraud would be $p_0 := (1 - \frac{C_W}{C})(1 - p)$. Thus, with a probability of p_0 and $1 - p_0$, $\mathcal{A}$'s revenue and her loss in an eltoo channel would be $C_{\mathcal{A}} - f$ and f , respectively. Thus, $\mathcal{A}$ is discouraged to attack iff:

$$(C_{\mathcal{A}} - f) \cdot p_0 - f \cdot (1 - p_0) < 0 \Leftrightarrow p > 1 - \frac{\frac{f}{C_{\mathcal{A}}}}{1 - \frac{C_W}{C}}.$$

Similarly, for a Daric channel, we have:

$$0.99 \cdot C_{\mathcal{A}} \cdot p_0 - 0.01 \cdot C_{\mathcal{A}} \cdot (1 - p_0) < 0 \Leftrightarrow p > 1 - \frac{0.01}{1 - \frac{C_W}{C}}.$$

As explained earlier, the threshold value for eltoo depends on $C_{\mathcal{A}}$ and is typically more significant than that of Daric.

7.7 Performance Analysis

Table 7.2 shows the total number of weight units of transactions, published on the blockchain for different payment channels in different channel closure scenarios. Since the weight units of a transaction directly impact its fee, we use this parameter to compare different schemes. Payment channels perform similarly in the collaborative channel closure, so we do not consider this scenario in our analysis. Since the funding transaction

²Each update transaction in eltoo contains 332 byte of witness data and 125 bytes of non-witness data leading to 208 VBytes. See Appendix H in the full version of the Daric paper [63] for more details

is the same in all schemes, we do not involve it in our comparison results either. To do a consistent comparison, we assume that each transaction output is either P2WSH³ or P2WPKH⁴, each public key and signature are respectively 33 bytes and 73 bytes, shared outputs are implemented using the OP_CHECKMULTISIG opcode (rather than using multi-party signing), and each state contain m HTLC outputs with $0 \leq m \leq 966$ [72] where each party is the payer for $\frac{m}{2}$ HTLC outputs and the payee for the rest.

Once a dishonest party in a Lightning channel publishes a revoked commit transaction, $m+1$ revoked outputs are created. For simplicity, we assume that the victim claims all the revoked outputs through one transaction. Cerberus [23], Sleepy [28] and Outpost [22] have not explained ways HTLC is added to these schemes and discussing it is out of the scope of this chapter, so Table 7.2 contains their figures with $m = 0$.

As Table 7.2 shows, in the dishonest closure scenario, (1) the weight units for Lightning and eltoo increase linearly with the number of HTLC outputs m compared to Daric, Generalized and FPPW and (2) Daric (with weight unit equal to 1239) is more cost-effective than other schemes with $m \geq 1$. In the non-collaborative closure scenario with $m \neq 0$, Daric outperforms Generalized, eltoo and FPPW channels with any value of m and Lightning channel with $m > 6$.

Table 7.2 also compares the number of operations performed by each party for a channel update. To count the operations, we additionally assume that (1) channel parties delegate the monitoring task to a watchtower and (2) they do not compute a signature unless it is supposed to be sent to their counterparty or their watchtower. Appendix H in the full version of the Daric paper [63] provides complete details regarding the way figures of Table 7.2 have been computed (The full version paper is available at <https://eprint.iacr.org/2022/1295.pdf>).

7.8 Daric Transactions Scripts

Funding transaction has one output with the following script where Com_pubkeyA and Com_pubkeyB are public keys of A and B, respectively:

2 <Com_pubkeyA> <Com_pubkeyB> 2 OP_CHECKMULTISIG

Commit transactions have one input that takes the output of the funding transaction with the witness script of $\emptyset$ <Com_pubkeyA_sig> <Com_pubkeyB_sig>. The commit transaction $TX_{CM,i}^A$ has one output with the following script:

<absolute time $S_0 + i$ > OP_CHECKLOCKTIMEVERIFY OP_DROP

³Pay-to-Witness-Script-Hash: Used to lock bitcoin to a SegWit script hash.

⁴Pay-to-Witness-Public-Key-Hash: Used to lock bitcoin to a SegWit public key hash.

Table 7.2: On-chain cost of different closure scenarios and the number of operations performed by each party for a channel update for different payment channels with m HTLC outputs ($0 \leq m \leq 966$).

Scheme	dishonest closure		non-coll. closure		num. of operations		
	#Tx	weight units	#Tx	weight units	Sign	Verify	Exp.
Lightning [15]	≥ 2	$\geq 1209+582.5m$	$1+m$	$724+793m$	$2+2m$	$1+\frac{m}{2}$	2
Generalized [14]	2	1342	$2+m$	$1432+696m$	3	2	1
FPPW [54]	2	2045	$2+m$	$1562+696m$	6	10	1
Cerberus [23]	2	1798	1	772	3	6	0
Outpost [22]	3	2632	3	3018	4	4	0
Sleepy [28]	3	2172	3	2558	5	5	0
eltoo [1]	3	$2268+696m$	$2+m$	$1588+696m$	2	2	1
Daric (this work)	2	1239	$2+m$	$1363+696m$	4	3	0

OP_IF

Revocation

2 $\langle \text{Rev_pubkeyA} \rangle \langle \text{Rev_pubkeyB} \rangle$ 2 OP_CHECKMULTISIG

OP_ELSE

Split

$\langle \text{delay } t \rangle$ OP_CHECKSEQUENCEVERIFY OP_DROP

2 $\langle \text{Spl_pubkeyA} \rangle \langle \text{Spl_pubkeyB} \rangle$ 2 OP_CHECKMULTISIG

OP_ENDIF

where $\langle \text{Rev_pubkeyA} \rangle$ and $\langle \text{Spl_pubkeyA} \rangle$ are public keys of A and $\langle \text{Rev_pubkeyB} \rangle$ and $\langle \text{Spl_pubkeyB} \rangle$ are public keys of B . The script of the commit transaction $\text{TX}_{\text{CM},i}^B$ is similar to that of $\text{TX}_{\text{CM},i}^A$, but its revocation keys are $\langle \text{Rev}'_pubkeyA \rangle$ and $\langle \text{Rev}'_pubkeyB \rangle$.

The split transaction $\text{TX}_{\text{SP},i}$ spends the output of $\text{TX}_{\text{CM},i}^A$ or $\text{TX}_{\text{CM},i}^B$ with the witness script:

$$0 \langle \text{Spl_pubkeyA_Sig} \rangle \langle \text{Spl_pubkeyB_Sig} \rangle 0$$

The revocation transactions $\text{TX}_{\text{RV},i}^A$ and $\text{TX}_{\text{RV},i}^B$ spend the output of a revoked commit transaction with the witness scripts

$$0 \langle \text{Rev}'_pubkeyA_sig \rangle \langle \text{Rev}'_pubkeyB_sig \rangle 1$$

and

$$0 \langle \text{Rev_pubkeyA_sig} \rangle \langle \text{Rev_pubkeyB_sig} \rangle 1,$$

respectively. The transactions $\text{TX}_{\text{RV},i}^A$ and $\text{TX}_{\text{RV},i}^B$ have one output with the scripts $\langle \text{pubkeyA} \rangle \text{OP_CHECKSIG}$ and $\langle \text{pubkeyB} \rangle \text{OP_CHECKSIG}$, respectively.

7.9 UC Framework

We model the security of our protocol in the synchronous version of the global UC framework (GUC) [70] which is an extension of the standard UC framework [55]. In the synchronous version, we can have a global setup that is used to model the ledger. The model in this work closely follows that of some works on layer-2 solutions to the scalability of blockchains [14, 68, 69].

Let π be a protocol executed among parties of a set $\mathcal{P} = \{P_1, \dots, P_n\}$. Assume that there exists an adversary $\mathcal{A}$ that takes as input a security parameter $\lambda \in \mathbb{N}$ and an auxiliary input $z \in \{0, 1\}^*$. Before execution of the protocol, π , the adversary $\mathcal{A}$ can select any party $P_i \in \mathcal{P}$ to learn their internal state and fully control them. Anything outside the protocol execution is modelled by the environment $\mathcal{E}$. Each protocol party as well as the adversary take their inputs from the environment. Outputs of all parties are also observed by the environment. There are also ideal functionalities $\mathcal{F}_1, \dots, \mathcal{F}_m$ whose functions might be called by parties. Then, the protocol π is denoted by $\pi^{\mathcal{F}_1, \dots, \mathcal{F}_m}$.

To simplify our model, we assume that the communication network is synchronous, meaning that we let the protocol be executed through several rounds. The ideal functionality $\mathcal{F}_{\text{clock}}$ represents a global clock that increases by one unit once all parties are prepared to proceed to the next round. All entities know the value of the current round.

We assume that parties of the protocol are connected to each other via an authenticated communication channel which guarantees that messages are delivered to recipient parties after exactly one round. In other words, if party P sends a message m to the party Q in round τ , the message reaches Q in the beginning of round $\tau + 1$ and Q can ensure that the sender is P . The adversary $\mathcal{A}$ observes the message m and can even change the order of messages that are sent in the same round. However, the adversary cannot drop, delay or change any transmitted message or insert new messages. The ideal functionality $\mathcal{F}_{\text{GDC}}$ models such a communication channel between the channel parties. Other communications in which some other entities, e.g. $\mathcal{A}$ or $\mathcal{E}$, are involved take zero rounds to complete. Moreover, to simplify the model, we assume that any required computation is also performed within zero rounds.

In this thesis, we focus on UTXO-based cryptocurrencies such as Bitcoin. The *global* ideal functionality $\mathcal{L}$ models such cryptocurrencies where it is parameterised by two parameters: (1) a digital signature scheme Σ , and (2) a delay parameter Δ which is an upper bound on the number of rounds it takes for a valid transaction that has been

posted to the blockchain network to be published on the blockchain. To simplify the model, we assume that the set of parties $\mathcal{P}$ is fixed and transactions are published one by one rather than being published on the blockchain in blocks. Badertscher et. al [73] provided a more accurate model of the Bitcoin blockchain.

The environment $\mathcal{E}$ initiates the ledger functionality $\mathcal{L}$ by (1) instructing $\mathcal{L}$ to generate the public parameters of the signature scheme Σ , (2) instructing each party $P \in \mathcal{P}$ to create a key pair (pk_P, sk_P) and submit its public key to $\mathcal{L}$, and (3) creating an initial state TX that contains all the accepted transactions. The set TX is accessible to everyone including the protocol parties, the environment and the adversary. Once a party $P \in \mathcal{P}$ posts a transaction tx to the blockchain, $\mathcal{L}$ waits for $\tau \leq \Delta$ rounds where τ is selected by the adversary. Then, if the validity of tx is successfully verified, it is added to the set TX.

Ideal Functionality $\mathcal{L}(\Delta, \Sigma)$

The set $\mathcal{P}$ defines the set of all parties who can send messages to the functionality. The functionality maintains the set PKI for the parties in $\mathcal{P}$. The sets TX and UTXO respectively define all the transactions accepted to date and all the unspent transaction outputs. The set of valid output conditions is represented by $\mathcal{V}$.

Public key Registration: Upon $(\text{register}, pk_P) \xleftrightarrow{\tau_0} P$, check if it is the first registration message received from $P \in \mathcal{P}$. If not drop the message, else add (pk_P, P) to PKI.

Post transaction: Upon $(\text{post}, \text{tx}) \xleftrightarrow{\tau_0} P$, check if $|\text{PKI}| = |\mathcal{P}|$, If not drop the message, else wait for $\tau \leq \Delta$ rounds where τ is selected by the adversary. Then, check if:

1. id uniqueness: For all $(t, \text{tx}') \in \text{TX}$, $\text{tx}'.\text{txid} \neq \text{tx.txid}$ holds.
2. Input and witness validity: For each $(\text{txid} \| i) \in \text{tx.Input}$, there exists $(t, \text{txid}, i, \theta) \in \text{UTXO}$ s.t. tx.Witness with inputs tx.nLT, the current round $\tau_0 + \tau$ and t satisfies $\theta.\varphi$.
3. Output validity: For each $\theta \in \text{tx.Output}$, $\theta.\text{Cash} > 0$ and $\theta.\varphi \in \mathcal{V}$ hold.
4. Value validity: Let $I := \{\text{utxo} := (t, \text{txid}, i, \theta) \mid \text{utxo} \in \text{UTXO} \wedge (\text{txid} \| i) \in \text{tx.Input}\}$. Then, $\sum_{\theta' \in \text{tx.Output}} \theta'.\text{cash} \leq \sum_{\text{utxo} \in I} \text{utxo}.\theta.\text{cash}$ holds.
5. Absolute time-lock validity: For the transaction tx, $\text{tx.nLT} \leq \tau_0 + \tau$ holds.

If any of the above checks fail, drop the message. Else, set $\text{TX} := \text{TX} \cup (\tau_0 + \tau, \text{tx})$, $\text{UTXO} := \text{UTXO} \setminus \text{I}$ and $\text{UTXO} := \text{UTXO} \cup \{(\tau_0 + \tau, \text{tx.txid}, i, \theta_i)\}_{i \in [n]}$ for $(\theta_1, \dots, \theta_n) := \text{tx.Output}$.

Let π be a protocol that has access to the global ledger $\mathcal{L}(\Delta, \Sigma)$ as well as the global clock $\mathcal{F}_{\text{clock}}$ and $\varphi_{\mathcal{F}}$ denote the ideal protocol for an ideal functionality $\mathcal{F}$ with access to the same global functionalities. Let $\text{EXE}_{\pi, \mathcal{A}, \mathcal{E}}^{\mathcal{L}(\Delta, \Sigma), \mathcal{F}_{\text{clock}}}(\lambda, z)$ denote the output of the environment $\mathcal{E}$ which interacts with a protocol π and an adversary $\mathcal{A}$ on input a security parameter λ and an auxiliary input z and similarly $\text{EXE}_{\varphi_{\mathcal{F}}, \mathcal{S}, \mathcal{E}}^{\mathcal{L}(\Delta, \Sigma), \mathcal{F}_{\text{clock}}}(n, z)$ denote the output of the environment $\mathcal{E}$ which interacts with a protocol $\varphi_{\mathcal{F}}$ and an adversary $\mathcal{S}$ (also called the simulator) on input a security parameter λ and an auxiliary input z .

The following definition is informally saying that should a protocol π UC-realise $\mathcal{F}$, any attack against the protocol π can be transformed into an attack against the ideal protocol $\varphi_{\mathcal{F}}$ and vice versa.

Definition 7.1. *A protocol π UC-realises an ideal functionality $\mathcal{F}$ with respect to a global ledger $\mathcal{L}(\Delta, \Sigma)$ and a global clock $\mathcal{F}_{\text{clock}}$ if for every adversary $\mathcal{A}$ there exists an adversary $\mathcal{S}$ such that we have*

$$\begin{aligned} \{\text{EXE}_{\pi, \mathcal{A}, \mathcal{E}}^{\mathcal{L}(\Delta, \Sigma), \mathcal{F}_{\text{clock}}}(\lambda, z)\}_{n \in \mathbb{N}, z \in \{0, 1\}^*} &\approx \\ \{\text{EXE}_{\varphi_{\mathcal{F}}, \mathcal{S}, \mathcal{E}}^{\mathcal{L}(\Delta, \Sigma), \mathcal{F}_{\text{clock}}}(\lambda, z)\}_{n \in \mathbb{N}, z \in \{0, 1\}^*} &\end{aligned} \quad (7.1)$$

where $\approx$ denotes computational indistinguishability.

7.10 Ideal Functionality

This section defines an ideal functionality $\mathcal{F}(t)$ with $t > \Delta$ that achieves the desired properties stated in Section 7.5.2. To simplify the notations, we abbreviate $\mathcal{F} := \mathcal{F}(t)$. The ideal functionality $\mathcal{F}$ stores a set Γ of all the created channels and their corresponding funding transactions. The set Γ can also be treated as a function s.t. $\Gamma(id) = (\gamma, \text{TX})$ with $\gamma.id = id$ if γ exists and $\Gamma(id) = \perp$ otherwise. Before presenting the ideal functionality $\mathcal{F}$ in detail, we briefly introduce its different phases and explain the way $\mathcal{F}$ achieves the desired properties.

a) Create: In this phase, $\mathcal{F}$ receives messages $(\text{INTRO}, \gamma, \text{tid}_P)$ and $(\text{CREATE}, \gamma.id)$ from both parties in rounds τ_0 and $\tau_0 + 1$, respectively, where tid_P specifies the funding source of the user P . Then, if the corresponding funding transaction appears on the ledger $\mathcal{L}$ within $2 + \Delta$ rounds, $\mathcal{F}$ sends the message $(\text{CREATED}, \gamma.id)$ to both parties and stores γ and the funding transaction in $\Gamma(\gamma.id)$. If the CREATE message is not received from both

parties but the funding transaction appears on $\mathcal{L}$ within $2 + \Delta$ rounds, $\mathcal{F}$ outputs Error. Since the message CREATED might be sent to the parties only if they both have sent the message CREATE to $\mathcal{F}$, the ideal functionality achieves consensus on creation.

b) Update: One of the parties, denoted by P , initiates this phase by sending the message $(\text{UPDATE}, id, \vec{\theta}, t_{stp})$ to $\mathcal{F}$, where id is the channel identifier, $\vec{\theta}$ is the new channel state and t_{stp} is the number of rounds needed to prepare prerequisites of the channel update (e.g. preparing the needed HTLCs). Due to disagreeing with the new state or failure in preparing its prerequisites, party Q can stop it by not sending the message $(\text{UPDATE} - \text{OK}, id)$ in step 2. Abort by P or Q in the next steps causes the procedure $\text{ForceClose}(id)$ to be executed. The property optimistic update is satisfied because if both parties act honestly, the channel can be updated without any blockchain interaction. Furthermore, if P or Q disagree to update the channel, they can stop sending the UPDATE or UPDATE - OK messages, respectively. This stops the channel update process without changing the latest channel state. Also, in cases where either P or Q stop cooperating, the procedure $\text{ForceClose}(id)$ is executed. This procedure takes at most Δ rounds to complete. This also guarantees consensus on update.

c) Close: If $\mathcal{F}$ receives the message (CLOSE, id) from both parties, a transaction TX is expected to appear on $\mathcal{L}$ within $\Delta + 1$ rounds. This transaction spends the output of the funding transaction and its outputs equal the latest channel state $\gamma.st$. If the CLOSE message is received only from one of the parties, $\mathcal{F}$ executes the procedure $\text{ForceClose}(id)$. In both cases, the output of the funding transaction must be spent within $\Delta + 1$ rounds. Otherwise, $\mathcal{F}$ outputs Error.

d) Punish: If a transaction TX spends the funding transaction's output of a channel γ , one of the following events is expected to occur: (1) another transaction appears on $\mathcal{L}$ within Δ rounds where this transaction spends output of TX and sends $\gamma.cash$ coins to the honest party P ; or (2) another transaction whose outputs correspond to the channel state $\gamma.st$ or $\gamma.st'$ appears on $\mathcal{L}$ within $t + \Delta$ rounds. Otherwise, $\mathcal{F}$ outputs Error. According to its definition, bounded closure with punish is achieved, if $\mathcal{F}$ returns no Error in the close and punish phases.

We describe the ideal functionality below. Normally, once $\mathcal{F}$ receives a message, it performs several validations on the message. But to simplify the description, we assume that messages are well-formed. Data exchange between $\mathcal{F}$ and other parties is represented by directed arrows. If $\mathcal{F}$ sends the message m to party P in round τ_0 , we denote it with $m \xrightarrow{\tau_0} P$. Similarly, if $\mathcal{F}$ is supposed to receive the message m from party P in round τ_0 , we denote it with $m \xleftarrow{\tau_0} P$.

Create

upon (INTRO, γ , tid_P) $\xleftrightarrow{\tau_0} P$:

- If (INTRO, γ , tid_Q) $\xleftrightarrow{\tau_0} Q$, then continue. Else stop.
- If (CREATE, id) $\xleftrightarrow{\tau_0+1} \gamma.users$:
 - Wait if in round $\tau_1 \leq \tau_0 + 3 + \Delta$ a transaction TX_{FU} with $TX_{FU}.Input = (tid_P, tid_Q)$ and $TX_{FU}.Output = \{(\gamma.cash, \phi)\}$ appears on the ledger $\mathcal{L}$. If yes, set $\Gamma(\gamma.id) := (\gamma, TX_{FU})$ and (CREATED, $\gamma.id$) $\xrightarrow{\tau_1} \gamma.users$. Else stop.

Otherwise:

- Wait if in round $\tau_1 \leq \tau_0 + 3 + \Delta$ a transaction TX_{FU} with $TX_{FU}.Input = (tid_P, tid_Q)$ and $TX_{FU}.Output = \{(\gamma.cash, \phi)\}$ appears on the ledger $\mathcal{L}$. If yes, Output Error $\xrightarrow{\tau_1} \gamma.users$. Else, stop.
-

Update

Upon (UPDATE, $id, \vec{\theta}, t_{stp}$) $\xleftrightarrow{\tau_0} P$, parse $(\gamma, TX) := \Gamma(id)$ and proceed as follows:

1. Send (UPDATE – REQ, $id, \vec{\theta}, t_{stp}$) $\xrightarrow{\tau_0+1} Q$.
2. If (UPDATE – OK, id) $\xleftrightarrow{\tau_1 \leq \tau_0+1+t_{stp}} Q$, then set $\gamma.flag := 2$ and $\gamma.st' := \vec{\theta}$ and send (SETUP, id) $\xrightarrow{\tau_1+1} P$. Else stop.
3. If (SETUP – OK, id) $\xrightarrow{\tau_1+1} P$, then (SETUP', id) $\xrightarrow{\tau_1+2} Q$. Else ForceClose(id) and stop.
4. If (SETUP' – OK, id) $\xrightarrow{\tau_1+2} Q$, then (UPDATE – OK, id) $\xrightarrow{\tau_1+3} P$. Else execute ForceClose(id) and stop.
5. If (REVOKE, id) $\xrightarrow{\tau_1+3} P$, then (REVOKE – REQ, id) $\xrightarrow{\tau_1+4} Q$. Else execute ForceClose(id) and stop.
6. If (REVOKE', id) $\xrightarrow{\tau_1+4} Q$, set $\gamma.st := \vec{\theta}$, $\gamma.flag := 1$, $\gamma.st' := \perp$, $\gamma.sn := \gamma.sn + 1$, $\Gamma(id) := (\gamma, TX)$, (UPDATED, id) $\xrightarrow{\tau_1+5} \gamma.Users$ and stop. Else execute ForceClose(id) and stop.

Close

upon $(\text{CLOSE}, id) \xleftrightarrow{\tau_0} P$, distinguish:

Both agreed: If $(\text{CLOSE}, id) \xleftrightarrow{\tau_0} Q$, let $(\gamma, \text{TX}_{\text{FU}}) := \Gamma(id)$ and distinguish:

- If in round $\tau_1 \leq \tau_0 + 1 + \Delta$, $\text{TX}_{\overline{\text{SP}}}$, with $\text{TX}_{\overline{\text{SP}}}.\text{Output} = \gamma.\text{st}$ and $\text{TX}_{\overline{\text{SP}}}.\text{Input} = \text{TX}_{\text{FU}}.\text{txid}\|1$ appears on $\mathcal{L}$, set $\Gamma(id) := (\perp, \text{TX}_{\text{FU}})$, $(\text{CLOSED}, id) \xrightarrow{\tau_1} \gamma.\text{users}$ and stop.
- If in round $\tau_0 + 1 + \Delta$, the TX_{FU} is still unspent, output $\text{Error} \xrightarrow{\tau_0+1+\Delta} \gamma.\text{users}$ and stop.

Q disagreed: Else, execute $\text{ForceClose}(id)$ in round $\tau_0 + 1$.

Punish (executed at the end of every round τ_0)

For each $(\gamma_i, \text{TX}_i) \in \Gamma$ check if there is a transaction TX on the ledger $\mathcal{L}$ s.t. $\text{TX}.\text{Input} = \text{TX}_i.\text{txid}\|1$ and $\gamma_i \neq \perp$. If yes, distinguish:

1. **Punish:** For the honest $P \in \gamma_i.\text{users}$, in round $\tau_1 \leq \tau_0 + \Delta$, a transaction TX_j with $\text{TX}_j.\text{Input} = \text{TX}.\text{txid}\|1$ and $\text{TX}_j.\text{Output} = (\gamma.\text{cash}, pk_P)$ appears on $\mathcal{L}$. Then, $(\text{PUNISHED}, id) \xrightarrow{\tau_1} P$, set $\Gamma(id) := (\perp, \text{TX}_i)$ and stop.
 2. **Close:** In round $\tau_1 \leq \tau_0 + t + \Delta$ a transaction TX_j appears on $\mathcal{L}$ where one of the following two sets of conditions hold: (1) $\gamma.\text{flag} = 1$, $\text{TX}_j.\text{Input} = \text{TX}.\text{txid}\|1$ and $\text{TX}_j.\text{Output} = \gamma.\text{st}$ or (2) $\gamma.\text{flag} = 2$, $\text{TX}_j.\text{Input} = \text{TX}.\text{txid}\|1$ and either $\text{TX}_j.\text{Output} = \gamma.\text{st}$ or $\text{TX}_j.\text{Output} = \gamma.\text{st}'$. Then, set $\Gamma(id) := (\perp, \text{TX}_i)$ and $(\text{CLOSED}, id) \xrightarrow{\tau_1} \gamma.\text{users}$.
 3. **Error:** Otherwise, $\text{Error} \xrightarrow{\tau_0+t+\Delta} \gamma.\text{users}$.
-

Subprocedure ForceClose(id)

Let τ_0 be the current round and $(\gamma, \text{TX}_{\text{FU}}) := \Gamma(id)$. If within Δ rounds, $\text{TX}_{\text{FU}}.\text{Output}$ is still an unspent output on $\mathcal{L}$, then output $\text{Error} \xrightarrow{\tau_0+\Delta} \gamma.\text{users}$.

7.10.1 Functionality Wrapper

The functionality $\mathcal{F}$ is supposed to perform several checks once he receives a message from another party. The functionality $\mathcal{F}$ must perform those checks in order to ensure that the received messages are well-formed. The following wrapper summarises those checks.

Functionality Wrapper: $\mathcal{W}_{\mathcal{F}}$

Create

Upon $(\text{INTRO}, \gamma, \text{tid}_P) \xleftrightarrow{\tau_0} P$ check if: $P \in \gamma.\text{users}$; $\Gamma(\gamma.\text{id}) \neq \perp$; there is no channel γ' with $\gamma.\text{id} = \gamma'.\text{id}$, $\gamma.\text{sn} = 0$; $\gamma.\text{st} = \{(c_P, \text{One} - \text{Sig}_{pk_P}), (c_Q, \text{One} - \text{Sig}_{pk_Q})\}$ with $c_P, c_Q \in \mathbb{R}^{>0}$ and $c_P + c_Q = \gamma.\text{cash}$; there exist $(t, \text{id}, i, \theta) \in \mathcal{L}.\text{UTXO}$ such that $\theta = (c_P, \text{One} - \text{Sig}_P)$ with $\text{id} \parallel i = \text{tid}$; and none of the other channels that are being created at the moment, must use tid_P . Drop the message if any above checks fail. Else proceed as $\mathcal{F}$.

Upon $(\text{CREATE}, \text{id}) \xleftrightarrow{\tau} P$ check if: you accepted messages $(\text{INTRO}, \gamma, \text{tid}_P) \xleftrightarrow{\tau_0} P$ and $(\text{INTRO}, \gamma, \text{tid}_Q) \xleftrightarrow{\tau_0} Q$ with $P, Q \in \gamma.\text{users}$, $\tau_0 + 1 = \tau$ and $\text{id} = \gamma.\text{id}$. Else proceed as $\mathcal{F}$.

Update

Upon $(\text{UPDATE}, \text{id}, \vec{\theta}, t_{stP}) \xleftrightarrow{\tau_0} P$ set $(\gamma, \text{TX}_{\text{FU}}) := \Gamma(\text{id})$ and check if: $\gamma \neq \perp$, there is no other update being preformed; let $\vec{\theta} = (\theta_1, \dots, \theta_l) = ((c_1, \varphi_1), \dots, (c_l, \varphi_l))$, then $\sum_{i \in [l] c_i = \gamma.\text{cash}}$ and $\varphi_i \in \mathcal{L}.\mathcal{V}$. Drop the message if any above checks fail. Else proceed as $\mathcal{F}$.

Upon $(\text{UPDATE} - \text{OK}, \text{id}) \xleftrightarrow{\tau} P$ check if: the message is a reply to the message $(\text{UPDATE} - \text{REQ}, \text{id}, \vec{\theta}, t_{stP})$ sent to P in round τ . If not, drop the message. Else proceed as $\mathcal{F}$.

Upon $(\text{SETUP} - \text{OK}, \text{id}) \xleftrightarrow{\tau} P$ check if: the message is a reply to the message $(\text{SETUP}, \text{id})$ sent to P in round τ_0 where $\tau = \tau_0 + t_{stP}$. If not, drop the message. Else proceed as $\mathcal{F}$.

Upon $(\text{SETUP}' - \text{OK}, \text{id}) \xleftrightarrow{\tau} P$ check if: the message is a reply to the message $(\text{SETUP}', \text{id})$ sent to P in round τ . If not, drop the message. Else proceed as $\mathcal{F}$.

Upon $(\text{REVOKE}, \text{id}) \xleftrightarrow{\tau} P$ check if: the message is a reply to the message $(\text{UPDATE} - \text{OK}, \text{id})$ sent to P in round τ . If not, drop the message. Else proceed as $\mathcal{F}$.

Upon $(\text{REVOKE}', id) \xleftarrow{\tau} P$ check if: the message is a reply to the message $(\text{REVOKE} - \text{REQ}, id)$ sent to P in round τ . If not, drop the message. Else proceed as $\mathcal{F}$.

Close

Upon $(\text{CLOSE}, id) \xleftarrow{\tau} P$, set $(\gamma, \text{TX}_{\text{FU}}) := \Gamma(id)$ and check if: $P \in \gamma.\text{users}$, $\Gamma(\gamma.\text{id}) \neq \perp$ and $\gamma.\text{flag} = 1$. Drop the message if any above checks fail. Else proceed as $\mathcal{F}$.

7.11 Daric Protocol

In this section, details of different phases of Daric will be presented. Before presenting the protocol, some notations are introduced. In different steps of the protocol, channel participants generate (or verify) some signatures on protocol transactions. When a signature with SIGHASH of type SIGHASH_ALL or ANYPREVOUT is going to be generated (or verified) for the transaction TX, the input message to the signing (or verification) algorithm is denoted by $f(\text{TX})$ or $\hat{f}(\overline{[\text{TX}]})$, respectively [58].

The set Γ^P , maintained by each party $P \in \mathcal{P}$, stores information of the latest channel state for all the open channels that P is involved in, where $\Gamma^P(id)$ corresponds with the channel with the identifier of id . In the channel update phase, while $\gamma.\text{flag} = 2$, the channel has two active states where the information about the new state is maintained by $P \in \mathcal{P}$ in $\Gamma'^P(id)$ with $id = \gamma.\text{id}$. To refer to the i^{th} element of $\Gamma'^P(id)$ we use $\Gamma'^P(id)[i]$ with $i \geq 1$. The party P also maintains a signature from his counterparty on the latest revocation transaction for each open channel that P is a party of. These signatures are maintained by P in the set Θ^P where $\Theta^P(id)$ corresponds with the channel with the identifier of id . We use directional arrows to show an exchange of messages. To simplify the protocol description, we remove some validations that must be normally done by channel parties. The protocol wrapper $\mathcal{W}_{\mathcal{P}}$ in Section 7.11.1 defines those validations.

Daric protocol: π

Create

Party P upon $(\text{INTRO}, \gamma, \text{tid}_p) \xleftarrow{\tau_0} \mathcal{E}$

1. Set $id := \gamma.id$, generate $(pk_{SP}^P, sk_{SP}^P) \leftarrow \text{Gen}$, $(pk_{RV}^P, sk_{RV}^P) \leftarrow \text{Gen}$ and $(pk_{RV}'^P, sk_{RV}'^P) \leftarrow \text{Gen}$ and send $(\text{createInfo}, id, tid_P, pk_{SP}^P, pk_{RV}^P, pk_{RV}'^P) \xrightarrow{\tau_0} Q$.
2. If $(\text{createInfo}, id, tid_Q, pk_{SP}^Q, pk_{RV}^Q, pk_{RV}'^Q) \xleftarrow{\tau_0+1} Q$, create:

$$\begin{aligned} [TX_{FU}] &:= \text{GenFund}((tid_P, tid_Q), \gamma) \\ ([TX_{CM,0}^P], [TX_{CM,0}^Q]) &:= \\ &\text{GenCommit}([TX_{FU}].txid \| 1, I_P, I_Q, 0) \\ \overline{[TX_{SP,0}]} &:= \text{GenSplit}(\gamma.st, 0) \end{aligned}$$

for $I_P := (pk_{SP}^P, pk_{RV}^P, pk_{RV}'^P)$ and $I_Q := (pk_{SP}^Q, pk_{RV}^Q, pk_{RV}'^Q)$. Else stop.

3. Compute $\hat{\sigma}_{TX_{SP,0}}^P := \text{Sign}_{sk_{SP}^P}(\hat{f}(\overline{[TX_{SP,0}]}))$ and $\sigma_{TX_{CM,0}^Q}^P := \text{Sign}_{sk_P}(f([TX_{CM,0}^Q]))$ and send $(\text{createCom}, id, \hat{\sigma}_{TX_{SP,0}}^P, \sigma_{TX_{CM,0}^Q}^P) \xrightarrow{\tau_0+1} Q$.
4. If $(\text{createCom}, id, \hat{\sigma}_{TX_{SP,0}}^Q, \sigma_{TX_{CM,0}^P}^Q) \xleftarrow{\tau_0+2} Q$, s.t. $\text{Vrfy}_{pk_{SP}^Q}(\hat{f}(\overline{[TX_{SP,0}]}); \hat{\sigma}_{TX_{SP,0}}^Q) = 1$ and also $\text{Vrfy}_{pk_Q}(f([TX_{CM,0}^P]); \sigma_{TX_{CM,0}^P}^Q) = 1$, then $\sigma_{TX_{FU}}^P := \text{Sign}_{sk_P}(f([TX_{FU}]))$ and send $(\text{createFund}, id, \sigma_{TX_{FU}}^P) \xrightarrow{\tau_0+2} Q$. Else stop.
5. If $(\text{createFund}, id, \sigma_{TX_{FU}}^Q) \xleftarrow{\tau_0+3} Q$, s.t. $\text{Vrfy}_{pk_Q}(f([TX_{FU}]); \sigma_{TX_{FU}}^Q) = 1$, create the transaction $TX_{FU} := (\mathcal{H}([TX_{FU}]), [TX_{FU}], ((x, \sigma_{TX_{FU}}^P), (y, \sigma_{TX_{FU}}^Q)))$ and $(\text{post}, TX_{FU}) \xrightarrow{\tau_0+3} \mathcal{L}$. Else, create a transaction TX with $TX.\text{Input} := tid_P$ and $TX.\text{Output}.\varphi := pk_P$ and $(\text{post}, TX) \xrightarrow{\tau_0+3} \mathcal{L}$.
6. If TX_{FU} is accepted by $\mathcal{L}$ in round $\tau_1 \leq \tau_0 + 3 + \Delta$, compute $\sigma_{TX_{CM,0}^P}^P = \text{Sign}_{sk_P}(f([TX_{CM,0}^P]))$, create $TX_{CM,0}^P := (\mathcal{H}([TX_{CM,0}^P]), [TX_{CM,0}^P], (1, \{\sigma_{TX_{CM,0}^P}^P, \sigma_{TX_{CM,0}^Q}^Q\}))$, set $\overline{TX_{SP,0}} := (\overline{[TX_{SP,0}]}, (1, \{\hat{\sigma}_{TX_{SP,0}}^P, \hat{\sigma}_{TX_{SP,0}}^Q\}))$, store $\Gamma^P(\gamma.id) := (\gamma, TX_{FU}, TX_{CM,0}^P, [TX_{CM,0}^Q], \overline{TX_{SP,0}})$ and $(\text{CREATED}, id) \xrightarrow{\tau_1} \mathcal{E}$. Else $\Gamma^P(\gamma.id) := (\perp, TX_{FU}, \perp, \perp, \perp)$ and stop.

Update

Party P upon $(\text{UPDATE}, id, \vec{\theta}, t_{stp}) \xleftarrow{\tau_0} \mathcal{E}$

1. Send (updateReq, $id, \vec{\theta}, t_{stp}$) $\xrightarrow{\tau_0} Q$.

Party Q upon (updateReq, $id, \vec{\theta}, t_{stp}$) $\xleftarrow{t_0} P$

2. Send (UPDATE – REQ, $id, \vec{\theta}, t_{stp}$) $\xrightarrow{t_0} \mathcal{E}$.
3. If (UPDATE – OK, id) $\xleftarrow{t_1 \leq t_0 + t_{stp}} \mathcal{E}$, then extract TX_{FU} and $i := \gamma.\text{sn}$ from $\Gamma^Q(id)$, create $([\text{TX}_{\text{CM}, i+1}^P], [\text{TX}_{\text{CM}, i+1}^Q]) := \text{GenCommit}([\text{TX}_{\text{FU}}].\text{txid} \| 1, I_P, I_Q, i+1)$, and $[\overline{\text{TX}_{\text{SP}, i+1}}] := \text{GenSplit}(\vec{\theta}, i+1)$ for $I_P := (pk_{\text{SP}}^P, pk_{\text{RV}}^P, pk_{\text{RV}}^{'P})$ and $I_Q := (pk_{\text{SP}}^Q, pk_{\text{RV}}^Q, pk_{\text{RV}}^{'Q})$, compute $\hat{\sigma}_{\text{TX}_{\text{SP}, i+1}}^Q = \text{Sign}_{sk_{\text{SP}}^Q}(\hat{f}([\overline{\text{TX}_{\text{SP}, i+1}}]))$, and send (updateInfo, $id, \hat{\sigma}_{\text{TX}_{\text{SP}, i+1}}^Q$) $\xrightarrow{t_1} P$.

Party P upon (updateInfo, $id, \hat{\sigma}_{\text{TX}_{\text{SP}, i+1}}^Q$) $\xleftarrow{\tau_1 \leq \tau_0 + 2 + t_{stp}} Q$

4. Extract TX_{FU} and $i := \gamma.\text{sn}$ from $\Gamma^P(id)$ and create the transactions $([\text{TX}_{\text{CM}, i+1}^P], [\text{TX}_{\text{CM}, i+1}^Q]) := \text{GenCommit}([\text{TX}_{\text{FU}}].\text{txid} \| 1, I_P, I_Q, i+1)$ and $[\overline{\text{TX}_{\text{SP}, i+1}}] := \text{GenSplit}(\vec{\theta}, i+1)$ for $I_P := (pk_{\text{SP}}^P, pk_{\text{RV}}^P, pk_{\text{RV}}^{'P})$ and $I_Q := (pk_{\text{SP}}^Q, pk_{\text{RV}}^Q, pk_{\text{RV}}^{'Q})$. If $\text{Vrfy}_{pk_{\text{SP}}^Q}(\hat{f}([\overline{\text{TX}_{\text{SP}, i+1}}]); \hat{\sigma}_{\text{TX}_{\text{SP}, i+1}}^Q) = 1$, compute $\hat{\sigma}_{\text{TX}_{\text{SP}, i+1}}^P := \text{Sign}_{sk_{\text{SP}}^P}(\hat{f}([\overline{\text{TX}_{\text{SP}, i+1}}]))$, set $[\overline{\text{TX}_{\text{SP}, i+1}}] := ([\overline{\text{TX}_{\text{SP}, i+1}}], (1, \{\hat{\sigma}_{\text{TX}_{\text{SP}, i+1}}^P, \hat{\sigma}_{\text{TX}_{\text{SP}, i+1}}^Q\}))$, store $\Gamma^P(id) = (\perp, [\text{TX}_{\text{CM}, i+1}^Q], [\overline{\text{TX}_{\text{SP}, i+1}}])$, set $\gamma.\text{flag} = 2$ and $\gamma.\text{st}' = \vec{\theta}$ and send (SETUP, id) $\xrightarrow{\tau_1} \mathcal{E}$. Else stop.
5. If (SETUP – OK, id) $\xleftarrow{\tau_1} \mathcal{E}$, sign $\sigma_{\text{TX}_{\text{CM}, i+1}^Q}^P := \text{Sign}_{sk_P}(f([\text{TX}_{\text{CM}, i+1}^Q]))$ and output (updateComP, $id, \hat{\sigma}_{\text{TX}_{\text{SP}, i+1}}^P, \sigma_{\text{TX}_{\text{CM}, i+1}^Q}^P$) $\xrightarrow{\tau_1} Q$. Else, execute $\text{ForceClose}^P(id)$ and stop.

Party Q

6. If (updateComP, $id, \hat{\sigma}_{\text{TX}_{\text{SP}, i+1}}^P, \sigma_{\text{TX}_{\text{CM}, i+1}^Q}^P$) $\xleftarrow{t_1+2} P$, such that $\text{Vrfy}_{pk_{\text{SP}}^P}(\hat{f}([\overline{\text{TX}_{\text{SP}, i+1}}]); \hat{\sigma}_{\text{TX}_{\text{SP}, i+1}}^P) = 1$ and also $\text{Vrfy}_{pk_P}(f([\text{TX}_{\text{CM}, i+1}^Q]); \sigma_{\text{TX}_{\text{CM}, i+1}^Q}^P) = 1$, compute $\sigma_{\text{TX}_{\text{CM}, i+1}^Q}^Q = \text{Sign}_{sk_Q}(f([\text{TX}_{\text{CM}, i+1}^Q]))$, set $[\overline{\text{TX}_{\text{SP}, i+1}}] := ([\overline{\text{TX}_{\text{SP}, i+1}}], (1, \{\hat{\sigma}_{\text{TX}_{\text{SP}, i+1}}^P, \hat{\sigma}_{\text{TX}_{\text{SP}, i+1}}^Q\}))$, set $\text{TX}_{\text{CM}, i+1}^Q = ([\text{TX}_{\text{CM}, i+1}^Q], (1, \{\sigma_{\text{TX}_{\text{CM}, i+1}^Q}^P, \sigma_{\text{TX}_{\text{CM}, i+1}^Q}^Q\}))$, $\gamma.\text{flag} = 2$ and $\gamma.\text{st}' = \vec{\theta}$, store $\Gamma^Q(id) = (\text{TX}_{\text{CM}, i+1}^Q, [\text{TX}_{\text{CM}, i+1}^P], [\overline{\text{TX}_{\text{SP}, i+1}}])$, and output (SETUP', id) $\xrightarrow{t_1+2} \mathcal{E}$. Else, execute $\text{ForceClose}^Q(id)$ and stop.
7. If (SETUP' – OK, id) $\xleftarrow{t_1+2} \mathcal{E}$, sign $\sigma_{\text{TX}_{\text{CM}, i+1}^Q}^Q = \text{Sign}_{sk_Q}(f([\text{TX}_{\text{CM}, i+1}^P]))$ and send (updateComQ, $id, \sigma_{\text{TX}_{\text{CM}, i+1}^Q}^Q$) $\xrightarrow{t_1+2} P$. Else, execute $\text{ForceClose}^Q(id)$ and stop.

Party P

8. If $(\text{updateComQ}, id, \sigma_{\text{TX}_{\text{CM},i+1}^P}^Q) \xleftarrow{\tau_1+2} Q$, s.t. $\text{Vrfy}_{pk_Q}(f([\text{TX}_{\text{CM},i+1}^P]); \sigma_{\text{TX}_{\text{CM},i+1}^P}^Q) = 1$, compute $\sigma_{\text{TX}_{\text{CM},i+1}^P}^P := \text{Sign}_{sk_P}(f([\text{TX}_{\text{CM},i+1}^P], 1))$, set $\text{TX}_{\text{CM},i+1}^P := (\mathcal{H}([\text{TX}_{\text{CM},i+1}^P]), [\text{TX}_{\text{CM},i+1}^P], (1, \{\sigma_{\text{TX}_{\text{CM},i+1}^P}^P, \sigma_{\text{TX}_{\text{CM},i+1}^Q}^Q\}))$, set $\Gamma^P(id)[1] := \text{TX}_{\text{CM},i+1}^P$ and then output $(\text{UPDATE} - \text{OK}, id) \xrightarrow{\tau_1+2} \mathcal{E}$. Else, execute the procedure $\text{ForceClose}^P(id)$ and stop.
9. If $(\text{REVOKE}, id) \xleftarrow{\tau_1+2} \mathcal{E}$, create $([\text{TX}_{\text{RV},i}^P], [\text{TX}_{\text{RV},i}^Q]) := \text{GenRevoke}(pk_P, pk_Q, \gamma.\text{cash}, i + 1)$, compute $\hat{\sigma}_{\text{TX}_{\text{RV},i}^Q}^P := \text{Sign}_{sk_{\text{RV}}}^P(\hat{f}([\text{TX}_{\text{RV},i}^Q]))$ if $P = A$ or $\hat{\sigma}_{\text{TX}_{\text{RV},i}^Q}^P = \text{Sign}_{sk_{\text{RV}}}^{P'}(\hat{f}([\text{TX}_{\text{RV},i}^Q]))$ otherwise and send $(\text{revokeP}, id, \hat{\sigma}_{\text{TX}_{\text{RV},i}^Q}^P) \xrightarrow{\tau_1+2} Q$. Else, execute the procedure $\text{ForceClose}^P(id)$ and stop.

Party Q

10. Create $([\text{TX}_{\text{RV},i}^P], [\text{TX}_{\text{RV},i}^Q]) := \text{GenRevoke}(pk_P, pk_Q, \gamma.\text{cash}, i + 1)$. If $(\text{revokeP}, id, \hat{\sigma}_{\text{TX}_{\text{RV},i}^Q}^P) \xleftarrow{t_1+4} P$, such that $\text{Vrfy}_{pk_{\text{RV}}}^P(\hat{f}([\text{TX}_{\text{RV},i}^Q]); \hat{\sigma}_{\text{TX}_{\text{RV},i}^Q}^P) = 1$ if $Q = B$ or $\text{Vrfy}_{pk_{\text{RV}}}^{P'}(\hat{f}([\text{TX}_{\text{RV},i}^Q]); \hat{\sigma}_{\text{TX}_{\text{RV},i}^Q}^P) = 1$ otherwise, set $\Theta^Q(id) := (\hat{\sigma}_{\text{TX}_{\text{RV},i}^Q}^P), \gamma.\text{sn} := i + 1, \gamma.\text{st} := \vec{\theta}$, $\Gamma^Q(id) := (\gamma, \text{TX}_{\text{FU}}, \text{TX}_{\text{CM},i+1}^Q, [\text{TX}_{\text{CM},i+1}^P], \text{TX}_{\text{SP},i+1}), \gamma.\text{flag} = 1$, and $\Gamma^Q(id) = (\perp, \perp, \perp)$ and send $(\text{REVOKE} - \text{REQ}, id) \xrightarrow{t_1+4} \mathcal{E}$. Else, execute the procedure $\text{ForceClose}^Q(id)$ and stop.
11. If $(\text{REVOKE}', id) \xleftarrow{t_1+4} \mathcal{E}$, then compute $\hat{\sigma}_{\text{TX}_{\text{RV},i}^Q}^Q := \text{Sign}_{sk_{\text{RV}}}^Q(\hat{f}([\text{TX}_{\text{RV},i}^P]))$ if $Q = A$ or compute $\hat{\sigma}_{\text{TX}_{\text{RV},i}^Q}^Q := \text{Sign}_{sk_{\text{RV}}}^{Q'}(\hat{f}([\text{TX}_{\text{RV},i}^P]))$ otherwise, output $(\text{revokeQ}, id, \hat{\sigma}_{\text{TX}_{\text{RV},i}^Q}^Q) \xrightarrow{t_1+4} P$ and $(\text{UPDATED}, id) \xrightarrow{t_1+5} \mathcal{E}$. Else, execute the procedure $\text{ForceClose}^Q(id)$ and stop.

Party P

12. If $(\text{revokeQ}, id, \hat{\sigma}_{\text{TX}_{\text{RV},i}^Q}^Q) \xleftarrow{\tau_1+4} Q$, such that it holds that $\text{Vrfy}_{pk_{\text{RV}}}^Q(\hat{f}([\text{TX}_{\text{RV},i}^P]); \hat{\sigma}_{\text{TX}_{\text{RV},i}^Q}^Q) = 1$ if $P = B$ or $\text{Vrfy}_{pk_{\text{RV}}}^{Q'}(\hat{f}([\text{TX}_{\text{RV},i}^P]); \hat{\sigma}_{\text{TX}_{\text{RV},i}^Q}^Q) = 1$ otherwise, assign $\Theta^P(id) := (\hat{\sigma}_{\text{TX}_{\text{RV},i}^Q}^Q), \gamma.\text{sn} := i + 1, \gamma.\text{st} := \vec{\theta}, \Gamma^P(id) := (\gamma, \text{TX}_{\text{FU}}, \text{TX}_{\text{CM},i+1}^P, [\text{TX}_{\text{CM},i+1}^Q], \overline{\text{TX}_{\text{SP},i+1}}), \gamma.\text{flag} := 1$, and $\Gamma^P(id) := (\perp, \perp, \perp)$ and send $(\text{UPDATED}, id) \xrightarrow{\tau_1+4} \mathcal{E}$. Else, execute the procedure $\text{ForceClose}^P(id)$ and stop.

Close

Party P upon $(\text{CLOSE}, id) \xleftarrow{\tau_0} \mathcal{E}$

1. Extract $\text{TX}_{\text{FU}}, i := \gamma.sn, \text{TX}_{\text{SP},i}$ from $\Gamma^P(id)$ and create $[\text{TX}_{\text{SP}}] := \text{GenFinSplit}(\text{TX}_{\text{FU}}.\text{txid}\|1, \gamma.st)$.
2. Compute $\sigma_{\text{TX}_{\text{SP}}}^P := \text{Sign}_{sk_P}(f([\text{TX}_{\text{SP}}]))$ and send $(\text{CloseP}, id, \sigma_{\text{TX}_{\text{SP}}}^P) \xrightarrow{\tau_0} Q$.
3. If $(\text{CloseQ}, id, \sigma_{\text{TX}_{\text{SP}}}^Q) \xleftarrow{\tau_0+1} Q$ s.t. $\text{Vrfy}_{pk_Q}(f([\text{TX}_{\text{SP}}])); \sigma_{\text{TX}_{\text{SP}}}^Q = 1$, create the transaction $\text{TX}_{\text{SP}} := (\mathcal{H}([\text{TX}_{\text{SP}}]), [\text{TX}_{\text{SP}}], (1, \{\sigma_{\text{TX}_{\text{SP}}}^P, \sigma_{\text{TX}_{\text{SP}}}^Q\}))$ and send $(\text{post}, \text{TX}_{\text{SP}}) \xrightarrow{\tau_0+1} \mathcal{L}$. Else, execute the procedure $\text{ForceClose}^P(id)$ and stop.
4. If in round $\tau_1 \leq \tau_0 + 1 + \Delta$, the transaction TX_{SP} is accepted by $\mathcal{L}$, set $\Gamma^P(id) := \perp$, $\Theta^P(id) := \perp$ and send $(\text{CLOSED}, id) \xrightarrow{\tau_1} \mathcal{E}$.

Punish (executed at the end of every round τ_0)

Party P

For each $id \in \{0, 1\}^*$, extract $i := \gamma.sn$ and $flag := \gamma.flag$ from $\Gamma^P(id)$.

If $flag = 1$:

- Set $(\gamma, \text{TX}_{\text{FU}}, \text{TX}_{\text{CM},i}^P, [\text{TX}_{\text{CM},i}^Q], \overline{\text{TX}_{\text{SP},i}}) := \Gamma^P(id)$ and $I := \{[\text{TX}_{\text{CM},i}^P], [\text{TX}_{\text{CM},i}^Q]\}$. Check if $\text{TX}_{\text{FU}}.\text{Output}$ is spent by a transaction TX s.t. $[\text{TX}] \notin I$. If yes:
 - Create $([\text{TX}_{\text{RV},i-1}^P], [\text{TX}_{\text{RV},i-1}^Q]) := \text{GenRevoke}(pk_P, pk_Q, \gamma.cash, i)$ and then set $[\text{TX}_{\text{RV},i-1}^P] := (\text{TX}.\text{txid}\|1, [\text{TX}_{\text{RV},i-1}^P])$.
 - compute $\hat{\sigma}_{\text{TX}_{\text{RV},i-1}^P}^P = \text{Sign}_{sk_{\text{RV}}}^P(\hat{f}([\text{TX}_{\text{RV},i-1}^P]))$ if $P = B$ or $\hat{\sigma}_{\text{TX}_{\text{RV},i-1}^P}^P = \text{Sign}_{sk_{\text{RV}}}^P(\hat{f}([\text{TX}_{\text{RV},i-1}^P]))$ otherwise.
 - Set $\hat{\sigma}_{\text{TX}_{\text{RV},i-1}^Q}^Q := \Theta^P(id)$, and create $\text{TX}_{\text{RV},i-1}^P := (\mathcal{H}([\text{TX}_{\text{RV},i-1}^P]), [\text{TX}_{\text{RV},i-1}^P], (2, \{\hat{\sigma}_{\text{TX}_{\text{RV},i-1}^P}^P, \hat{\sigma}_{\text{TX}_{\text{RV},i-1}^Q}^Q\}))$.

- Post (post, $\text{TX}_{\text{RV},i-1}^P$) $\xrightarrow{\tau_0} \mathcal{L}$.
- Let $\text{TX}_{\text{RV},i-1}^P$ be accepted by $\mathcal{L}$ in round $\tau_1 \leq \tau_0 + \Delta$. Set $\Theta^P(id) := \perp$, $\Gamma^P(id) := \perp$ and output (PUNISHED, id) $\xrightarrow{\tau_1} \mathcal{E}$.

If no, set $[\text{TX}_{\text{SP},i}] := (\text{TX.txid} \| 1, \overline{\text{TX}_{\text{SP},i}.\text{nLT}}, \overline{\text{TX}_{\text{SP},i}.\text{Output}})$, set $\text{TX}_{\text{SP},i} := (\mathcal{H}([\text{TX}_{\text{SP},i}]), \text{TX.txid} \| 1, \overline{\text{TX}_{\text{SP},i}})$ and then post (post, $\text{TX}_{\text{SP},i}$) $\xrightarrow{\tau_0+t} \mathcal{L}$. Let TX be spent in round $\tau_1 \leq \tau_0 + t + \Delta$. Set $\Theta^P(id) := \perp$, $\Gamma^P(id) := \perp$ and output (CLOSED, id) $\xrightarrow{\tau_1} \mathcal{E}$.

If $flag = 2$:

- Set $(\gamma, \text{TX}_{\text{FU}}, \text{TX}_{\text{CM},i}^P, [\text{TX}_{\text{CM},i}^Q], \overline{\text{TX}_{\text{SP},i}}) := \Gamma^P(id)$, and $(\text{TX}_{\text{CM},i+1}^P, [\text{TX}_{\text{CM},i+1}^Q], \overline{\text{TX}_{\text{SP},i+1}}) := \Gamma'^P(id)$ and also $I := \{[\text{TX}_{\text{CM},i}^P], [\text{TX}_{\text{CM},i}^Q], [\text{TX}_{\text{CM},i+1}^P], [\text{TX}_{\text{CM},i+1}^Q]\}$. Check if $\text{TX}_{\text{FU}}.\text{Output}$ is spent by a transaction TX s.t. $[\text{TX}] \notin I$. If yes:

- Create $([\text{TX}_{\text{RV},i-1}^P], [\text{TX}_{\text{RV},i-1}^Q]) := \text{GenRevoke}(pk_P, pk_Q, \gamma.\text{cash}, i)$. and then set $[\text{TX}_{\text{RV},i-1}^P] := (\text{TX.txid} \| 1, [\text{TX}_{\text{RV},i-1}^P])$.
- compute $\hat{\sigma}_{\text{TX}_{\text{RV},i-1}^P}^P := \text{Sign}_{sk_{\text{RV}}^P}(\hat{f}([\text{TX}_{\text{RV},i-1}^P]))$ if $P = B$ or $\hat{\sigma}_{\text{TX}_{\text{RV},i-1}^P}^P := \text{Sign}_{sk_{\text{RV}}^P}(\hat{f}([\text{TX}_{\text{RV},i-1}^P]))$ otherwise.
- Set $\hat{\sigma}_{\text{TX}_{\text{RV},i-1}^Q}^Q := \Theta^P(id)$, and create $\text{TX}_{\text{RV},i-1}^P := (\mathcal{H}([\text{TX}_{\text{RV},i-1}^P]), [\text{TX}_{\text{RV},i-1}^P], (2, \{\hat{\sigma}_{\text{TX}_{\text{RV},i-1}^P}^P, \hat{\sigma}_{\text{TX}_{\text{RV},i-1}^Q}^Q\}))$.
- Post (post, $\text{TX}_{\text{RV},i-1}^P$) $\xrightarrow{\tau_0} \mathcal{L}$.
- Let $\text{TX}_{\text{RV},i-1}^P$ be accepted by $\mathcal{L}$ in round $\tau_1 \leq \tau_0 + \Delta$. Set $\Theta^P(id) := \perp$, $\Gamma^P(id) := \perp$ and output (PUNISHED, id) $\xrightarrow{\tau_1} \mathcal{E}$.

If no, set $[\text{TX}_{\text{SP},i+1}] := (\text{TX.txid} \| 1, \overline{\text{TX}_{\text{SP},i+1}.\text{nLT}}, \overline{\text{TX}_{\text{SP},i+1}.\text{Output}})$, $\text{TX}_{\text{SP},i+1} := (\mathcal{H}([\text{TX}_{\text{SP},i+1}]), \text{TX.txid} \| 1, \overline{\text{TX}_{\text{SP},i+1}})$, wait for t rounds and post (post, $\text{TX}_{\text{SP},i+1}$) $\xrightarrow{\tau_0+t} \mathcal{L}$. Let $\text{TX}_{\text{SP},i+1}$ be accepted by $\mathcal{L}$ in round $\tau_1 \leq \tau_0 + t + \Delta$. Then, set $\Theta^P(id) := \perp$, $\Gamma^P(id) := \perp$ and output (CLOSED, id) $\xrightarrow{\tau_1} \mathcal{E}$.

Subprocedures

$\text{ForceClose}^P(id)$:

Let τ_0 be the current round. Extract $i := \gamma.\text{sn}$, $flag := \gamma.\text{flag}$, $\text{TX}_{\text{CM},i}^P$ and $\overline{\text{TX}_{\text{SP},i}}$ from $\Gamma^P(id)$. Also, extract $\text{TX}_{\text{CM},i+1}^P$ and $\overline{\text{TX}_{\text{SP},i+1}}$ from $\Gamma'^P(id)$ if $flag = 2$.

If (1) $flag = 1$ or (2) $flag = 2$ and $TX_{CM,i+1}^P = \perp$, then $\text{post}(\text{post}, TX_{CM,i}^P) \xrightarrow{\tau_0} \mathcal{L}$. Otherwise, $\text{post}(\text{post}, TX_{CM,i+1}^P) \xrightarrow{\tau_0} \mathcal{L}$.

(Publishing the corresponding split transaction takes place in the Punish phase.)

$\text{GenFund}((tid_P, tid_Q), \gamma)$:

Return $[TX_{FU}]$ where $[TX_{FU}].\text{Input} := (tid_P, tid_Q)$, $[TX_{FU}].\text{nLT} := 0$ and $[TX_{FU}].\text{Output} := (\gamma.\text{Cash}, pk_P \wedge pk_Q)$

$\text{GenCommit}([TX_{FU}].\text{txid}\|1, (pk_{SP}^P, pk_{RV}^P, pk_{RV}'^P), (pk_{SP}^Q, pk_{RV}^Q, pk_{RV}'^Q), i)$:

Return $[TX_{CM,i}^P]$ and $[TX_{CM,i}^Q]$ where $TX_{CM,i}^P.\text{nLT} := 0$, $TX_{CM,i}^P.\text{Input} := TX_{FU}.\text{txid}\|1$, and $TX_{CM,i}^P.\text{Output} := \{(TX_{FU}.\text{Output}.\text{cash}, (\varphi_1 \vee \varphi_2))\}$, $TX_{CM,i}^Q.\text{nLT} := 0$, $TX_{CM,i}^Q.\text{Input} := TX_{FU}.\text{txid}\|1$, and $TX_{CM,i}^Q.\text{Output} := \{(TX_{FU}.\text{Output}.\text{cash}, (\varphi_1 \vee \varphi_2'))\}$ with $\varphi_1 := (pk_{SP}^P \wedge pk_{SP}^Q \wedge CSV_t \wedge CLTV_{S_0+i})$, $\varphi_2 := (pk_{RV}^P \wedge pk_{RV}^Q \wedge CLTV_{S_0+i})$, $\varphi_2' := (pk_{RV}'^P \wedge pk_{RV}'^Q \wedge CLTV_{S_0+i})$.

$\text{GenSplit}(\vec{\theta}, i)$:

Return $\overline{[TX_{SP,i}]}$ where $\overline{[TX_{SP,i}]}.\text{nLT} := S_0 + i$ and $\overline{[TX_{SP,i}]}.\text{Output} := \vec{\theta}$.

$\text{GenRevoke}(pk_P, pk_Q, \gamma.\text{cash}, i + 1)$:

Return $\overline{[TX_{RV,i}^P]}$ and $\overline{[TX_{RV,i}^Q]}$ where $\overline{[TX_{RV,i}^P]}.\text{nLT} := S_0 + i$, $\overline{[TX_{RV,i}^P]}.\text{Output} := \{(\gamma.\text{cash}, pk_P)\}$, $\overline{[TX_{RV,i}^Q]}.\text{nLT} := S_0 + i$ and $\overline{[TX_{RV,i}^Q]}.\text{Output} := \{(\gamma.\text{cash}, pk_Q)\}$

7.11.1 Protocol Wrapper

Each party in Daric protocol π is supposed to perform several checks once he receives a message from another party. Parties perform those checks to ensure that the received messages are well-formed. The following wrapper summarises those checks.

Protocol Wrapper: $\mathcal{W}_{\mathcal{P}}$

Create

Upon $(\text{INTRO}, \gamma, \text{tid}_P) \xleftrightarrow{\tau_0} \mathcal{E}$ check if: $P \in \gamma.\text{users}$; $\Gamma(\gamma.\text{id}) \neq \perp$; there is no channel γ' with $\gamma.\text{id} = \gamma'.\text{id}$, $\gamma.\text{sn} = 0$; $\gamma.\text{st} = \{(c_P, \text{One} - \text{Sig}_{pk_P}), (c_Q, \text{One} - \text{Sig}_{pk_Q})\}$ with $c_P, c_Q \in \mathbb{R}^{>0}$ and $c_P + c_Q = \gamma.\text{cash}$; there exist $(t, \text{id}, i, \theta) \in \mathcal{L}.\text{UTXO}$ such that $\theta = (c_P, \text{One} - \text{Sig}_P)$ with $\text{id} \parallel i = \text{tid}$; and none of the other channels that are being created at the moment, must use tid_P . Drop the message if any above checks fails. Else proceed as P in Daric protocol.

Upon $(\text{CREATE}, \text{id}) \xleftrightarrow{\tau} \mathcal{E}$ check if: you accepted messages $(\text{INTRO}, \gamma, \text{tid}_P) \xleftrightarrow{\tau_0} P$ and $(\text{INTRO}, \gamma, \text{tid}_Q) \xleftrightarrow{\tau_0} Q$ with $P, Q \in \gamma.\text{users}$, $\tau_0 + 1 = \tau$ and $\text{id} = \gamma.\text{id}$. Else proceed as P in Daric protocol.

Update

Upon $(\text{UPDATE}, \text{id}, \theta, t_{stp}) \xleftrightarrow{\tau_0} \mathcal{E}$ set $(\gamma, \text{TX}_{\text{FU}}) := \Gamma(\text{id})$ and check if: $\gamma \neq \perp$, there is no other update being preformed; let $\theta = (\theta_1, \dots, \theta_l) = ((c_1, \varphi_1), \dots, (c_l, \varphi_l))$, then $\sum_{i \in [l]} c_i = \gamma.\text{cash}$ and $\varphi_i \in \mathcal{L}.\mathcal{V}$. Drop the message if any above checks fails. Else proceed as P in Daric protocol.

Upon $(\text{UPDATE} - \text{OK}, \text{id}) \xleftrightarrow{\tau} \mathcal{E}$ check if: the message is a reply to the message $(\text{UPDATE} - \text{REQ}, \text{id}, \theta, t_{stp}) \xrightarrow{\tau}$ sent to $\mathcal{E}$ in round τ . If not, drop the message. Else proceed as P in Daric protocol.

Upon $(\text{SETUP} - \text{OK}, \text{id}) \xleftrightarrow{\tau} \mathcal{E}$ check if: the message is a reply to the message $(\text{SETUP}, \text{id})$ sent to $\mathcal{E}$ in round τ_0 where $\tau = \tau_0 + t_{stp}$. If not, drop the message. Else proceed as P in Daric protocol.

Upon $(\text{SETUP}' - \text{OK}, \text{id}) \xleftrightarrow{\tau} \mathcal{E}$ check if: the message is a reply to the message $(\text{SETUP}', \text{id})$ sent to $\mathcal{E}$ in round τ . If not, drop the message. Else proceed as P in Daric protocol.

Upon $(\text{REVOKE}, \text{id}) \xleftrightarrow{\tau} \mathcal{E}$ check if: the message is a reply to the message $(\text{UPDATE} - \text{OK}, \text{id})$ sent to $\mathcal{E}$ in round τ . If not, drop the message. Else proceed as P in Daric protocol.

Upon $(\text{REVOKE}', \text{id}) \xleftrightarrow{\tau} \mathcal{E}$ check if: the message is a reply to the message $(\text{REVOKE} - \text{REQ}, \text{id})$ sent to $\mathcal{E}$ in round τ . If not, drop the message. Else proceed as P in Daric protocol.

Close

Upon $(\text{CLOSE}, \text{id}) \xleftrightarrow{\tau} \mathcal{E}$, set $(\gamma, \text{TX}_{\text{FU}}) := \Gamma(\text{id})$ and check if: $P \in \gamma.\text{users}$, $\Gamma(\gamma.\text{id}) \neq \perp$ and $\gamma.\text{flag} = 1$. Drop the message if any above checks fails. Else proceed as P in Daric protocol.

7.12 Security Analysis

In this section, we prove the Theorem 7.1. To do so, a simulator for the protocol π in the ideal world is provided and then we formally show that the Daric protocol (introduced in Section 7.11) UC-realises the ideal functionality $\mathcal{F}$ (introduced in Section 7.10).

Simulator:

Create

Case A is honest and B is corrupted.

Upon A sending $(\text{INTRO}, \gamma, \text{tid}_A) \xrightarrow{\tau_0} \mathcal{F}$,

1. Set $\text{id} := \gamma.\text{id}$, generate $(pk_{\text{SP}}^A, sk_{\text{SP}}^A) \leftarrow \text{Gen}$, $(pk_{\text{RV}}^A, sk_{\text{RV}}^A) \leftarrow \text{Gen}$ and $(pk_{\text{RV}}'^A, sk_{\text{RV}}'^A) \leftarrow \text{Gen}$ and send $(\text{createInfo}, \text{id}, \text{tid}_A, pk_{\text{SP}}^A, pk_{\text{RV}}^A, pk_{\text{RV}}'^A) \xrightarrow{\tau_0} B$.
2. If B sends $(\text{createInfo}, \text{id}, \text{tid}_B, pk_{\text{SP}}^B, pk_{\text{RV}}^B, pk_{\text{RV}}'^B) \xrightarrow{\tau_0} A$, then $(\text{INTRO}, \gamma, \text{tid}_B) \xrightarrow{\tau_0} \mathcal{F}$ on behalf of B . Else stop.
3. If A sends $(\text{CREATE}, \text{id}) \xrightarrow{\tau_0+1} \mathcal{F}$, then create $[\text{TX}_{\text{FU}}] := \text{GenFund}((\text{tid}_A, \text{tid}_B), \gamma)$, $([\text{TX}_{\text{CM},0}^A], [\text{TX}_{\text{CM},0}^B]) := \text{GenCommit}([\text{TX}_{\text{FU}}].\text{txid}\|1, I_A, I_B, 0)$, and $[\text{TX}_{\text{SP},0}] := \text{GenSplit}(\gamma.\text{st}, 0)$ for $I_A := (pk_{\text{SP}}^A, pk_{\text{RV}}^A, pk_{\text{RV}}'^A)$ and $I_B := (pk_{\text{SP}}^B, pk_{\text{RV}}^B, pk_{\text{RV}}'^B)$. Else stop.
4. Compute $\hat{\sigma}_{\text{TX}_{\text{SP},0}}^A = \text{Sign}_{sk_{\text{SP}}^A}(\hat{f}(\overline{[\text{TX}_{\text{SP},0}]}))$ and $\sigma_{\text{TX}_{\text{CM},0}}^A = \text{Sign}_{sk_{\text{CM}}^A}(f([\text{TX}_{\text{CM},0}^B]))$ and send $(\text{createCom}, \text{id}, \hat{\sigma}_{\text{TX}_{\text{SP},0}}^A, \sigma_{\text{TX}_{\text{CM},0}}^A) \xrightarrow{\tau_0+1} B$.
5. If B sends $(\text{createCom}, \text{id}, \hat{\sigma}_{\text{TX}_{\text{SP},0}}^B, \sigma_{\text{TX}_{\text{CM},0}}^B) \xrightarrow{\tau_0+1} A$, s.t. $\text{Vrfy}_{pk_{\text{SP}}^B}(\hat{f}(\overline{[\text{TX}_{\text{SP},0}]}); \hat{\sigma}_{\text{TX}_{\text{SP},0}}^B) = 1$ and $\text{Vrfy}_{pk_{\text{CM}}^B}(f([\text{TX}_{\text{CM},0}^A]); \sigma_{\text{TX}_{\text{CM},0}}^B) = 1$, send $(\text{CREATE}, \text{id}) \xrightarrow{\tau_0+1} \mathcal{F}$ on behalf of B . Else stop.
6. Compute $\sigma_{\text{TX}_{\text{FU}}}^A = \text{Sign}_{sk_A}(f([\text{TX}_{\text{FU}}]))$ and send $(\text{createFund}, \text{id}, \sigma_{\text{TX}_{\text{FU}}}^A) \xrightarrow{\tau_0+2} B$.
7. If B sends $(\text{createFund}, \text{id}, \sigma_{\text{TX}_{\text{FU}}}^B) \xrightarrow{\tau_0+2} A$, s.t. $\text{Vrfy}_{pk_B}(f([\text{TX}_{\text{FU}}]); \sigma_{\text{TX}_{\text{FU}}}^B) = 1$, create $\text{TX}_{\text{FU}} := (\mathcal{H}([\text{TX}_{\text{FU}}]), [\text{TX}_{\text{FU}}], ((x, \sigma_{\text{TX}_{\text{FU}}}^A), (y, \sigma_{\text{TX}_{\text{FU}}}^B)))$ and $(\text{post}, \text{TX}_{\text{FU}}) \xrightarrow{\tau_0+3} \mathcal{L}$. Else create a transaction TX with $\text{TX}.\text{input} := \text{tid}_A$ and $\text{TX}.\text{Output}.\varphi := pk_A$ and $(\text{post}, \text{TX}) \xrightarrow{\tau_0+3} \mathcal{L}$.

8. If TX_{FU} is accepted by $\mathcal{L}$ in round $\tau_1 \leq \tau_0 + 3 + \Delta$, compute $\sigma_{\text{TX}_{\text{CM},0}}^A = \text{Sign}_{sk_{\text{CM}}^A}(f([\text{TX}_{\text{CM},0}^A]))$, create $\text{TX}_{\text{CM},0}^A := (\mathcal{H}([\text{TX}_{\text{CM},0}^A]), [\text{TX}_{\text{CM},0}^A], (1, \{\sigma_{\text{TX}_{\text{CM},0}}^A, \sigma_{\text{TX}_{\text{CM},0}}^B\}))$, set $\overline{\text{TX}_{\text{SP},0}} := (\overline{[\text{TX}_{\text{SP},0}]}, (1, \{\hat{\sigma}_{\text{TX}_{\text{SP},0}}^A, \hat{\sigma}_{\text{TX}_{\text{SP},0}}^B\}))$, store $\Gamma^A(\gamma.id) := (\gamma, \text{TX}_{\text{FU}}, \text{TX}_{\text{CM},0}^A, [\text{TX}_{\text{CM},0}^B], \overline{\text{TX}_{\text{SP},0}})$ and $(\text{CREATED}, id) \xrightarrow{\tau_1} \mathcal{E}$. Else $\Gamma^A(\gamma.id) := (\perp, \text{TX}_{\text{FU}}, \perp, \perp, \perp)$ and stop.

Update

Case A is honest and B is corrupted.

Upon A sending $(\text{UPDATE}, id, \vec{\theta}, t_{stp}) \xrightarrow{\tau_0} \mathcal{F}$, proceed as follows:

1. Send $(\text{updateReq}, id, \vec{\theta}, t_{stp}) \xrightarrow{\tau_0} B$.
2. If B sends $(\text{updateInfo}, id, \hat{\sigma}_{\text{SP},i+1}^B) \xrightarrow{\tau_1 \leq \tau_0 + 1 + t_{stp}} A$, extract TX_{FU} and $i := \gamma.sn$ from $\Gamma^A(id)$ and create $([\text{TX}_{\text{CM},i+1}^A], [\text{TX}_{\text{CM},i+1}^B]) := \text{GenCommit}([\text{TX}_{\text{FU}}].\text{txid}\|1, I_A, I_B, i+1)$, and $\overline{[\text{TX}_{\text{SP},i+1}]} := \text{GenSplit}(\gamma.st, i+1)$ for $I_A := (pk_{\text{SP}}^A, pk_{\text{RV}}^A, pk_{\text{RV}}'^A)$ and $I_B := (pk_{\text{SP}}^B, pk_{\text{RV}}^B, pk_{\text{RV}}'^B)$. If $\text{Vrfy}_{pk_{\text{SP}}^B}(\hat{f}(\overline{[\text{TX}_{\text{SP},i+1}]}) ; \hat{\sigma}_{\text{TX}_{\text{SP},i+1}}^B) = 1$, then compute $\hat{\sigma}_{\text{TX}_{\text{SP},i+1}}^A := \text{Sign}_{sk_{\text{SP}}^A}(\hat{f}(\overline{[\text{TX}_{\text{SP},i+1}]})$, set $\overline{\text{TX}_{\text{SP},i+1}} := (\overline{[\text{TX}_{\text{SP},i+1}]}, (1, \{\hat{\sigma}_{\text{TX}_{\text{SP},i+1}}^A, \hat{\sigma}_{\text{TX}_{\text{SP},i+1}}^B\}))$, store $\Gamma^A(id) := (\perp, [\text{TX}_{\text{CM},i+1}^B], \overline{\text{TX}_{\text{SP},i+1}})$, set $\gamma.\text{flag} := 2$ and $\gamma.st' := \vec{\theta}$ and send $(\text{UPDATE} - \text{OK}, id) \xrightarrow{\tau_1} \mathcal{F}$ on behalf of B. Else stop.
3. If A sends $(\text{SETUP} - \text{OK}, id) \xrightarrow{\tau_1+1} \mathcal{F}$, compute $\sigma_{\text{TX}_{\text{CM},i+1}}^A := \text{Sign}_{sk_{\text{SP}}^A}(f([\text{TX}_{\text{CM},i+1}^B]))$ and send $(\text{updateComA}, id, \hat{\sigma}_{\text{TX}_{\text{SP},i+1}}^A, \sigma_{\text{TX}_{\text{CM},i+1}}^A) \xrightarrow{\tau_1+1} B$. Else, execute the procedure $\text{ForceClose}^A(id)$ and stop.
4. If B sends $(\text{updateComB}, id, \sigma_{\text{TX}_{\text{CM},i+1}}^B) \xrightarrow{\tau_1+2} A$, s.t. $\text{Vrfy}_{pk_{\text{CM}}^B}(f([\text{TX}_{\text{CM},i+1}^A]); \sigma_{\text{TX}_{\text{CM},i+1}}^B) = 1$, compute $\sigma_{\text{TX}_{\text{CM},i+1}}^A = \text{Sign}_{sk_{\text{SP}}^A}(f([\text{TX}_{\text{CM},i+1}^A]))$, set $\text{TX}_{\text{CM},i+1}^A := (\mathcal{H}([\text{TX}_{\text{CM},i+1}^A]), [\text{TX}_{\text{CM},i+1}^A], (1, \{\sigma_{\text{TX}_{\text{CM},i+1}}^A, \sigma_{\text{TX}_{\text{CM},i+1}}^B\}))$, store $\Gamma^A(id)[1] := \text{TX}_{\text{CM},i+1}^A$, and send $(\text{SETUP}' - \text{OK}, id) \xrightarrow{\tau_1+2} \mathcal{F}$ on behalf of B. Else, execute the procedure $\text{ForceClose}^A(id)$ and stop.
5. If A sends $(\text{REVOKE}, id) \xrightarrow{\tau_1+3} \mathcal{F}$, then create $(\overline{[\text{TX}_{\text{RV},i}^A]}, \overline{[\text{TX}_{\text{RV},i}^B]}]) := \text{GenRevoke}(pk_A, pk_B, \gamma.\text{cash}, i+1)$, compute $\hat{\sigma}_{\text{TX}_{\text{RV},i}}^A = \text{Sign}_{sk_{\text{RV}}^A}(\hat{f}([\text{TX}_{\text{RV},i}^B]))$ and send

- (revokeA, $id, \hat{\sigma}_{\text{TX}_{\text{RV},i}}^A$) $\xrightarrow{\tau_1+3}$ B . Else, execute the procedure $\text{ForceClose}^A(id)$ and stop.
6. If B sends (revokeB, $id, \hat{\sigma}_{\text{TX}_{\text{RV},i}}^B$) $\xrightarrow{\tau_1+4}$ A , s.t. $\text{Vrfy}_{pk_{\text{RV}}^B}(\hat{f}([\text{TX}_{\text{RV},i+1}^A]); \hat{\sigma}_{\text{TX}_{\text{RV},i}}^B) = 1$, then set $\Theta^A(id) := (\hat{\sigma}_{\text{TX}_{\text{RV},i}}^B)$, $\gamma.sn := i + 1$, $\gamma.st := \vec{\theta}$, $\Gamma^A(id) := (\gamma, \text{TX}_{\text{FU}}^A, \text{TX}_{\text{CM},i+1}^A, [\text{TX}_{\text{CM},i+1}^B], \overline{\text{TX}_{\text{SP},i+1}})$, $\gamma.flag = 1$, and $\Gamma'^A(id) = (\perp, \perp, \perp)$ and send (REVOKE', id) $\xrightarrow{\tau_1+4}$ $\mathcal{F}$ on behalf of B . Else, execute $\text{ForceClose}^A(id)$ and stop.

Case B is honest and A is corrupted.

Upon A sending ($updateReq, id, \vec{\theta}, t_{stp}$) $\xrightarrow{\tau_0}$ B , proceed as follows:

1. Send (UPDATE, $id, \vec{\theta}, t_{stp}$) $\xrightarrow{\tau_0}$ $\mathcal{F}$ on behalf of A .
2. If B sends (UPDATE – OK, id) $\xrightarrow{\tau_1 \leq \tau_0 + 1 + t_{stp}}$ $\mathcal{F}$, extract TX_{FU} and $i := \gamma.sn$ from $\Gamma^B(id)$, create $([\text{TX}_{\text{CM},i+1}^A], [\text{TX}_{\text{CM},i+1}^B]) := \text{GenCommit}([\text{TX}_{\text{FU}}].\text{txid} \| 1, I_A, I_B, i + 1)$, and $[\text{TX}_{\text{SP},i+1}] := \text{GenSplit}(\gamma.st, i + 1)$ for $I_A := (pk_{\text{SP}}^A, pk_{\text{RV}}^A, pk_{\text{RV}}'^A)$ and $I_B := (pk_{\text{SP}}^B, pk_{\text{RV}}^B, pk_{\text{RV}}'^B)$, compute $\hat{\sigma}_{\text{TX}_{\text{SP},i+1}}^B = \text{Sign}_{sk_{\text{SP}}^B}(\hat{f}([\text{TX}_{\text{SP},i+1}]))$ and send (updateInfo, $id, \hat{\sigma}_{\text{TX}_{\text{SP},i+1}}^B$) $\xrightarrow{\tau_1}$ A .
3. If A sends (updateComA, $id, \hat{\sigma}_{\text{TX}_{\text{SP},i+1}}^A, \sigma_{\text{TX}_{\text{CM},i+1}}^A$) $\xrightarrow{\tau_1+1}$ B , such that $\text{Vrfy}_{pk_{\text{SP}}^A}(\hat{f}([\text{TX}_{\text{SP},i+1}]); \hat{\sigma}_{\text{TX}_{\text{SP},i+1}}^A) = 1$ and $\text{Vrfy}_{pk_{\text{CM}}^A}(\hat{f}([\text{TX}_{\text{CM},i+1}^B]); \sigma_{\text{TX}_{\text{CM},i+1}}^A) = 1$, compute $\sigma_{\text{TX}_{\text{CM},i+1}}^B = \text{Sign}_{sk_{\text{CM}}^B}(\hat{f}([\text{TX}_{\text{CM},i+1}^B]))$, set $\overline{\text{TX}_{\text{SP},i+1}} = ([\text{TX}_{\text{SP},i+1}], (1, \{\hat{\sigma}_{\text{TX}_{\text{SP},i+1}}^A, \hat{\sigma}_{\text{TX}_{\text{SP},i+1}}^B\}))$, $\text{TX}_{\text{CM},i+1}^B = ([\text{TX}_{\text{CM},i+1}^B], (1, \{\sigma_{\text{TX}_{\text{CM},i+1}}^A, \sigma_{\text{TX}_{\text{CM},i+1}}^B\}))$, $\gamma.flag = 2$ and $\gamma.st' = \vec{\theta}$, store $\Gamma'^B(id) = (\text{TX}_{\text{CM},i+1}^B, [\text{TX}_{\text{CM},i+1}^A], \overline{\text{TX}_{\text{SP},i+1}})$, set $\gamma.flag = 2$ and $\gamma.st' = \vec{\theta}$ and send (SETUP – OK, id) $\xrightarrow{\tau_1+1}$ $\mathcal{F}$ on behalf of A . Else execute the procedure $\text{ForceClose}^B(id)$ and stop.
4. If (SETUP' – OK, id) $\xleftarrow{\tau_1+2}$ B , compute $\sigma_{\text{TX}_{\text{CM},i+1}}^B = \text{Sign}_{sk_{\text{SP}}^B}(\hat{f}([\text{TX}_{\text{CM},i+1}^A]))$ and send (updateComB, $id, \sigma_{\text{TX}_{\text{CM},i+1}}^B$) $\xrightarrow{\tau_1+2}$ A . Else, execute $\text{ForceClose}^B(id)$ and stop.
5. Create $([\text{TX}_{\text{RV},i}^A], [\text{TX}_{\text{RV},i}^B]) := \text{GenRevoke}(pk_A, pk_B, \gamma.cash, i + 1)$. If A sends (revokeA, $id, \hat{\sigma}_{\text{TX}_{\text{RV},i}}^A$) $\xrightarrow{\tau_1+3}$ B , s.t. $\text{Vrfy}_{pk_{\text{RV}}^A}(\hat{f}([\text{TX}_{\text{RV},i+1}^B]); \hat{\sigma}_{\text{TX}_{\text{RV},i}}^A) = 1$, set $\Theta^B(id) := (\hat{\sigma}_{\text{TX}_{\text{RV},i}}^A)$, $\gamma.sn := i + 1$, $\gamma.st := \vec{\theta}$, $\Gamma^B(id) := (\gamma, \text{TX}_{\text{FU}}, \text{TX}_{\text{CM},i+1}^B,$

- $[\text{TX}_{\text{CM},i+1}^A], \text{TX}_{\text{SP},i+1}), \gamma.\text{flag} = 1$, and $\Gamma'^B(id) = (\perp, \perp, \perp)$ and send $(\text{REVOKE}, id) \xrightarrow{\tau_1+3} \mathcal{F}$ on behalf of A . Else, execute the procedure $\text{ForceClose}^B(id)$ and stop.
6. If B sends $(\text{REVOKE}', id) \xrightarrow{\tau_1+4} \mathcal{F}$, compute $\hat{\sigma}_{\text{TX}_{\text{RV},i}^A}^B = \text{Sign}_{sk_{\text{RV}}^B}(\hat{f}([\text{TX}_{\text{RV},i}^A], 1))$, send $(\text{revokeB}, id, \hat{\sigma}_{\text{TX}_{\text{RV},i}^A}^B) \xrightarrow{\tau_1+4} A$. Else, execute $\text{ForceClose}^B(id)$ and stop.

Close

Case A is honest and B is corrupted.

1. Upon A sending $(\text{CLOSE}, id) \xrightarrow{\tau_0} \mathcal{F}$, extract $\text{TX}_{\text{FU}}, i := \gamma.sn, \text{TX}_{\text{SP},i}$ from $\Gamma^A(id)$ and create $[\text{TX}_{\text{SP}}] := \text{GenFinSplit}(\text{TX}_{\text{FU}}.\text{txid}\|1, \gamma.st)$.
2. Compute $\sigma_{\text{TX}_{\text{SP}}}^A := \text{Sign}_{sk_{\text{CM}}^A}(f(\overline{[\text{TX}_{\text{SP}}]}))$ and send $(\text{CloseA}, id, \sigma_{\text{TX}_{\text{SP}}}^A) \xrightarrow{\tau_0} B$.
3. If B sends $(\text{CloseB}, id, \sigma_{\text{TX}_{\text{SP}}}^B) \xrightarrow{\tau_0} A$ s.t. $\text{Vrfy}_{pk_{\text{CM}}^B}(f([\text{TX}_{\text{SP}}]); \sigma_{\text{TX}_{\text{SP}}}^B) = 1$, then send $(\text{CLOSE}, id) \xrightarrow{\tau_0} \mathcal{F}$ on behalf of B . Otherwise, execute the simulator code of the procedure $\text{ForceClose}^A(id)$ and stop.
4. Create $\text{TX}_{\text{SP}} := ([\text{TX}_{\text{SP}}], (1, \{\sigma_{\text{TX}_{\text{SP}}}^A, \sigma_{\text{TX}_{\text{SP}}}^B\}))$ and send $(\text{post}, \text{TX}_{\text{SP}}) \xrightarrow{\tau_0+1} \mathcal{L}$.
5. If $\tau_1 \leq \tau_0 + 1 + \Delta$ is the round in which TX_{SP} is accepted by $\mathcal{L}$, set $\Gamma^A(id) = \perp$, $\Theta^A(id) = \perp$.

Punish

Case A is honest and B is corrupted.

For each $id \in \{0, 1\}^*$, extract $i := \gamma.sn$ and $flag := \gamma.flag$ from $\Gamma^A(id)$.

If $flag = 1$:

- Parse $(\gamma, \text{TX}_{\text{FU}}, \text{TX}_{\text{CM},i}^A, [\text{TX}_{\text{CM},i}^B], \overline{[\text{TX}_{\text{SP},i}]}) := \Gamma^A(id)$ and set $I := \{[\text{TX}_{\text{CM},i}^A], [\text{TX}_{\text{CM},i}^B]\}$. Check if $\text{TX}_{\text{FU}}.\text{Output}$ is spent by a transaction TX s.t. $[\text{TX}] \notin I$ and $\text{TX}.\text{Output} \neq \gamma.st$. If yes:

1. Create $(\overline{[TX_{RV,i-1}^A]}, \overline{[TX_{RV,i-1}^B]}) := \text{GenRevoke}(pk_{SP}^A, pk_{SP}^B, \gamma.\text{cash}, i)$ and then set $[TX_{RV,i-1}^A] := (\text{TX.txid}\|1, \overline{[TX_{RV,i-1}^A]})$.
2. compute $\hat{\sigma}_{TX_{RV,i-1}^A}^A = \text{Sign}_{sk_{RV}^A}(\hat{f}(\overline{[TX_{RV,i-1}^A]}))$ (for the case where B is honest and A is corrupted, sk_{RV}^B is used to compute the signature).
3. Set $\hat{\sigma}_{TX_{RV,i-1}^A}^B := \Theta^A(id)$, and create $TX_{RV,i-1}^A := (\mathcal{H}(\overline{[TX_{RV,i-1}^A]}), [TX_{RV,i-1}^A], (1, \{\hat{\sigma}_{TX_{RV,i-1}^A}^A, \hat{\sigma}_{TX_{RV,i-1}^A}^B\}))$.
4. Post $(\text{post}, TX_{RV,i-1}^A) \xrightarrow{\tau_0} \mathcal{L}$.
5. Let $TX_{RV,i-1}^A$ be accepted by $\mathcal{L}$ in round $\tau_1 \leq \tau_0 + \Delta$. Set $\Theta^A(id) := \perp$, $\Gamma^A(id) := \perp$ and output $(\text{PUNISHED}, id) \xrightarrow{\tau_1} \mathcal{E}$.

Otherwise, if $\text{TX.Output} = \gamma.\text{st}$, then set $\Theta^A(id) := \perp$, $\Gamma^A(id) := \perp$ and output $(\text{CLOSED}, id) \xrightarrow{\tau_0} \mathcal{E}$. Else, set $[TX_{SP,i}] := (\text{TX.txid}\|1, \overline{TX_{SP,i}.\text{nLT}}, \overline{TX_{SP,i}.\text{Output}})$, $TX_{SP,i} := (\mathcal{H}(\overline{[TX_{SP,i}]})$, $\text{TX.txid}\|1, \overline{TX_{SP,i}})$ and post $(\text{post}, TX_{SP,i}) \xrightarrow{\tau_0+t} \mathcal{L}$. Let $TX_{SP,i}$ be accepted by $\mathcal{L}$ in round $\tau_1 \leq \tau_0 + t + \Delta$. Set $\Theta^A(id) := \perp$, $\Gamma^A(id) := \perp$ and output $(\text{CLOSED}, id) \xrightarrow{\tau_1} \mathcal{E}$.

If $flag = 2$:

- Parse $(\gamma, TX_{FU}, TX_{CM,i}^A, [TX_{CM,i}^B], \overline{TX_{SP,i}}) := \Gamma^A(id)$ as well as $(TX_{CM,i+1}^A, [TX_{CM,i+1}^B], \overline{TX_{SP,i+1}}) := \Gamma^A(id)$ and set $I = \{[TX_{CM,i}^A], [TX_{CM,i}^B], [TX_{CM,i+1}^A], [TX_{CM,i+1}^B]\}$. Check if $TX_{FU}.\text{Output}$ is spent by a transaction TX s.t. $[TX] \notin I$, $TX.\text{Output} \neq \gamma.\text{st}$, and $TX.\text{Output} \neq \gamma.\text{st}'$. If yes:

1. Create $(\overline{[TX_{RV,i-1}^A]}, \overline{[TX_{RV,i-1}^B]}) := \text{GenRevoke}(pk_A, pk_B, \gamma.\text{cash}, i)$ and then set $[TX_{RV,i-1}^A] := (\text{TX.txid}\|1, \overline{[TX_{RV,i-1}^A]})$.
2. compute $\hat{\sigma}_{TX_{RV,i-1}^A}^A := \text{Sign}_{sk_{RV}^A}(\hat{f}(\overline{[TX_{RV,i-1}^A]}))$ (for the case where B is honest and A is corrupted, sk_{RV}^B is used to compute the signature).
3. Set $\hat{\sigma}_{TX_{RV,i-1}^A}^B := \Theta^A(id)$, and create $TX_{RV,i-1}^A := (\mathcal{H}(\overline{[TX_{RV,i-1}^A]}), [TX_{RV,i-1}^A], (2, \{\hat{\sigma}_{TX_{RV,i-1}^A}^A, \hat{\sigma}_{TX_{RV,i-1}^A}^B\}))$.
4. Post $(\text{post}, TX_{RV,i-1}^A) \xrightarrow{\tau_0} \mathcal{L}$.
5. Let $TX_{RV,i-1}^A$ be accepted by $\mathcal{L}$ in round $\tau_1 \leq \tau_0 + \Delta$. Set $\Theta^A(id) := \perp$, $\Gamma^A(id) := \perp$ and output $(\text{PUNISHED}, id) \xrightarrow{\tau_1} \mathcal{E}$.

Otherwise, if $TX.\text{Output} = \gamma.\text{st}$ or $TX.\text{Output} = \gamma.\text{st}'$ hold, then set $\Theta^A(id) := \perp$, $\Gamma^A(id) := \perp$ and output $(\text{CLOSED}, id) \xrightarrow{\tau_0} \mathcal{E}$. Else, set $[TX_{SP,i+1}] = (\text{TX.txid}\|1,$

$\overline{\text{TX}_{\text{SP},i+1}.\text{nLT}}, \overline{\text{TX}_{\text{SP},i+1}.\text{Output}}, \text{TX}_{\text{SP},i+1} := (\mathcal{H}([\text{TX}_{\text{SP},i+1}]), \text{TX.txid}\|1, \overline{\text{TX}_{\text{SP},i+1}})$ and $\text{post}(\text{post}, \text{TX}_{\text{SP},i+1}) \xrightarrow{\tau_0+t} \mathcal{L}$. Let $\text{TX}_{\text{SP},i+1}$ be accepted by $\mathcal{L}$ in round $\tau_1 \leq \tau_0 + t + \Delta$. Then, set $\Theta^P(id) := \perp$, $\Gamma^A(id) := \perp$ and $(\text{CLOSED}, id) \xrightarrow{\tau_1} \mathcal{E}$.

Subprocedure ForceClose^P(id)

Let τ_0 be the current round. Extract $i := \gamma.\text{sn}$, $\text{flag} := \gamma.\text{flag}$, $\text{TX}_{\text{CM},i}^P$ and $\overline{\text{TX}_{\text{SP},i}}$ from $\Gamma^P(id)$ and $\text{TX}_{\text{CM},i+1}^P$ and $\overline{\text{TX}_{\text{SP},i+1}}$ from $\Gamma'^P(id)$.

If $\text{flag} = 1$:

1. $\text{Post}(\text{post}, \text{TX}_{\text{CM},i}^P) \xrightarrow{\tau_0} \mathcal{L}$.
2. Let $\tau_1 \leq \tau_0 + \Delta$ be the round in which $\text{TX}_{\text{CM},i}^P$ is accepted by the blockchain. Wait for t rounds, set $[\text{TX}_{\text{SP},i}] = (\text{TX}_{\text{CM},i}^P.\text{txid}\|1, \overline{\text{TX}_{\text{SP},i}.\text{nLT}}, \overline{\text{TX}_{\text{SP},i}.\text{Output}})$, $\text{TX}_{\text{SP},i} := (\mathcal{H}([\text{TX}_{\text{SP},i}]), \text{TX}_{\text{CM},i}^P.\text{txid}\|1, \overline{\text{TX}_{\text{SP},i}})$ and $\text{post}(\text{post}, \text{TX}_{\text{SP},i}) \xrightarrow{\tau_2:=\tau_1+t} \mathcal{L}$.
3. Once $\text{TX}_{\text{SP},i}$ is accepted by $\mathcal{L}$ in round $\tau_3 \leq \tau_2 + \Delta$ set $\Theta^P(id) := \perp$ and $\Gamma^P(id) := \perp$ and output $(\text{CLOSED}, id) \xrightarrow{\tau_3} \gamma.\text{users}$.

Otherwise, extract $\text{TX}_{\text{CM},i+1}^P$ and $\overline{\text{TX}_{\text{SP},i+1}}$ from $\Gamma'^P(id)$:

1. If $\text{TX}_{\text{CM},i+1}^P = \perp$, Send $(\text{post}, \text{TX}_{\text{CM},i}^P) \xrightarrow{\tau_0} \mathcal{L}$. Else, Send $(\text{post}, \text{TX}_{\text{CM},i+1}^P) \xrightarrow{\tau_0} \mathcal{L}$.
2. Let $\tau_1 \leq \tau_0 + \Delta$ be the round in which either $\text{TX}_{\text{CM},i}^P$ or $\text{TX}_{\text{CM},i+1}^P$ is accepted by the blockchain. Wait for t rounds, set $[\text{TX}_{\text{SP},i+1}] = (\text{TX.txid}\|1, \overline{\text{TX}_{\text{SP},i+1}.\text{nLT}}, \overline{\text{TX}_{\text{SP},i+1}.\text{Output}})$, $\text{TX}_{\text{SP},i+1} := (\mathcal{H}([\text{TX}_{\text{SP},i+1}]), \text{TX.txid}\|1, \overline{\text{TX}_{\text{SP},i+1}})$ and then $\text{post}(\text{post}, \text{TX}_{\text{SP},i+1}) \xrightarrow{\tau_2:=\tau_1+t} \mathcal{L}$.
3. Once $\text{TX}_{\text{SP},i+1}$ is accepted by $\mathcal{L}$ in round $\tau_3 \leq \tau_2 + \Delta$ set $\Theta^P(id) := \perp$ and $\Gamma^P(id) := \perp$ and output $(\text{CLOSED}, id) \xrightarrow{\tau_3} \gamma.\text{users}$.

Lemma 7.1. *Let Σ be a secure signature scheme. Then, the Create phase of protocol π GUC-emulates the Create phase of functionality $\mathcal{F}$.*

Proof. We define the following message.

- $m_0 := (\text{createCom}, id, \hat{\sigma}_{\text{TX}_{\text{SP},0}}^B, \sigma_{\text{TX}_{\text{CM},0}}^B)$,

The proof is composed of multiple hybrids, where we gradually modify the initial experiment.

Hybrid $\mathcal{H}_0^{Cr}$: This corresponds to the Create phase of protocol π .

Hybrid $\mathcal{H}_1^{Cr}$: For the honest party A in hybrid $\mathcal{H}_0^{Cr}$, if the corrupted party B publishes the funding transaction TX_{FU} on the ledger $\mathcal{L}$ without sending the message m_0 to A in round $\tau_0 + 1$, then the experiment outputs Error and fails.

Simulator $\mathcal{S}^{Cr}$: This corresponds to the Create phase of the simulator, as defined in beginning of this section.

Lemma 7.2 proves the indistinguishability of the neighbouring experiments Hybrid $\mathcal{H}_0^{Cr}$ and Hybrid $\mathcal{H}_1^{Cr}$. Lemma 7.3 also proves the indistinguishability of the neighbouring experiments Hybrid $\mathcal{H}_1^{Cr}$ and Simulator $\mathcal{S}^{Cr}$. This concludes the proof of Lemma 7.1. $\square$

Lemma 7.2. *For all PPT distinguishers $\mathcal{E}$ it holds that*

$$\begin{aligned} & \{\text{EXE}_{\mathcal{H}_0^{Cr}, \mathcal{A}, \mathcal{E}}^{\mathcal{L}(\Delta, \Sigma), \mathcal{F}_{\text{clock}}}(\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*} \approx \\ & \{\text{EXE}_{\mathcal{H}_1^{Cr}, \mathcal{A}, \mathcal{E}}^{\mathcal{L}(\Delta, \Sigma), \mathcal{F}_{\text{clock}}}(\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*} \end{aligned}$$

Proof. In addition to the message m_0 defined in Lemma 7.1, we define m_1 as follows:

- $m_1 := (\text{createFund}, id, \sigma_{\text{TX}_{\text{FU}}}^A)$,

Two hybrids $\mathcal{H}_0^{Cr}$ and $\mathcal{H}_1^{Cr}$ differ if the experiment outputs Error. Thus, we must bound the probability that this event occurs. The experiment outputs Error, if and only if the corrupted party B publishes the funding transaction TX_{FU} on the ledger $\mathcal{L}$ without sending the message m_0 to A . Furthermore, according to the protocol π , if A does not receive the message m_0 , he does not send the message m_1 including the signature $\sigma_{\text{TX}_{\text{FU}}}^A$ to B . However, this signature must be part of TX_{FU} . Witness if TX_{FU} is published on $\mathcal{L}$ and hence the signature $\sigma_{\text{TX}_{\text{FU}}}^A$ must be created by the adversary. Thus, given that the experiment outputs Error with non-negligible probability, as will be shown in the next paragraph, we construct a reduction against the existential unforgeability of the underlying signature scheme Σ with non-negligible success probability which contradicts with our assumption regarding the security of Σ .

Assume that $\Pr[\text{Error} \mid \mathcal{H}_0^{Cr}] \geq \frac{1}{\text{poly}(\lambda)}$. The reduction receives as input a public key pk from the challenger and registers it by sending the message $(\text{register}, pk)$ to $\mathcal{L}$. Now assume that the honest party A , upon receiving the message $(\text{INTRO}, \gamma, \text{tid}_A)$ from $\mathcal{E}$,

initiates the Create phase of protocol π . If in this process, m_0 is received from the adversary, the hybrid $\mathcal{H}_1^{Cr}$ does not output Error and the reduction aborts. If the experiment outputs Error, meaning that for this channel the corresponding funding transaction TX_{FU} is accepted by $\mathcal{L}$, then the reduction outputs (m^*, σ^*) with $m^* = [\text{TX}_{\text{FU}}]$ and $\sigma^* \in \text{TX}_{\text{FU}}.\text{Witness}$. This reduction is clearly efficient, and whenever $\mathcal{H}_1^{Cr}$ outputs Error, the reduction succeeds in forging the signature. Moreover, the reduction has never called the signing oracle for any messages. Therefore, the reduction outputs a valid forgery with probability at least $\frac{1}{\text{poly}(\lambda)}$, which contradicts with our assumption regarding the security of the signature scheme Σ . This proves that $\Pr[\text{Error} \mid \mathcal{H}_0^{Cr}] < \frac{1}{\text{poly}(\lambda)}$. $\square$

Lemma 7.3. *For all PPT distinguishers $\mathcal{G}$ it holds that*

$$\begin{aligned} \{\text{EXE}_{\mathcal{H}_1^{Cr}, \mathcal{A}, \mathcal{G}}^{\mathcal{L}(\Delta, \Sigma), \mathcal{F}_{\text{clock}}}(\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*} &\approx \\ \{\text{EXE}_{\varphi_{\mathcal{F}}, \mathcal{S}^{Cr}, \mathcal{G}}^{\mathcal{L}(\Delta, \Sigma), \mathcal{F}_{\text{clock}}}(\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*}. \end{aligned}$$

Proof. The two experiments are identical, and hence, indistinguishability follows. $\square$

Lemma 7.4. *The Update phase of protocol π GUC-emulates the Update phase of functionality $\mathcal{F}$.*

Proof. The two experiments are identical, and hence, indistinguishability follows. $\square$

Lemma 7.5. *The Close phase of protocol π GUC-emulates the Close phase of functionality $\mathcal{F}$.*

Proof. The proof is composed of multiple hybrids, where we gradually modify the initial experiment.

Hybrid $\mathcal{H}_0^{Cl}$: This corresponds to the Close phase of protocol π .

Hybrid $\mathcal{H}_1^{Cl}$: If the honest party A initiates hybrid $\mathcal{H}_0^{Cl}$ in round τ_0 and in round $\tau_0 + 1 + \Delta$, the output $\text{TX}_{\text{FU}}.\text{Output}$ is still unspent, then the experiment outputs Error and fails.

Simulator $\mathcal{S}^{Cl}$: This corresponds to the Close phase of the simulator, as defined at the beginning of this section.

Lemma 7.6 proves the indistinguishability of the neighbouring experiments Hybrid $\mathcal{H}_0^{Cl}$ and Hybrid $\mathcal{H}_1^{Cl}$. Lemma 7.7 also proves the indistinguishability of the neighbouring experiments Hybrid $\mathcal{H}_1^{Cl}$ and Simulator $\mathcal{S}^{Cl}$. This concludes the proof of Lemma 7.5. $\square$

Lemma 7.6. *For all PPT distinguishers $\mathcal{E}$ it holds that*

$$\begin{aligned} \{\text{EXE}_{\mathcal{H}_0^{Cl}, \mathcal{A}, \mathcal{E}}^{\mathcal{L}(\Delta, \Sigma), \mathcal{F}_{clock}}(\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*} &\approx \\ \{\text{EXE}_{\mathcal{H}_1^{Cl}, \mathcal{A}, \mathcal{E}}^{\mathcal{L}(\Delta, \Sigma), \mathcal{F}_{clock}}(\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*}. \end{aligned}$$

Proof. Similar to the proof of Lemma 7.2, we must bound the probability that $\mathcal{H}_1^{Cl}$ outputs Error. According to $\mathcal{H}_1^{Cl}$, the honest party A sends either $(\text{post}, \text{TX}_{\overline{\text{SP}}}) \xrightarrow{\tau_0+1} \mathcal{L}$ or $(\text{post}, \text{TX}_{\text{CM},i}^A) \xrightarrow{\tau_0+1} \mathcal{L}$. Since both $\text{TX}_{\overline{\text{SP}}}$ and $\text{TX}_{\text{CM},i}^A$ are valid and both take output of TX_{FU} as their input, based on the ledger functionality $\mathcal{L}$, TX_{FU} .Output becomes spent within at most Δ rounds. Therefore, $\mathcal{H}_1^{Cl}$ will never output Error. This completes the proof. $\square$

Lemma 7.7. *For all PPT distinguishers $\mathcal{E}$ it holds that*

$$\begin{aligned} \{\text{EXE}_{\mathcal{H}_1^{Cl}, \mathcal{A}, \mathcal{E}}^{\mathcal{L}(\Delta, \Sigma), \mathcal{F}_{clock}}(\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*} &\approx \\ \{\text{EXE}_{\phi_{\mathcal{F}}, \mathcal{S}^{Cl}, \mathcal{E}}^{\mathcal{L}(\Delta, \Sigma), \mathcal{F}_{clock}}(\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*}. \end{aligned}$$

Proof. The two experiments are identical, and hence, indistinguishability follows. $\square$

Lemma 7.8. *Let Σ be a secure signature scheme. Then, the Punish phase of protocol π GUC-emulates the Punish phase of functionality $\mathcal{F}$.*

Proof. The proof is composed of multiple hybrids, where we gradually modify the initial experiment.

- Hybrid $\mathcal{H}_0^{Pu}$: This corresponds to the Punish phase of the protocol π .
- Hybrid $\mathcal{H}_1^{Pu}$: For the honest party A in hybrid $\mathcal{H}_0^{Pu}$, if a transaction TX is published s.t. $[\text{TX}] = [\text{TX}_{\text{CM},j}^B]$ with $j \in [0, i-1]$ but $\text{TX}_{\text{RV},i-1}^A$ is not accepted by $\mathcal{L}$ within Δ rounds, then the experiment outputs Error and fails.
- Hybrid $\mathcal{H}_2^{Pu}$: For the honest party A in hybrid $\mathcal{H}_1^{Pu}$, if a transaction TX is published s.t. $[\text{TX}] = [\text{TX}_{\text{CM},i}^B]$ given that $\gamma.\text{flag} = 1$ or $[\text{TX}] \in \{[\text{TX}_{\text{CM},i}^B], [\text{TX}_{\text{CM},i+1}^B]\}$ given that $\gamma.\text{flag} = 2$, and then a transaction TX' is published s.t. $\text{TX}'.\text{Input} = \text{TX}.\text{txid} \parallel 1$ and $\text{TX}'.\text{Witness}.\eta = 2$ (i.e. $\text{TX}'.\text{Witness}$ satisfies the second sub-condition of $\text{TX}.\text{Output}$), then the experiment outputs Error and fails.
- Hybrid $\mathcal{H}_3^{Pu}$: For the honest party A in hybrid $\mathcal{H}_2^{Pu}$, if either of the following cases occurs, the experiment outputs Error and fails.

- While $\gamma.\text{flag} = 1$, a transaction TX is published s.t. $[\text{TX}] = [\text{TX}_{\text{CM},i}^A]$, and then a transaction TX' is published s.t. $\text{TX}.\text{Input} = \text{TX}.\text{txid}\|1$ and $\text{TX}.\text{Witness}.\eta = 2$ (i.e. $\text{TX}.\text{Witness}$ satisfies the second sub-condition of $\text{TX}.\text{Output}$).
- While $\gamma.\text{flag} = 2$, a transaction TX is published s.t. $[\text{TX}] = [\text{TX}_{\text{CM},i}^A]$ given that $\Gamma^A(\text{id})[1] = \perp$ or $[\text{TX}] = [\text{TX}_{\text{CM},i+1}^A]$ otherwise, and then a transaction TX' is published s.t. $\text{TX}.\text{Input} = \text{TX}.\text{txid}\|1$ and $\text{TX}.\text{Witness}.\eta = 2$ (i.e. $\text{TX}.\text{Witness}$ satisfies the second sub-condition of $\text{TX}.\text{Output}$).
- Hybrid $\mathcal{H}_4^{\text{Pu}}$: For the honest party A in hybrid $\mathcal{H}_3^{\text{Pu}}$, if either of the following cases occurs, the experiment outputs Error and fails.
 - While $\gamma.\text{flag} = 1$, a transaction TX with $[\text{TX}] \in \{[\text{TX}_{\text{CM},i}^A], [\text{TX}_{\text{CM},i}^B]\}$ is published, but $\text{TX}_{\text{SP},i}$ is not accepted by $\mathcal{L}$ within $t + \Delta$ rounds.
 - While $\gamma.\text{flag} = 2$, a transaction TX with $[\text{TX}] \in \{[\text{TX}_{\text{CM},i}^A], [\text{TX}_{\text{CM},i}^B], [\text{TX}_{\text{CM},i+1}^B]\}$ is published, but $\text{TX}_{\text{SP},i}$ or $\text{TX}_{\text{SP},i+1}$ is not accepted by $\mathcal{L}$ within $t + \Delta$ rounds.
- Hybrid $\mathcal{H}_5^{\text{Pu}}$: For the honest party A in hybrid $\mathcal{H}_4^{\text{Pu}}$, if either of the following cases occurs, the experiment outputs Error and fails.
 - While $\gamma.\text{flag} = 1$, a transaction TX with $\text{TX}.\text{Input} = \text{TX}_{\text{FU}}.\text{txid}\|1$ is published s.t. $[\text{TX}] \notin \{[\text{TX}_{\text{CM},i}^A], [\text{TX}_{\text{CM},j}^B], j = [0, i]\}$ and $\text{TX}.\text{Output} \neq \gamma.\text{st}$.
 - While $\gamma.\text{flag} = 2$, a transaction TX with $\text{TX}.\text{Input} = \text{TX}_{\text{FU}}.\text{txid}\|1$ is published s.t. $[\text{TX}] \notin \{[\text{TX}_{\text{CM},i}^A], [\text{TX}_{\text{CM},i+1}^A], [\text{TX}_{\text{CM},j}^B], j = [0, i + 1]\}$ and $\text{TX}.\text{Output} \neq \gamma.\text{st}$.
- Simulator $\mathcal{S}^{\text{Pu}}$: This corresponds to the Punish phase of the simulator, as defined at the beginning of this section.

Lemmas 7.9 to 7.14 prove the indistinguishability of the above neighbouring experiments. This concludes the proof of Lemma 7.8. $\square$

Lemma 7.9. *For all PPT distinguishers $\mathcal{G}$ it holds that*

$$\begin{aligned} & \{\text{EXE}_{\mathcal{H}_0^{\text{Pu}}, \mathcal{A}, \mathcal{G}}^{\mathcal{L}(\Delta, \Sigma), \mathcal{F}^{\text{clock}}}(\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*} \approx \\ & \{\text{EXE}_{\mathcal{H}_1^{\text{Pu}}, \mathcal{A}, \mathcal{G}}^{\mathcal{L}(\Delta, \Sigma), \mathcal{F}^{\text{clock}}}(\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*}. \end{aligned}$$

Proof. Two hybrids differ if the experiment outputs Error. Thus, we must bound the probability that this event occurs. The hybrid $\mathcal{H}_1^{\text{Pu}}$ does not output Error unless $\text{TX}_{\text{FU}}.\text{Output}$ is spent by a transaction TX s.t. $[\text{TX}] = [\text{TX}_{\text{CM},j}^B]$ with $j \in [0, i - 1]$. According to $\mathcal{H}_1^{\text{Pu}}$, once this transaction is observed by A at the end of round τ_0 , A posts $(\text{post}, \text{TX}_{\text{RV},i-1}^A) \xrightarrow{\tau_0} \mathcal{L}$. Since $\text{TX}_{\text{RV},i-1}^A$ with $\text{TX}_{\text{RV},i-1}^A.\text{Input} = \text{TX}.\text{txid}\|1$ is a valid transaction, it is accepted by $\mathcal{L}$ within Δ rounds unless another valid transaction

TX' is published within this Δ -round interval with $\text{TX}'.\text{Input} = \text{TX}.\text{txid}\|1$. The output $\text{TX}.\text{Output}$ has two sub-conditions, one of which must be satisfied by TX' . The first sub-condition $\varphi_1 = pk_{\text{SP}}^A \wedge pk_{\text{SP}}^B \wedge \text{CLTV}_{S_0+j} \wedge \text{CSV}_t$ cannot be satisfied within t rounds and since we have $t > \Delta$, φ_1 cannot be met within Δ rounds. Satisfying the second sub-condition $\varphi_2 = pk_{\text{RV}}'^A \wedge pk_{\text{RV}}'^B \wedge \text{CLTV}_{S_0+j}$ requires A 's signature and according to the protocol π , A does not grant such an authorisation to anyone. Thus, if the experiment outputs **Error** with non-negligible probability, we construct a reduction against the existential unforgeability of the underlying signature scheme Σ with non-negligible success probability which contradicts our assumption regarding the security of Σ .

Assume that $\Pr[\text{Error} \mid \mathcal{H}_0^{\text{Pu}}] \geq \frac{1}{\text{poly}(\lambda)}$. The reduction receives a public key pk from the challenger as input, and in the channel creation phase sets $pk_{\text{RV}}'^A := pk$. The channel might be updated any arbitrary number of times and might be closed at any time using any method (peacefully or forcefully) selected by the adversary. Assume that the channel has been updated i times. If the output of the funding transaction is spent by a transaction TX s.t. $[\text{TX}] \neq [\text{TX}_{\text{CM},j}^B]$ with $j = [0, i-1]$, the hybrid $\mathcal{H}_1^{\text{Pu}}$ does not output **Error** and hence the reduction aborts. If a transaction TX with $[\text{TX}] \in [\text{TX}_{\text{CM},j}^B]$ with $j = [0, i-1]$ is published, the reduction calls the signing oracle for the message $[\text{TX}_{\text{RV},i-1}^A]$, creates $\text{TX}_{\text{RV},i-1}^A$ and posts it to the ledger $\mathcal{L}$. If $\text{TX}_{\text{RV},i-1}^A$ is published on $\mathcal{L}$ within Δ rounds, the hybrid $\mathcal{H}_1^{\text{Pu}}$ does not output **Error** and hence the reduction aborts. Otherwise, since $\text{TX}_{\text{RV},i-1}^A$ is a valid transaction, based on our assumptions on $\mathcal{L}$, another transaction TX' with $[\text{TX}'] \neq [\text{TX}_{\text{RV},i-1}^A]$ and $\text{TX}'.\text{Input} = \text{TX}.\text{txid}\|1$ appears on the ledger within this Δ round interval. This causes $\mathcal{H}_1^{\text{Pu}}$ to output **Error**. As mentioned earlier, $\text{TX}'.\text{Witness}$ satisfies the condition $pk_{\text{RV}}'^A \wedge pk_{\text{RV}}'^B \wedge \text{CLTV}_{S_0+j}$ of the output $\text{TX}.\text{Output}$. Now, the reduction outputs (m^*, σ^*) with $m^* = [\text{TX}']$ and $\sigma^* \in \text{TX}'.\text{Witness}.\zeta$. Moreover, the reduction has never called the signing oracle for m^* before, because the signing oracle was called only once for $[\text{TX}_{\text{RV},i-1}^A] \neq m^*$. Therefore, the reduction outputs a valid forgery with probability at least $\frac{1}{\text{poly}(\lambda)}$, which contradicts our assumption regarding the security of Σ . This contradiction proves that $\Pr[\text{Error} \mid \mathcal{H}_0^{\text{Pu}}] < \frac{1}{\text{poly}(\lambda)}$.

□

Lemma 7.10. *For all PPT distinguishers $\mathcal{E}$ it holds that*

$$\begin{aligned} \{\text{EXE}_{\mathcal{H}_1^{\text{Pu}}, \mathcal{A}, \mathcal{E}}^{\mathcal{L}(\Delta, \Sigma), \mathcal{F}^{\text{clock}}}(\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*} &\approx \\ \{\text{EXE}_{\mathcal{H}_2^{\text{Pu}}, \mathcal{A}, \mathcal{E}}^{\mathcal{L}(\Delta, \Sigma), \mathcal{F}^{\text{clock}}}(\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*}. \end{aligned}$$

Proof. Similar to the proof of Lemma 7.9, we show that if the experiment outputs **Error** with non-negligible probability we construct a reduction against the existential unforgeability of the underlying signature scheme Σ with non-negligible success probability. Assume that $\Pr[\text{Error} \mid \mathcal{H}_1^{\text{Pu}}] \geq \frac{1}{\text{poly}(\lambda)}$. The reduction receives as input a public key

pk , and in the channel creation phase sets $pk'_{RV}^A := pk$. The channel is updated any arbitrary number of times and might be closed at any time using any method (peacefully or forcefully) selected by the adversary. Assume that the channel has been updated i times. If the output of the funding transaction is spent by a transaction TX with $[TX] \neq [TX_{CM,i}^B]$ given that $\gamma.\text{flag} = 1$ or $[TX] \notin \{[TX_{CM,i}^B], [TX_{CM,i+1}^B]\}$ given that $\gamma.\text{flag} = 2$, the experiment does not output Error and the reduction aborts. Otherwise, the reduction waits for t rounds and then publishes the latest split transaction. If $TX.\text{Output}$ is spent by a transaction TX' s.t. $TX'.\text{Witness}.\eta = 1$, the experiment does not output Error and the reduction aborts. However, if $TX'.\text{Witness}.\eta = 1$, the experiment outputs Error. Since $TX'.\text{Witness}$ satisfies the condition $pk'_{RV}^A \wedge pk'_{RV}^B \wedge CLTV_{S_0+i}$ of the output $TX.\text{Output}$, reduction outputs (m^*, σ^*) with $m^* = [TX']$ and $\sigma^* \in TX'.\text{Witness}.\zeta$. Moreover, the reduction has never called the signing oracle. Therefore, the reduction outputs a valid forgery with probability at least $\frac{1}{\text{poly}(\lambda)}$, which contradicts our assumption regarding the security of Σ . This contradiction proves that $\Pr[\text{Error} \mid \mathcal{H}_1^{Pu}] < \frac{1}{\text{poly}(\lambda)}$. $\square$

Lemma 7.11. *For all PPT distinguishers $\mathcal{E}$ it holds that*

$$\begin{aligned} & \{\text{EXE}_{\mathcal{H}_2^{Pu}, \mathcal{A}, \mathcal{E}}^{\mathcal{L}(\Delta, \Sigma), \mathcal{F}_{\text{clock}}}(\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*} \approx \\ & \{\text{EXE}_{\mathcal{H}_3^{Pu}, \mathcal{A}, \mathcal{E}}^{\mathcal{L}(\Delta, \Sigma), \mathcal{F}_{\text{clock}}}(\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*}. \end{aligned}$$

Proof. We bound the probability that the experiment outputs Error. Assume that $\Pr[\text{Error} \mid \mathcal{H}_2^{Pu}] \geq \frac{1}{\text{poly}(\lambda)}$. The reduction receives as input a public key pk from the challenger and in the channel creation phase sets $pk'_{RV}^A := pk$. The channel is updated any arbitrary number of times and might be closed at any time using any method (peacefully or forcefully) selected by the adversary. For the j^{th} channel update, to generate the signature on $TX_{RV,j-1}^B$, a call to the signing oracle for the message $[TX_{RV,j-1}^B]$ is performed. We know that the channel has been updated i times. Assume that $TX_{FU}.\text{Output}$ is spent by a transaction TX . If one of the following cases occurs, the experiment does not output Error and the reduction aborts: (1) $\gamma.\text{flag} = 1$ and $[TX] \neq [TX_{CM,i}^A]$, (2) $\gamma.\text{flag} = 2$, and $[TX] \neq [TX_{CM,i}^A]$ given that $\Gamma'^A(id)[1] = \perp$ or $[TX] \neq [TX_{CM,i+1}^A]$ otherwise. Otherwise, the reduction waits for t rounds and then publishes the latest split transaction. If $TX.\text{Output}$ is spent by a transaction TX' s.t. $TX'.\text{Witness}.\eta = 1$, the experiment does not output Error and hence the reduction aborts. If $TX'.\text{Witness}.\eta = 2$, the experiment outputs Error. Now either of the following cases might have occurred:

- We have $\gamma.\text{flag} = 1$ or $\gamma.\text{flag} = 2$ and $\Gamma'^A(id)[1] = \perp$ and hence $[TX] = [TX_{CM,i}^A]$ holds. Also, since $TX'.\text{Witness}.\eta = 2$, $TX'.\text{Witness}.\zeta$ satisfies the condition $pk'_{RV}^A \wedge pk'_{RV}^B \wedge CLTV_{S_0+i}$. The reduction outputs (m^*, σ^*) with $m^* = [TX']$ and $\sigma^* \in TX'.\text{Witness}.\zeta$. Moreover, the signing oracle was only called for messages $[TX_{RV,j}^B]$ with $j < i$ and

since for these transactions we have $\overline{[\text{TX}_{\text{RV},j}^B]}.n\text{LT} = S_0 + j < S_0 + i$, according to $\mathcal{L}$, $\text{TX}_{\text{RV},j}^B$ can not spend $\text{TX}.\text{Output}$ and hence $m^* \neq [\text{TX}_{\text{RV},j}^B]$ with $j < i$.

- We have $\gamma.\text{flag} = 2$ and $\Gamma^A(\text{id})[1] \neq \perp$ and hence $[\text{TX}] = [\text{TX}_{\text{CM},i+1}^A]$ holds. Also, since $\text{TX}'.\text{Witness}.\eta = 2$, $\text{TX}'.\text{Witness}.\zeta$ satisfies the condition $pk_{\text{RV}}^A \wedge pk_{\text{RV}}^B \wedge \text{CLTV}_{S_0+i+1}$. The reduction outputs (m^*, σ^*) with $m^* = [\text{TX}']$ and $\sigma^* \in \text{TX}'.\text{Witness}.\zeta$. Moreover, the signing oracle was only called for messages $\overline{[\text{TX}_{\text{RV},j}^B]}$ with $j \leq i$ and since for these transactions we have $\overline{[\text{TX}_{\text{RV},j}^B]}.n\text{LT} = S_0 + j < S_0 + i + 1$, according to $\mathcal{L}$, $\text{TX}_{\text{RV},j}^B$ cannot spend $\text{TX}.\text{Output}$ and hence $m^* \neq [\text{TX}_{\text{RV},j}^B]$ with $j \leq i$.

Therefore, the reduction has never called the signing oracle for the message m^* and hence the reduction outputs a valid forgery with probability at least $\frac{1}{\text{poly}(\lambda)}$, which contradicts with our assumption regarding the security of Σ . This contradiction proves that $\Pr[\text{Error} \mid \mathcal{H}_2^{\text{Pu}}] < \frac{1}{\text{poly}(\lambda)}$.

□

Lemma 7.12. *For all PPT distinguishers $\mathcal{E}$ it holds that*

$$\begin{aligned} & \{\text{EXE}_{\mathcal{H}_3^{\text{Pu}}, \mathcal{A}, \mathcal{E}}^{\mathcal{L}(\Delta, \Sigma), \mathcal{F}^{\text{clock}}}(\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*} \approx \\ & \{\text{EXE}_{\mathcal{H}_4^{\text{Pu}}, \mathcal{A}, \mathcal{E}}^{\mathcal{L}(\Delta, \Sigma), \mathcal{F}^{\text{clock}}}(\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*}. \end{aligned}$$

Proof. Similar to previous proofs, assume that $\Pr[\text{Error} \mid \mathcal{H}_3^{\text{Pu}}] \geq \frac{1}{\text{poly}(\lambda)}$. The reduction receives as input a public key pk , and in the channel creation phase, sets $pk_{\text{SP}}^A = pk$. The channel is updated any arbitrary number of times and might be closed at any time using any method (peacefully or forcefully) selected by the adversary. Once the channel is created or once it is updated for the j^{th} time, a call to the signing oracle is performed to receive the signature on $\text{TX}_{\text{SP},0}$ and $\text{TX}_{\text{SP},j}$, respectively. Assume that the channel has been updated i times. If the output of the funding transaction is spent by a transaction TX with $[\text{TX}] \in \{[\text{TX}_{\text{CM},i}^A], [\text{TX}_{\text{CM},i}^B]\}$ given that $\gamma.\text{flag} = 1$ or $[\text{TX}] \in \{[\text{TX}_{\text{CM},i}^A], [\text{TX}_{\text{CM},i}^B], [\text{TX}_{\text{CM},i+1}^A], [\text{TX}_{\text{CM},i+1}^B]\}$ given that $\gamma.\text{flag} = 2$, the experiment does not output Error and the reduction aborts. Otherwise, the reduction waits for t rounds and then posts the transaction TX' on the ledger $\mathcal{L}$ with $\text{TX}' = \text{TX}_{\text{SP},i}$ given that $\gamma.\text{flag} = 1$ or $\text{TX}' = \text{TX}_{\text{SP},i+1}$ given that $\gamma.\text{flag} = 2$.

If TX' is accepted by $\mathcal{L}$, the experiment does not output Error and the reduction aborts. If $\text{TX}.\text{Output}$ is spent by a transaction TX'' with $[\text{TX}''] \neq [\text{TX}']$ we know that either the first or the second sub-condition of $\text{TX}.\text{Output}$ is satisfied by $\text{TX}''.\text{Witness}$. The second sub-condition of $\text{TX}.\text{Output}$ is not satisfied by $\text{TX}''.\text{Witness}$ otherwise $\mathcal{H}_2^{\text{Pu}}$ or $\mathcal{H}_3^{\text{Pu}}$

would output Error which contradicts with our assumptions. Thus, $\text{TX}''.$ Witness satisfies the first sub-condition of $\text{TX}.$ Output. Therefore, the reduction outputs (m^*, σ^*) with $m^* = \overline{[\text{TX}'']}$ and $\sigma^* \in \text{TX}''.$ Witness. ζ . Moreover, the signing oracle was only called for messages $\overline{[\text{TX}_{\text{SP},j}]}$ with $j = [0, i]$ given that $\gamma.\text{flag} = 1$ or $j = [0, i + 1]$ given that $\gamma.\text{flag} = 2$. However, $m^* \notin \{\overline{[\text{TX}_{\text{SP},i}], \overline{[\text{TX}_{\text{SP},i+1}]}}\}$. Otherwise, $\mathcal{H}_4^{\text{Pu}}$ would not output Error. Also, $m^* \notin \{\overline{[\text{TX}_{\text{SP},j}]}$ with $j = [0, i - 1]$ because for these transactions we have $\overline{[\text{TX}_{\text{SP},j}].\text{nLT}} = S_0 + j < S_0 + i$, and hence according to $\mathcal{L}$, $\text{TX}_{\text{SP},j}$ can not spend $\text{TX}.$ Output. Therefore, the reduction has never called the signing oracle for the message m^* and hence the reduction outputs a valid forgery with probability at least $\frac{1}{\text{poly}(\lambda)}$, which contradicts with our assumption regarding the security of Σ . This contradiction proves that $\Pr[\text{Error} \mid \mathcal{H}_3^{\text{Pu}}] < \frac{1}{\text{poly}(\lambda)}$.

□

Lemma 7.13. *For all PPT distinguishers $\mathcal{E}$ it holds that*

$$\begin{aligned} \{\text{EXE}_{\mathcal{H}_4^{\text{Pu}}, \mathcal{A}, \mathcal{E}}^{\mathcal{L}(\Delta, \Sigma), \mathcal{F}^{\text{clock}}}(\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*} &\approx \\ \{\text{EXE}_{\mathcal{H}_5^{\text{Pu}}, \mathcal{A}, \mathcal{E}}^{\mathcal{L}(\Delta, \Sigma), \mathcal{F}^{\text{clock}}}(\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*}. \end{aligned}$$

Proof. Similar to previous proofs, we construct a reduction against the existential unforgeability of the underlying signature scheme Σ . Assume that $\Pr[\text{Error} \mid \mathcal{H}_4^{\text{Pu}}] \geq \frac{1}{\text{poly}(\lambda)}$. The reduction receives as input a public key pk from the challenger and in the channel creation phase sets $pk_A := pk$. The channel is updated any arbitrary number of times and might be closed at any time using any method (peacefully or forcefully) selected by the adversary. All calls to the signing algorithm are redirected to the signing oracle. Assume that the channel has been updated i times. Now assume that the output of the funding transaction is spent by a transaction TX s.t. either of the following two sets of conditions hold:

- $\gamma.\text{flag} = 1$, $[\text{TX}] \notin \{[\text{TX}_{\text{CM},i}^A], [\text{TX}_{\text{CM},j}^B]\}$, $j = [0, i]$ and $\text{TX}.$ Output $\neq \gamma.\text{st}$.
- $\gamma.\text{flag} = 2$, $[\text{TX}] \notin \{[\text{TX}_{\text{CM},i}^A], [\text{TX}_{\text{CM},i+1}^A], [\text{TX}_{\text{CM},j}^B]\}$, $j = [0, i + 1]$ and $\text{TX}.$ Output $\neq \gamma.\text{st}$.

Then, the hybrid outputs Error. The reduction also outputs (m^*, σ^*) with $m^* = [\text{TX}]$ and $\sigma^* \in \text{TX}.$ Witness. ζ . The signature σ^* is a valid signature on m^* because $\text{TX}.$ Witness satisfies the condition of $\text{TX}_{\text{FU}}.$ Output which is $pk_A \wedge pk_B$. Moreover, as we will show in the next paragraph, the signing oracle was never called for the message m^* .

Once the channel is created or each time it is updated, a call to the signing oracle is performed to receive the signature on the funding transaction as well as the new commit transaction held by B , i.e. $[\text{TX}_{\text{CM},j}^B]$ with $j = [0, i]$ if $\gamma.\text{flag} = 1$ or $j = [0, i + 1]$ otherwise. Also, if the channel is closed forcefully, A calls the signing oracle to receive the signature

on $[TX_{CM,i}^A]$ if $\gamma.\text{flag} = 1$ or either $[TX_{CM,i}^A]$ or $[TX_{CM,i+1}^A]$ otherwise. If the channel is closed peacefully, A calls the signing oracle to receive the signature on $[TX]$ with $TX.\text{Output} = \gamma.\text{st}$. Therefore, the reduction has never called the oracle for the message m^* . $\square$

Lemma 7.14. *For all PPT distinguishers $\mathcal{G}$ it holds that*

$$\begin{aligned} \{\text{EXE}_{\mathcal{H}_5^{Pu}, \mathcal{A}, \mathcal{G}}^{\mathcal{L}(\Delta, \Sigma), \mathcal{F}^{clock}}(\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*} &\approx \\ \{\text{EXE}_{\varphi_{\mathcal{F}}, \mathcal{G}^{Pu}, \mathcal{G}}^{\mathcal{L}(\Delta, \Sigma), \mathcal{F}^{clock}}(\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*}. \end{aligned}$$

Proof. The two experiments are identical, and hence, indistinguishability follows. $\square$

7.13 Discussions

Compatibility with P2WSH transactions: Let the output of commit transactions be of type P2WSH, meaning that it can be spent based on the fulfilment of the script whose hash is included in the condition part of the output. Now, assume that while the channel is in state n , party A publishes $TX_{CM,i}^A$ with $i < n$. According to the protocol, party B is supposed to instantly publish the latest revocation transaction $TX_{RV,n-1}^B$. To do so, he must create the original script of the main output of $TX_{CM,i}^A$ and add it to the witness data of $TX_{RV,n-1}^B$. Since the parameter i is a part of the script, B must extract the value of i from the published commit transaction. However, i varies in different commit transactions and its value cannot be directly derived from the hash of the script in the commit transaction output. Therefore, the value of i must be encoded in $TX_{CM,i}^A.\text{nLT}$ or in the parameter *sequence* of $TX_{CM,i}^A.\text{Input}$.

Fee handling: Once a dishonest channel party publishes a revoked commit transaction, her counterparty has T rounds time to publish the revocation transaction on the ledger. However, the time it takes for a transaction to be published depends on two factors: (1) network congestion at the time when the transaction is submitted to the blockchain network, and (2) the transaction fee. Revocation transactions in Daric have a single input and a single output. Since based on BIP 143 [58], the ANYPREVOUT flag may be combined with SINGLE flag, it is possible for a channel party to add a new input and a new output to the latest revocation transaction before submitting it to the blockchain. The difference between the value of the new output and the new input can be collected by miners. A similar approach can also be used for commit transactions.

Compatibility with any digital signature scheme: Generalized and FPPW payment channels leverage adaptor signatures and hence may not work if the current Bitcoin digital signature scheme changes to BLS [20] or a post-quantum digital signature. However, Daric is compatible with any digital signature and can benefit from their properties.

Extending Daric to multi-hop payments: Payment channels typically use HTLC to establish a PCN, where parties who do not have a shared channel can still exchange coins by using other nodes as relays. For Daric, since state duplication is avoided, there is no complications in adding HTLC outputs to split transactions.

Other applications: To have a new application on top of a Daric channel, parties must update the channel state such that the new split transaction has one or multiple outputs for the new application. For example, assume that channel parties want to create multiple channels on top of their existing channel. To do so, they update their channel such that the new split transaction of the channel consists of multiple outputs where each output is a 2-of-2 multisignature address shared between channel parties and acts like the output of a funding transaction for a new Daric channel. The only difference between this new channel and the original one is that since the split transaction for the original channel is floating, its transaction identifier depends on its input and so cannot be determined in advance. Hence, the commit transactions of new established channels must be also floating. The only important criteria for new channels is that each channel must have its own set of public keys. Otherwise, for example, a commit transaction from one channel can spend funding transaction output of another channel.

Channel reset: If the lifetime of a Daric channel is close to its end (which occurs if the channel update rate is more than once per second), channel parties can reset the channel off-chain. To do so, they update the channel such that output of the split transaction in the latest state acts like the output of the funding transaction for a new channel. All the state numbers also reset and the new established channel can be updated at least for about 1 billion times again. Along with required data from the new established channel, each party must also maintain the last commit, split and revocation transactions from the original channel.

7.14 Conclusion and Future Work

In this work, we presented an efficient payment channel with unlimited lifetime for Bitcoin, called Daric, that achieves constant storage. Moreover, the new scheme allows the honest channel party to penalise her dishonest counterparty by taking all the channel funds. Daric also guarantees that channel parties can close the channel within a bounded time. Furthermore, the new scheme is compatible with any digital signature algorithm and simultaneously avoids state duplication. We proved Daric is secure in the Universal Composability model.

An interesting open topic to study is extending Daric to an efficient m -party scheme with $m > 2$. Moreover, one of the main advantages of Daric is that the storage requirements

of the watchtower for each channel could be constant over time. However, there are also other factors for a watchtower (e.g. privacy, fairness, and coverage [54]) that must be carefully taken into account. Designing a watchtower for Daric which can achieve optimality in terms of the above mentioned properties could be another subject for future research.

Chapter 8

Conclusion and future work

Conclusion

In this thesis, we focused on different aspects of the payment channel which is a promising solution to the scalability issue of Bitcoin. Since the deployment of the payment channel does not require any changes in the Bitcoin protocol, this idea is already being used in the Lightning Network. To prevent their counterparties from closing the channel with an old state, each party in the Lightning Network should frequently monitor the blockchain. Alternatively, channel parties might employ a third-party service provider, called the watchtower, to do the monitoring task.

In this thesis, first, we formally defined the watchtower and its desired properties. Moreover, we proved a trade-off between two of those properties, i.e. the fairness towards the channel party and the coverage. This trade-off prevents a watchtower scheme from achieving both optimal fairness and coverage. We also compared all the existing watchtower schemes with respect to all the defined properties. This comparison showed none of the current watchtower schemes for Bitcoin achieve fairness towards the channel party and simultaneously protect the channel privacy against the watchtower. These contributions answer our first research question, RQ1, about the formal definition of different properties of a watchtower and the limitations of the existing watchtower schemes in meeting those properties.

We solved the mentioned fairness-privacy problem by designing a watchtower scheme called FPPW which can achieve both fairness and privacy at the same time. Coverage for this scheme is also the highest possible value for a fully-fair watchtower scheme. Focusing on performance features, we observed that for all current watchtower schemes as well as FPPW, the storage costs of the channel parties or their watchtowers increase linearly with each channel update. Thus, we also focused on improving the efficiency of

storage costs of the channel parties and their watchtower and designed a second watchtower scheme called Garrison. The storage costs of all participants for this new scheme increase logarithmically with the number of channel updates. Both FPPW and Garrison avoid state duplication (i.e. both parties store the same version of transactions) and can be implemented without any update in the Bitcoin blockchain. Designing these two watchtower schemes for existing payment channels answer our second research question, RQ2.

At the next step, rather than designing watchtowers for the existing payment channels, we focused on current Bitcoin payment channels and analysed their limitations in achieving all the properties required for a payment channel. Then, relying on the deployment of an already proposed signature type (called ANYPREVOUT), we designed a new payment channel with an unlimited lifetime called Daric. This new scheme with desired computation and communication complexity is optimal in storage, provides a penalisation mechanism and avoids state duplication without relying on any particular property of the underlying digital signature. We also proved the security of Daric in the UC framework. These contributions answer our third research question, RQ3, about the limitations of the existing payment channels and how to design a provably secure payment channel to mitigate those limitations.

In general, our work benefits payment channel research by providing a rigorous study on different aspects of this important type of solution to cryptocurrency scalability. We defined some properties based on which the future payment channels and watchtower schemes can be evaluated. We also designed two watchtower schemes for current payment channels, each focusing on different aspects and hence applicable in different scenarios. All the desirable features of our proposed payment channel might also urge the Bitcoin community to deploy the ANYPREVOUT signature type and benefit from its applications.

Future Work

Despite their advantages, FPPW and Garrison still suffer from some limitations: (i) the storage costs of channel parties and their watchtowers in FPPW increase linearly with each channel update and (ii) Garrison does not achieve fairness with respect to the channel party. Designing a watchtower scheme that mitigates the mentioned limitations of these two schemes has been left to future works. Furthermore, we have established the existence of a trade-off between fairness and coverage in watchtower schemes. Specifically, in order to ensure fairness towards the channel party, the watchtower is required to lock collateral as compensation for potential losses incurred by the hiring party. However, this collateral locking mechanism can weaken the overall coverage provided by the

watchtower. Exploring alternative methods to guarantee the watchtower's service without relying solely on collateral locking presents an intriguing avenue for future research. By investigating such approaches, we can address this trade-off and further enhance the effectiveness of watchtower schemes.

Our proposed payment channel performs well for payment applications when we have 2 parties involved. Exploring the extension of this idea to encompass more general multi-party applications, such as the channel factory [74, 75], represents an interesting and important open topic. The significance lies in the potential to leverage the concept's advantages, such as scalability and efficiency, in broader multi-party scenarios. For instance, by adapting and applying these principles to the channel factory setting, we can potentially achieve enhanced scalability, improved transaction throughput, and reduced costs. Further research and exploration in this direction are warranted to better understand the implications, challenges, and opportunities that arise when extending the idea to multi-party applications beyond the immediate payment channel context. Moreover, designing a watchtower for Daric with all the required properties is left as a future work.

References

- [1] Christian Decker, Rusty Russell, and Olaoluwa Osuntokun. eltoo: A simple layer2 protocol for bitcoin. URL: <https://blockstream.com/eltoo.pdf>, 2018.
- [2] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. *Decentralized business review*, 2008.
- [3] Tobias Bamert, Christian Decker, Lennart Elsen, Roger Wattenhofer, and Samuel Welten. Have a snack, pay with bitcoins. In *IEEE P2P 2013 Proceedings*, pages 1–5. IEEE, 2013.
- [4] Christian Decker and Roger Wattenhofer. A fast and scalable payment network with bitcoin duplex micropayment channels. In *Symposium on Self-Stabilizing Systems*, pages 3–18. Springer, 2015.
- [5] Yonatan Sompolinsky and Aviv Zohar. Accelerating bitcoin’s transaction processing. fast money grows on trees, not chains. *Cryptology ePrint Archive*, 2013.
- [6] Arthur Gervais, Ghassan O Karame, Karl Wüst, Vasileios Glykantzis, Hubert Ritzdorf, and Srdjan Capkun. On the security and performance of proof of work blockchains. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pages 3–16, 2016.
- [7] Kyle Croman, Christian Decker, Ittay Eyal, Adem Efe Gencer, Ari Juels, Ahmed Kosba, Andrew Miller, Prateek Saxena, Elaine Shi, Emin Gün Sirer, et al. On scaling decentralized blockchains. In *International conference on financial cryptography and data security*, pages 106–125. Springer, 2016.
- [8] Manny Trillo. Stress test prepares visanet for the most wonderful time of the year. URL: <http://www.visa.com/blogarchives/us/2013/10/10/stress-testprepares-visanet-for-the-most-wonderfultime-of-the-year/index.html>, 2013.
- [9] Aggelos Kiayias, Alexander Russell, Bernardo David, and Roman Oliynykov. Ouroboros: A provably secure proof-of-stake blockchain protocol. In *Annual International Cryptology Conference*, pages 357–388. Springer, 2017.

- [10] Ittay Eyal, Adem Efe Gencer, Emin Gün Sirer, and Robbert Van Renesse. Bitcoin-ng: A scalable blockchain protocol. In *13th USENIX symposium on networked systems design and implementation (NSDI 16)*, pages 45–59, 2016.
- [11] Loi Luu, Viswesh Narayanan, Chaodong Zheng, Kunal Baweja, Seth Gilbert, and Prateek Saxena. A secure sharding protocol for open blockchains. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 17–30, 2016.
- [12] Adem Efe Gencer, Robbert van Renesse, and Emin Gün Sirer. Service-oriented sharding with aspen. *arXiv preprint arXiv:1611.06816*, 2016.
- [13] Adam Back, Matt Corallo, Luke Dashjr, Mark Friedenbach, Gregory Maxwell, Andrew Miller, Andrew Poelstra, Jorge Timón, and Pieter Wuille. Enabling blockchain innovations with pegged sidechains. URL: <http://www.opensciencereview.com/papers/123/enablingblockchain-innovations-with-pegged-sidechains>, 72:201–224, 2014.
- [14] Lukas Aumayr, Oguzhan Ersoy, Andreas Erwig, Sebastian Faust, Kristina Hostáková, Matteo Maffei, Pedro Moreno-Sanchez, and Siavash Riahi. Generalized channels from limited blockchain scripts and adaptor signatures. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 635–664. Springer, 2021.
- [15] Joseph Poon and Thaddeus Dryja. The bitcoin lightning network: Scalable off-chain instant payments. 2016.
- [16] Lewis Gudgeon, Pedro Moreno-Sanchez, Stefanie Roos, Patrick McCorry, and Arthur Gervais. Sok: Off the chain transactions. *IACR Cryptology ePrint Archive*, 2019:360, 2019.
- [17] Andrew Miller, Iddo Bentov, Ranjit Kumaresan, and Patrick McCorry. Sprites: Payment channels that go faster than lightning. *CoRR abs/1702.05812*, 306, 2017.
- [18] Rami Khalil, Arthur Gervais, and Guillaume Felley. Nocust-a securely scalable commit-chain. *Cryptology ePrint Archive, Report 2018/642*, 2018.
- [19] Joseph Poon and Vitalik Buterin. Plasma: Scalable autonomous smart contracts. *White paper*, pages 1–47, 2017.
- [20] Dan Boneh, Ben Lynn, and Hovav Shacham. Short signatures from the weil pairing. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 514–532. Springer, 2001.

- [21] Thaddeus Dryja and Scaling Bitcoin Milano. Unlinkable outsourced channel monitoring. URL: <https://diyhpl.us/wiki/transcripts/scalingbitcoin/milan/unlinkable-outsourced-channel-monitoring>, 2016.
- [22] Majid Khabbazzian, Tejaswi Nadahalli, and Roger Wattenhofer. Outpost: A responsive lightweight watchtower. In *Proceedings of the 1st ACM Conference on Advances in Financial Technologies*, pages 31–40, 2019.
- [23] Zeta Avarikioti, Orfeas Stefanos Thyfronitis Litos, and Roger Wattenhofer. Cerberus channels: Incentivizing watchtowers for bitcoin. In *Financial Cryptography and Data Security: 24th International Conference, FC 2020, Kota Kinabalu, Malaysia, February 10–14, 2020 Revised Selected Papers 24*, pages 346–366. Springer, 2020.
- [24] Patrick McCorry, Surya Bakshi, Iddo Bentov, Sarah Meiklejohn, and Andrew Miller. Pisa: Arbitration outsourcing for state channels. In *Proceedings of the 1st ACM Conference on Advances in Financial Technologies*, pages 16–30, 2019.
- [25] Shehar Bano, Alberto Sonnino, Mustafa Al-Bassam, Sarah Azouvi, Patrick McCorry, Sarah Meiklejohn, and George Danezis. Sok: Consensus in the age of blockchains. In *Proceedings of the 1st ACM Conference on Advances in Financial Technologies*, pages 183–198, 2019.
- [26] J Spilman. [bitcoin-development] anti dos for tx replacement. URL: <https://lists.linuxfoundation.org/pipermail/bitcoin-dev/2013-April/002433.html>, 2013.
- [27] Lukas Aumayr, Matteo Maffei, Oğuzhan Ersoy, Andreas Erwig, Sebastian Faust, Siavash Riahi, Kristina Hostáková, and Pedro Moreno-Sanchez. Bitcoin-compatible virtual channels. In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 901–918. IEEE, 2021.
- [28] Lukas Aumayr, Sri AravindaKrishnan Thyagarajan, Giulio Malavolta, Pedro Moreno-Sanchez, and Matteo Maffei. Sleepy channels: Bitcoin-compatible bi-directional payment channels without watchtowers. *Cryptology ePrint Archive*, 2021.
- [29] Stefan Dziembowski, Lisa Ekey, Sebastian Faust, and Daniel Malinowski. Perun: Virtual payment hubs over cryptocurrencies. In *2019 IEEE Symposium on Security and Privacy (SP)*, pages 106–123. IEEE, 2019.
- [30] Matthew Green and Ian Miers. Bolt: Anonymous payment channels for decentralized currencies. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 473–489, 2017.

- [31] Christian Decker and Anthony Towns. Bip 118 - sighashnoinput. URL: <https://github.com/bitcoin/bips/blob/master/bip-0118.mediawiki>, 2017.
- [32] Georgia Avarikioti, Felix Laufenberg, Jakub Sliwinski, Yuyi Wang, and Roger Wattenhofer. Towards secure and efficient payment channels. *arXiv preprint arXiv:1811.12740*, 2018.
- [33] Sonbol Rahimpour and Majid Khabbazzian. Hashcached reputation with application in designing watchtowers. In *2021 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, pages 1–9. IEEE, 2021.
- [34] Adam Back. Hashcash-a denial of service counter-measure. 2002.
- [35] Bowen Liu, Pawel Szalachowski, and Siwei Sun. Fail-safe watchtowers and short-lived assertions for payment channels. *arXiv preprint arXiv:2003.06127*, 2020.
- [36] Marc Leinweber, Matthias Grundmann, Leonard Schönborn, and Hannes Hartenstein. Tee-based distributed watchtowers for fraud protection in the lightning network. In *Data Privacy Management, Cryptocurrencies and Blockchain Technology*, pages 177–194. Springer, 2019.
- [37] Georgia Avarikioti, Eleftherios Kokoris Kogias, and Roger Wattenhofer. Brick: Asynchronous state channels. *arXiv preprint arXiv:1905.11360*, 2019.
- [38] Ethan Heilman, Leen Alshenibr, Foteini Baldimtsi, Alessandra Scafuro, and Sharon Goldberg. Tumblebit: An untrusted bitcoin-compatible anonymous payment hub. In *Network and Distributed System Security Symposium*, 2017.
- [39] Erkan Tairi, Pedro Moreno-Sanchez, and Matteo Maffei. A²l: Anonymous atomic locks for scalability in payment channel hubs. In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 1834–1851. IEEE, 2021.
- [40] Noemi Glaeser, Matteo Maffei, Giulio Malavolta, Pedro Moreno-Sanchez, Erkan Tairi, and Sri Aravinda Krishnan Thyagarajan. Foundations of coin mixing services. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, pages 1259–1273, 2022.
- [41] Bitcoin Wiki. Hashed timelock contracts. URL: https://en.bitcoin.it/wiki/Hash_Time_Locked_Contracts, 2019.
- [42] Raiden network. URL: <http://raiden.network/>, 2017.
- [43] Vitalik Buterin. Ethereum white paper: a next generation smart contract & decentralized application platform. *First version*, 53, 2014.

- [44] Christoph Egger, Pedro Moreno-Sanchez, and Matteo Maffei. Atomic multi-channel updates with constant collateral in bitcoin-compatible payment-channel networks. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pages 801–815, 2019.
- [45] Maxim Jourenko, Mario Larangeira, and Keisuke Tanaka. Payment trees: Low collateral payments for payment channel networks. In *International Conference on Financial Cryptography and Data Security*, pages 189–208. Springer, 2021.
- [46] Giulio Malavolta, Pedro Moreno-Sanchez, Clara Schneidewind, Aniket Kate, and Matteo Maffei. Anonymous multi-hop locks for blockchain scalability and interoperability. *Cryptology ePrint Archive*, 2018.
- [47] Giulio Malavolta, Pedro Moreno-Sanchez, Aniket Kate, Matteo Maffei, and Srivatsan Ravi. Concurrency and privacy with payment-channel networks. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 455–471, 2017.
- [48] Bin Yu, Shabnam Kasra Kermanshahi, Amin Sakzad, and Surya Nepal. Chameleon hash time-lock contract for privacy preserving payment channel networks. In *Provable Security: 13th International Conference, ProvSec 2019, Cairns, QLD, Australia, October 1–4, 2019, Proceedings 13*, pages 303–318. Springer, 2019.
- [49] Sri Aravinda Krishnan Thyagarajan and Giulio Malavolta. Lockable signatures for blockchains: Scriptless scripts for all signatures. In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 937–954. IEEE, 2021.
- [50] Lukas Aumayr, Pedro Moreno-Sanchez, Aniket Kate, and Matteo Maffei. Blitz: Secure multi-hop payments without two-phase commits. In *30th USENIX Security Symposium*, pages 4043–4060, 2021.
- [51] Lukas Aumayr, Oguzhan Ersoy, Andreas Erwig, Sebastian Faust, Kristina Hostáková, Matteo Maffei, Pedro Moreno-Sanchez, and Siavash Riahi. Bitcoin-compatible virtual channels. *IACR Cryptol. ePrint Arch.*, 2020:554, 2020.
- [52] Don Johnson, Alfred Menezes, and Scott Vanstone. The elliptic curve digital signature algorithm (ecdsa). *International journal of information security*, 1:36–63, 2001.
- [53] Claus-Peter Schnorr. Efficient signature generation by smart cards. *Journal of cryptology*, 4:161–174, 1991.
- [54] Arash Mirzaei, Amin Sakzad, Jiangshan Yu, and Ron Steinfeld. Fppw: A fair and privacy preserving watchtower for bitcoin. In *International Conference on Financial Cryptography and Data Security*, pages 151–169. Springer, 2021.

- [55] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *Proceedings 42nd IEEE Symposium on Foundations of Computer Science*, pages 136–145. IEEE, 2001.
- [56] Juan Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol with chains of variable difficulty. In *Annual International Cryptology Conference*, pages 291–323. Springer, 2017.
- [57] libbitcoin. Sighash and tx signings. URL: <https://github.com/libbitcoin/libbitcoin-system/wiki/Sighash-and-TX-Signing>, 2018.
- [58] Johnson Lau and Pieter Wuille. Bip 143: Transaction signature verification for version 0 witness program. URL: <https://github.com/bitcoin/bips/blob/master/bip-0143.mediawiki>, 2016.
- [59] Andreas M Antonopoulos. *Mastering Bitcoin: unlocking digital cryptocurrencies*. ” O’Reilly Media, Inc.”, 2014.
- [60] Arash Mirzaei, Amin Sakzad, Jiangshan Yu, and Ron Steinfeld. Garrison: a novel watchtower scheme for bitcoin. In *Australasian Conference on Information Security and Privacy*, pages 489–508. Springer, 2022.
- [61] Lightning Developers. Bolt# 3: Bitcoin transaction and script formats, 2017.
- [62] Yehuda Lindell. Fast secure two-party ecdsa signing. In *Annual International Cryptology Conference*, pages 613–644. Springer, 2017.
- [63] Arash Mirzaei, Amin Sakzad, Jiangshan Yu, and Ron Steinfeld. Daric: a storage efficient payment channel with punishment mechanism. In *International Conference on Information Security*, pages 229–249. Springer, 2022.
- [64] Rene Pickhardt. Does eltoo eliminate the need to watch the blockchain/implement watchtowers. URL: <https://bitcoin.stackexchange.com/questions/84846/does-eltoo-eliminate-the-need-to-watch-the-blockchain-implement-watchtowers>, 2019.
- [65] Lightning channels - top capacity. URL: <https://1ml.com/channel?order=capacity>.
- [66] Using per-update credential to enable eltoo-penalty. URL: <https://lists.linuxfoundation.org/pipermail/lightning-dev/2019-July/002068.html>, 2019.
- [67] eltoo: A simplified update mechanism for lightning and off-chain contracts. URL: <https://lists.linuxfoundation.org/pipermail/lightning-dev/2018-June/001313.html>, 2018.

- [68] Stefan Dziembowski, Sebastian Faust, and Kristina Hostáková. General state channel networks. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pages 949–966, 2018.
- [69] Stefan Dziembowski, Lisa Ekey, Sebastian Faust, Julia Hesse, and Kristina Hostáková. Multi-party virtual state channels. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 625–656. Springer, 2019.
- [70] Ran Canetti, Yevgeniy Dodis, Rafael Pass, and Shabsi Walfish. Universally composable security with global setup. In *Theory of Cryptography Conference*, pages 61–85. Springer, 2007.
- [71] Peter Todd and David A. Harding. Bip 125: Opt-in full replace-by-fee signaling. URL: <https://github.com/bitcoin/bips/blob/master/bip-0125.mediawiki>, 2015.
- [72] Bolt 5: Recommendations for on-chain transaction handling. URL: <https://github.com/lightningnetwork/lightning-rfc/blob/master/05-onchain.md>.
- [73] Christian Badertscher, Ueli Maurer, Daniel Tschudi, and Vassilis Zikas. Bitcoin as a transaction ledger: A composable treatment. In *Annual international cryptology conference*, pages 324–356. Springer, 2017.
- [74] Conrad Burchert, Christian Decker, and Roger Wattenhofer. Scalable funding of bitcoin micropayment channel networks. *Royal Society open science*, 5(8):180089, 2018.
- [75] Alejandro Ranchal Pedrosa, Maria Potop-Butucaru, and Sara Tucci-Piergiovanni. Scalable lightning factories for bitcoin. In *Proceedings of the 34th ACM/SIGAPP Symposium on Applied Computing*, pages 302–309, 2019.

