

# P4TC—Provably-Secure yet Practical Privacy-Preserving Toll Collection

MAX HOFFMANN<sup>\*</sup>, Ruhr-Universität Bochum, Germany

VALERIE FETZER<sup>†</sup>, MATTHIAS NAGEL<sup>‡</sup>, ANDY RUPP<sup>§</sup>, and REBECCA SCHWERDT<sup>¶</sup>, Karlsruhe Institute of Technology, Germany

Electronic toll collection (ETC) is widely used all over the world not only to finance our road infrastructures, but also to realize advanced features like congestion management and pollution reduction by means of dynamic pricing. Unfortunately, existing systems rely on user identification and allow tracing a user’s movements. Several abuses of this personalized location data have already become public. In view of the planned European-wide interoperable tolling system EETS and the new EU General Data Protection Regulation, location privacy becomes of particular importance.

In this paper, we propose a flexible security model and crypto protocol framework designed for privacy-preserving toll collection in the most dominant setting, i.e., Dedicated Short Range Communication (DSRC) ETC. As opposed to our work, most related cryptographic proposals target a less popular type of toll collection based on Global Navigation Satellite Systems (GNSS), or do not come with a thorough security model and proof. A major challenge in designing the framework at hand was to combine provable security and practicality, where the latter includes practical performance figures and a suitable treatment of real-world issues, like broken on-board units etc. To the best of our knowledge, our work is the first in the DSRC setting with a rigorous security model and proof and arguably the most comprehensive formal treatment of ETC security and privacy overall.

For our ETC system, we make use of and significantly extend a payment protocol building block, called BBA+, introduced at ACM CCS 2017. Additionally, we provide a prototypical implementation on realistic hardware. This implementation already features fairly practical performance figures, even though there is still room for optimizations. An interaction between an on-board unit and a road-side unit is estimated to take less than a second allowing for toll collection at full speed assuming one road-side unit per lane.

## CONTENTS

1	Introduction	2
2	Security Model	10
3	System Definition	12
4	Protocol Overview	17
5	Security Theorem	22
6	Performance Evaluation	23
	References	26
A	Information Leakage and Discussion on Privacy Implications	29
B	Full System Definition	32
C	Setting and Building Blocks	43
D	Full Protocol Description	54
E	Adversarial Model	75

<sup>\*</sup>The author is supported by DFG grant PA 587/10-1.

<sup>†</sup>The author is supported by DFG grant RU 1664/3-1.

<sup>‡</sup>The author is supported by the German Federal Ministry of Education and Research within the framework of the project “Sicherheit vernetzter Infrastrukturen (SVI)” in the Competence Center for Applied Security Technology (KASTEL).

<sup>§</sup>The author is supported by DFG grant RU 1664/3-1 and the Competence Center for Applied Security Technology (KASTEL).

<sup>¶</sup>The author is supported by the German Research Foundation (DFG) as part of the Research Training Group GRK 2153: Energy Status Data-Informatics Methods for its Collection, Analysis and Exploitation.

## 1 INTRODUCTION

Electronic toll collection (ETC) is already deployed in many countries all over the world. A recent study [58] conducted by “Markets and Markets” predicts a CAGR for this market of 9.16% from 2017 to 2022 reaching 10.6 Billion USD by 2022. Europe plans to introduce the first implementation of a fully interoperable tolling system (EETS) until 2027 [30]. Hence, ETC has been and will be an important technology deserving a careful analysis of security and privacy issues as well as custom-tailored secure solutions.

As ETC will become the default option for paying tolls with no easy way to opt-out, privacy is a concern of particular importance. Unfortunately, the systems in use today do not protect the location privacy of their users. This encourages abuse: EZ-Pass records, for instance, have been used as evidence in divorce lawsuits [61], EZ-pass transponders are abused to track drivers throughout New York City [66], and the Norwegian AutoPASS system allowed anyone to obtain a transcript of which toll booths a car visited [32].

However, the only legitimate reason to store personalized location records in these systems is to bill customers. As opposed to other systems, as for instance loyalty systems, this data should not be used for other business relevant purposes like, e.g., personalized advertising. Thus, an efficient and cost-effective privacy-preserving mechanism which avoids data collection in the first place, but still enables the billing functionality, should be of interest to ETC providers as well. In this way, there is no need to deploy costly technical and organizational measures to protect a large amount of sensitive data and there is no risk of a data breach resulting in costly law suits, fines, and loss of customer trust. This is especially interesting in view of the new EU General Data Protection Regulation (GDPR) [38] which is in effect since May 2018. The GDPR stipulates comprehensive protection measures and heavy fines in case of non-compliance.

### 1.1 Classification of ETC

One can classify ETC systems based on what the user is charged for, where toll determination takes place, and how this determination is done [29, 62].

Concerning what the user is charged for, we distinguish between two major types of charging schemes:

*Distance-based:* The toll is calculated based on the distance traveled by the vehicle and adapted by other parameters like the type of vehicle, discounts, etc.

*Access-based:* Tolls apply to a specific geographic area, e.g. part of a city, segment of a highway, tunnel, etc. As before, it can also be dynamically adapted by other parameters. This charging scheme is typically used in urban areas—not only to finance the road infrastructure but also for congestion management and pollution reduction by adapting tolls dynamically.

There are two main types of toll determination environments:

*Toll plaza:* This is the traditional environment where cars pass toll booths on physically separated lanes which may be secured by barriers or cameras to enforce honest behavior.

*Open road:* In open road tolling the traffic is not disrupted as tolls are collected in a seamless fashion without forcing cars to slow down. In the DSRC setting, this is enabled by equipping roads with toll gantries and cameras enforcing participation.

Several key technologies define how toll is determined:

*Dedicated Short-Range Communication (DSRC):* Today, this is the most widely used ETC technology worldwide and the de facto standard in Europe [62]. It is based on bidirectional radio communication between a

road-side unit (RSU) and a mobile device aka on-board unit (OBU) installed in the vehicle. In today's systems, the OBU just identifies the user to trigger a payment. However, more complex protocols (like ours) between OBU and RSU can be implemented.

*Automatic Number Plate Recognition (ANPR)*: ANPR, aka video tolling, inherently violates privacy.

*Global Navigation Satellite System (GNSS)*: In a GNSS-based system the OBU keeps track of the vehicle's location (e.g., by means of GPS) and processes the necessary information to measure its road usage autonomously, i.e., without the aid of an RSU. GNSS is typically used in combination with GSM for communicating with the toll service provider.

## 1.2 Drawbacks of Existing Systems and Proposals

Independent of the technology used, systems deployed in practice build on identifying the user to charge him. Previous work on privacy-preserving toll collection mainly considers the GNSS setting and comes—apart from a few exceptions [6, 28, 31]—without any formal security analysis. Although DSRC is the most dominant setting in practice, it did not receive much attention in the literature so far. Moreover, practical issues like “what happens if an OBU breaks down”, are usually not taken into account by these proposals. We elaborate on related work in [Section 1.4](#).

## 1.3 P4TC

We propose a comprehensive security model as well as provably-secure and efficient protocols for privacy-friendly ETC in the DSRC setting. Our definitional framework and the system are very flexible. We cover access-based and distant-based charging schemes as well as combinations of those. Our protocols work for toll plaza environments, but are (due to offline precomputations) efficient enough for open road tolling. Additionally, we also cope with several issues that may arise in practice, e.g., broken/stolen OBUs, RSUs with non-permanent internet connection, imperfections of violation enforcement technology, etc. To the best of our knowledge, we arguably did the most comprehensive formal treatment of ETC security and privacy overall.

[Section 1.5](#) provides an overview on the toll collection scenario and the desired properties we consider. [Section 1.6](#) gives details on our contribution.

## 1.4 Related Work

In this section, we review proposals for privacy-preserving ETC. So far, more elaborate systems have been proposed for the GNSS setting in comparison to the actually more widely used DSRC setting.

*1.4.1 DSRC (OBU/RSU) Setting*. Previous work [33, 48–50] in this setting mainly focuses on a pre-pay scenario where some form of e-cash is used to spend coins when passing an RSU. This mode is not preferred by users due to its inconveniences. A user needs to ensure that he always has enough e-coins to pay the tolls. This is particularly inconvenient when prices change dynamically. In addition, using standard offline e-cash (e.g., [17]), the user may not overpay since he cannot get change from the RSU in a privacy-preserving way. Also transferrable e-cash [7] does not solve the problem of change as the impossibility result of [19] applies to the ETC setting. So a user would not only need enough money but the right denomination of e-coins.

In [48–50] multiple electronic road pricing systems specifically tailored for Low Emission Zones (LEZ) are proposed. In [48] a user drives by an RSU and directly pays some price depending on this RSU. For [49, 50] the price a user has to pay depends on the time spent inside the LEZ. To this end, the user receives some e-ticket from an Entry-RSU when entering the LEZ which he needs to present again at an Exit-RSU when leaving the LEZ. For the actual payment in all these systems, some untraceable e-cash scheme that supports dynamic pricing is assumed but not specified. The systems require tamper-proof hardware for the OBU and are claimed to provide fraud protection and privacy for honest drivers.

In [33], a simple access-based toll collection system based on RSA blind signatures is sketched. Users buy a set of tokens/coins during registration which are used to pay toll while driving. Double-spending detection can be detected with high probability and some ideas are presented to keep the double-spending database small.

In [10], the authors propose an interesting scheme for distance-based toll collection. Here, a user obtains a coin, used as an entry ticket, which is worth the maximum toll in the system and can be reused a fixed number of times. The actual toll is calculated at the exit RSU where the user is reimbursed for the difference. The system supports revocation of a user's anonymity and token revocation. The instantiation mixes cryptographic primitives from the Paillier and DLog setting. Zero-knowledge proofs for languages which mix statements (and share variables) of the two settings are also required. The actual proofs are not specified in the paper. Instead they are just claimed to be "quite standard", what we doubt. As opposed to ours, their system relies on online "over-spending" detection and does not come with any formal model or security proof.

*GNSS Setting.* A variety of GNSS-based system proposals can be found in the literature, including a series of papers [6, 59, 63] published at USENIX Security. Here, the OBU equipped with GPS and GSM typically collects location-time data or road segment prices and sends this data to the toll service provider (TSP). To ensure that the user behaves honestly and, e.g., does not omit or forge data, unpredictable spot checks are assumed which force a user to reveal the data he sent at a certain location and time. In a reconciliation phase, the user calculates his toll based on the data sent and proves that his calculation is correct.

One example for the GNSS approach is the simple and flexible electronic road pricing system in [34]. There the server collects hash values of the trip records from the OBUs in form of a hash tree. In the reconciliation phase the consistency of the hash tree is verified. However, the whole system is only described imprecisely and the security and privacy of the system are not proven.

In VPriv [63], the OBU anonymously sends tagged location-time data to the TSP while driving. The user previously committed to use exactly these random tags in a registration phase with the TSP. In an inefficient reconciliation phase, each user needs to download the database of all tagged location-time tuples, calculate his total fee, and prove that for this purpose, he correctly used the tuples belonging to his tags without leaking those. In [31] ProVerif is used to show that VPriv is privacy-preserving for honest-but-curious adversaries.

In [45] a road pricing system, where the OBU anonymously sends location data to a central server, is presented. The system is based on splitting a single trip into several unlinkable segments (called "legs") and on distributing the calculation of the fee for a trip between several distinct entities. During the reconciliation phase, every OBU learns the locations of the spot check devices that recorded its corresponding vehicle. The security and privacy requirements are only proved informally.

In the PrETP [6] scheme, the OBU non-anonymously sends payment tuples consisting of commitments to location, time, and the corresponding price that the OBU determined itself. During reconciliation with the TSP, the user presents his total toll and proves that this is indeed the sum of the individual prices he sent, using the homomorphic property of the commitment scheme.

The authors prove their system secure using the ideal/real world paradigm.

In [59], the authors identify large-scale driver collusion as a potential threat to the security of PrETP (and other systems): As spot check locations are leaked to the drivers in the reconciliation phase, they may collude to cheat the system by sharing these locations and only sending correct payment tuples nearby. To this end, the Milo system is constructed. In contrast to PrETP, the location of the spot checks is not revealed during the monthly reconciliation phase. Therefore, drivers are less motivated to collude and cheat. However, if a cheating user is caught, the corresponding spot check location is still revealed. Thus, Milo does not protect against mass-collusion of *dishonest* drivers. No security or privacy proofs are given.

In [27], the authors propose a system based on group signatures that achieves  $k$ -anonymity. A user is assigned to a group of drivers during registration. While driving, the user's OBU sends location, time, and group ID—signed

using one of the group’s signature keys—to the TSP. At the end of a billing period, each user is expected to pay his toll. If the sum of tolls payable by the group (calculated from the group’s location-time data) is not equal to the total toll actually paid by the group, a dispute solving protocol is executed. During dispute solving, the group manager recalculates each user’s toll and, thus, catches the cheater. Choosing an appropriate group size is difficult (the larger the size, the better the anonymity is protected, but the computation overhead rises), as well as choosing a suited group division policy (users in a group should have similar driving regions and similar driving patterns). In [28], the system’s security and privacy properties are verified using ProVerif.

Another cell-based road pricing scheme is presented in [39]. Here, a roadpricing area is divided into cells, where certain cells are randomly selected as spot check cells. A trusted platform module (TPM) inside each OBU is aware of this selection. While driving, the OBU tells its TPM the current location and time. The TPM updates the total toll and also generates a “proof of participation” which is sent to the TSP. This proof is the signed and encrypted location-time data under the TSP’s public key if the user is inside a spot check cell and 0 otherwise. In this way, the TSP can easily verify that a user behaved honestly at spot check cells without leaking their locations to honest users. A security proof is sketched.

A main issue with all systems described above is that their security relies on a strong assumption in practice: invisible spot checks at locations which are unpredictable in each billing period. Otherwise users could easily collude and cheat. On the other hand, spot checks reveal a user’s location. Hence, fixing the number of spot checks is a trade-off between ensuring honesty and preserving privacy. Clearly, the penalty a cheater faces influences the number of spot checks required to ensure a certain security level.

In [52], the authors argue that even mass surveillance, protecting no privacy at all, cannot prevent collusion under reasonable penalties. They present a protocol for a privacy-preserving spot checking device where the locations of these devices can be publicly known. Drivers interacting with the device identify themselves with a certain probability, but do not learn whether they do so, therefore being forced to honesty. To let a user not benefit from turning off his OBU between two spot check devices, such a device is needed at every toll segment. Furthermore, to encourage interaction with the device, an enforcement camera is required. However, since all these road-side devices are needed anyway, there is no advantage in terms of infrastructure requirements compared to the DSRC setting.

## 1.5 Considered Scenario

Here, we sketch the considered scenario, giving an idea of the required flexibility and complexity of our framework. We target a post-payment toll collection system in the DSRC setting which allows access-based as well as distance-based charging and can be deployed in a toll-plaza as well as an open-road environment. It involves the following entities:

- The *Toll Collection Service Provider (TSP)* which might be a privately owned company.
- The *State Authority (SA)*, e.g., the Department of Transportation, which outsourced the toll collection operations to the TSP but is responsible for violation enforcement.
- A *Dispute Resolver (DR)*, e.g., an NGO protecting privacy, which is involved in case of an incident or dispute.<sup>1</sup>
- *Users* who participate in the system by means of a (portable or mounted) On-Board Unit (OBU). The OBU is used for on-road transactions and in the scope of debt and dispute clearance periods. For the latter, it needs to establish a (3G) connection to the TSP/SA. Alternatively, a smartphone might be used for that purpose, which, however, needs access to the OBU’s data.
- *Road-Side Units (RSUs)* which interact with OBUs and are typically managed by the TSP. To enable fast and reliable transactions with OBUs, we do *not* require RSUs to have a permanent connection to the TSP. We

<sup>1</sup>Note that we assume the DR to be trusted by all other parties. It implements an *optional* “kind of key escrow” mechanism.

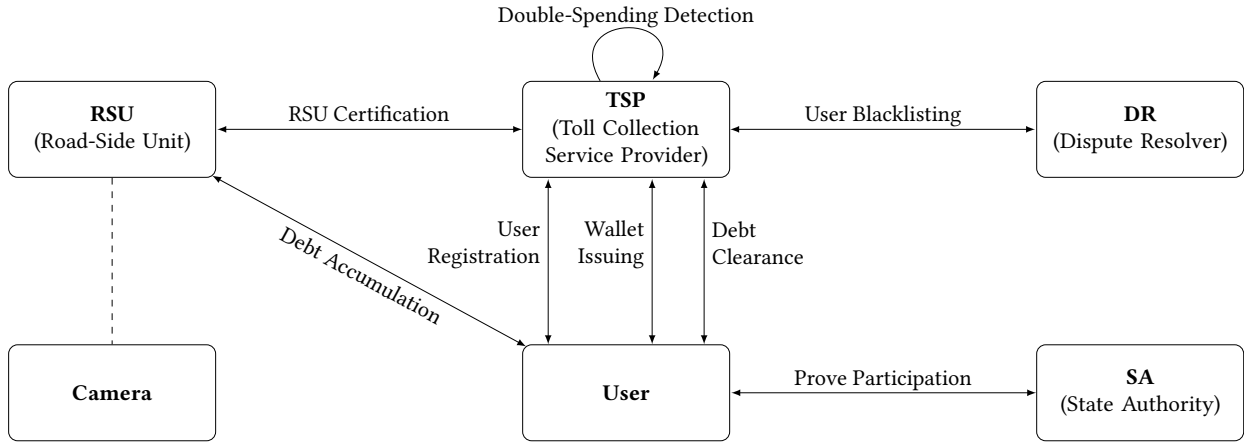


Fig. 1. The P4TC System Model

only assume that they are periodically online for a short duration (presumably at night when there is not much traffic) to exchange data for fraud detection with the TSP.

- *Enforcement Cameras* triggered by the RSUs which are typically owned by the SA and used to make photos of license plates (and possibly drivers) *only in case anything goes wrong*. Alternatively, there might be barriers in a toll-plaza environment which are controlled by the RSU.

**1.5.1 Main Protocols.** In the following, we sketch the main protocols of the system involving the parties from above. Fig. 1 provides an overview of these interactions. A more detailed description that also includes the remaining protocols can be found in Section 4. For simplicity, let us envision a system with (e.g., monthly) billing periods, although this is not mandatory but suggested for our framework.

*User Registration:* To participate in the system, a user needs to register with the TSP using some physical ID (e.g., passport number, SSN) which is verified out-of-band. This is done once and makes the user accountable in case he cheats or refuses to pay his bill.

*Wallet Issuing:* A wallet is issued to a user by the TSP which is bound to the user's ID and a set of user attributes (discussed later on). The wallet is used to accumulate debt.

*Debt Accumulation:* Every time a user passes an RSU, his OBU and the RSU execute the Debt Accumulation protocol. First, the toll the user owes is determined and then this toll is added to the user's wallet—possibly along with some public attributes of the RSU. The toll may be dynamic and depend on different factors like the current time and congestion, the number of axles (recognized by sensors attached to the RSU), some user attributes attached to the wallet as well as some attributes of the previous RSU the user drove by.

The camera takes a photo of the license plate(s) in case the RSU reports any protocol failure, aborts (e.g., if a wallet is outdated) or a car in the range of the RSU refuses to communicate at all. Due to technical limitations, it might not always be possible to determine exactly which car triggered the camera, especially in open-road tolling environments [46]. In these cases, photos of more than one car being in the range of the RSU are taken and transmitted to the SA.

*Prove Participation:* After the SA has identified the users behind the license plates involved in an incident, it determines who caused the incident. Since we demand that Debt Accumulation transactions are anonymous, the users need to interact with the SA to help with this matching. To this end, an instance of the Prove

Participation protocol is executed between the SA and each of these users consecutively. This prevents honest users who successfully ran Debt Accumulation from being penalized.

*Debt Clearance:* At the end of a billing period, a user is requested to participate in a (asynchronous) clearance phase with the TSP. As the user is not anonymous in this protocol, he can be penalized if he refuses to run the protocol within a certain time frame. In a protocol execution, he presents his wallet and the accumulated debt to the TSP. Then he may clear his debt immediately or within a certain grace period. After a successful protocol execution, his wallet is invalid. He can get a new one by running Wallet Issuing again.

*1.5.2 Attributes, Pricing Function and Privacy Leakage.* Our scenario involves two types of attribute vectors: user attributes as well as (previous) RSU attributes. To keep our framework flexible, we do not stipulate which kind of attributes or how many of them are used. Those details depend on the concrete pricing model with a pricing function depending on the user attributes, the current and previous RSU attributes and auxiliary, publicly verifiable input. However, we expect that for most scenarios very little information needs to be encoded into these attributes. For instance, one could realize access-based toll collection with a pricing function primarily depending on auxiliary input like location, time and congestion without any previous RSU attributes and only encoding the current billing period as a user attribute. Including the previous RSU's attributes into the toll calculation also allows to cover distance-based charging, where RSUs are installed at each entry and exit of a highway. Running Debt Accumulation at the Entry-RSU does not add any debt to the wallet but only encodes the previous RSU's ID as attribute. At the Exit-RSU, the appropriate toll is calculated and added but no RSU attribute is set.<sup>2</sup> To mitigate the damage of a stolen RSU, one might want RSUs to have a common "expiration date" which is periodically renewed and encoded as an RSU attribute. Likewise, to enforce that a user eventually pays his debt, the user attributes should at least encode the billing period they are valid for.<sup>3</sup>

Obviously, the concrete content of the attributes affects the "level" of user privacy in an instantiation of our system. Our goal is to provide provable privacy up to what can be possibly be deduced by an operator who *explicitly* learns (1) those attributes as part of its input to the pricing function and (2) the total debt of a user at the end of a billing period. Our framework guarantees that protocol runs of honest users do not leak anything (useful) beyond that. In [Appendix A](#), we analyze the privacy leakage of our system in more detail.

In order to allow users to assess the privacy of a particular instantiation of our framework, we assume that all attributes, all possible values for those attributes and how they are assigned, as well as the pricing function are public. In this way, the TSP is also discouraged from running trivial attacks by tampering with an individual's attribute values (e.g., by assigning a billing period value not assigned to any other user). To this end, a user needs to check if the assigned attribute values appear reasonable. Such checks could also be conducted (at random) by a regulatory authority or often also automatically by the user's OBU. Likewise, if a (corrupted) RSU tries to break privacy by charging a peculiar price that differs from the prescribed pricing function, the user is assumed to file a claim.

*1.5.3 Desired Properties.* The following list summarizes the desired security properties we had in mind in an informal manner. In [Section 2](#) we then propose an ideal world model formally covering what we achieve.

- (P1) A user may only use a wallet legitimately issued to him.
- (P2) In Debt Accumulation, a user may not pretend that he owes less by forging the attributes attached to his wallet.
- (P3) In Debt Clearance, a user should not be able to claim that he owes less than the amount added to his wallet, provided he never used an old copy of his wallet (double-spending) or aborted Debt Accumulation.

<sup>2</sup>In this way, the entry point of a user's trip can be linked to his exit point. However, our system ensures that the user is still anonymous and multiple entry/exit pairs are unlinkable.

<sup>3</sup>Clearly, for privacy reasons, unique expiration dates in attributes need to be avoided.

- (P4) If a user reuses an old copy of his wallet (double-spending) he will be identified.
- (P5) If a user fails to participate in or prematurely aborts Debt Accumulation he will be identified.
- (P6) A Debt Accumulation transaction of an honest user does not leak anything to an RSU/TSP beyond the user’s attributes and the attributes of the previous RSU.
- (P7) Prove Participation only enables the SA to deanonymize a single transaction of an honest user in case of an incident involving this user. His remaining transactions stay unlinkable.
- (P8) The TSP is—with a hint from the DR—able to efficiently blacklist wallets of individual users. This is important in practice, e.g., to mitigate the financial loss due to stolen or compromised OBUs or double-spending.
- (P9) The TSP is—with a hint from the DR—able to efficiently recalculate the debt for individual users during a billing period. This is important in practice, e.g., to mitigate the financial loss due to broken, stolen, or compromised OBUs. Furthermore, it allows to determine the actual debt of a double-spender. Also, in a dispute, a user may request a detailed invoice listing the toll points he visited and the amounts being charged.
- (P10) As the secrets of an RSU might enable a user to tamper with his debt, there is a mechanism to mitigate the financial loss due to a stolen or compromised RSU.

## 1.6 Our Contribution

Our contribution is threefold:

*1.6.1 Protocols.* While the overall idea underlying our construction might appear intuitive and we partly draw from known techniques, a major challenge was to twist and combine all these techniques to achieve simulation-based security and practicality *at the same time*.

To this end, we start from a payment system building-block called black-box accumulation (BBA+), recently introduced in [43]. BBA+ offers the core functionality of an unlinkable user wallet maintaining a balance. Values can be added to and subtracted from the wallet by an operator, where the use of an old wallet is detected offline by a double-spending mechanism. Besides unlinkability, the system guarantees that a wallet may only be used by its legitimate owner and with its legitimate balance. While BBA+ provides us with some ((P1), (P3), (P4) and (P6) partially) of the desired properties identified in subsection 1.5.3, significant modifications and enhancements had to be made to fully suit our needs.

For instance, the basic BBA+ mechanism does not allow for efficient blacklisting of individual wallets (P8) nor to recalculate individual balances (P9). We solve this by having an individual trapdoor for each wallet (accessible to the DR in case of an incident) which makes transactions involving this wallet forward and backward linkable. The trapdoor does not allow to link wallets of other users.<sup>4</sup> To realize this, we adopt and adjust an idea from the e-cash literature [18]. More precisely, we make use of a PRF applied to a counter value (which is bound to a wallet) to generate some fraud detection ID for a wallet state. To ensure security and privacy, we let both the user and the TSP jointly choose the PRF key with the key remaining unknown to the TSP. To make it accessible in case of an incident, the user is forced to deposit the key encrypted under the DR’s public key. The correctness of this deposit is ensured to the TSP by means of a NIZK proof. This part is tricky due to the use of Groth-Sahai NIZKs for efficiency reasons and the lack of a compatible (i.e., algebraic) encryption scheme with message space  $\mathbb{Z}_p$ . By letting a user choose a new PRF key for each billing period, his transactions from previous/upcoming periods stay unlinkable.

As a minor but very useful modification, we added (user and RSU) attributes to a wallet, which, of course, get signed along with the wallet to protect from forgery. This allows us to bind wallets to a billing period encoded as attribute. By making RSUs only accept wallets from the current period, the size of the blacklist checked by the

<sup>4</sup>To be precise, in the BBA+ system, there is a TTP-owned trapdoor that allows to link a user’s transactions. However, this is not a user- or wallet-specific trapdoor, which means, that if handed to the TSP, it could track each and every user in the system.



RSU can be limited to enable fast transactions. Similarly, the database needed to recalculate balances can be kept small.

Another problem is the use of a single shared wallet certification key in the BBA+ scheme. Translated to our setting, the TSP and all RSUs would share the same secret key. Hence, if an adversary corrupted a single RSU, he could create new wallets for fake users, forge user attributes, balances, etc. In order to mitigate this problem (P10), we take the following measures: First, we separate user identity and attribute information, i.e., the fixed part of a wallet, from balance information, i.e., the updatable part. The first part is signed by a signature key only held by the TSP when the wallet is issued. The second part is signed by an RSU each time the balance of a wallet is updated. This prevents a corrupted RSU to issue new wallets or fake user attributes. Furthermore, the individual key of each RSU is certified by the TSP along with its attributes. In this way, a RSU may not forge its attributes (P2) but may still fake the balance. By including an expiration date into the RSU attributes, one can limit the potential damage involved with the latter issue. In view of the fact that RSUs are usually not easily accessible physically and key material is usually protected by a HSM, we believe that these measures are sufficient. We intentionally refrain from using key revocation mechanisms like cryptographic accumulators [56] in order to retain real-time interactions.

Finally, we added mechanisms to prove the participation in Debt Accumulation interactions (P7) and to enable a simulation-based security proof.

*1.6.2 System Definition, Security Model and Proof.* Having the scenario from Section 1.5 at the back of our mind, we propose a system definition and security model for post-payment toll collection systems based on the real/ideal-paradigm. More precisely, we build upon the GUC framework [23].

Our work is one of very few that combines a complex, yet practical crypto system with a thorough UC security analysis. Typically, the standard approach is to cast a complex system as an MPC problem and then resort to a *generic but inefficient* UC-secure MPC protocol [24, 47].

The security of BBA+ has been modeled by formalizing each security property from a list of individual properties as it is usually done in the game-based setting. This approach bears the intrinsic risk that important and expedient security aspects are overlooked, e.g., the list is incomplete. This danger is eliminated by our UC-approach where we do not aim to formalize a list of individual properties but rather how an ideal system should look like.

A challenging task was to find a formalization of such an ideal system (aka ideal functionality) such that a reasonable trade-off between various aspects is accomplished: On the one hand, it needs to be sufficiently abstract to represent the semantics of the eventual goal “toll collection” while it should still admit a realization. On the other hand, keeping it too close to the concrete realization and declaring aspects as out-of-scope only provides weak security guarantees.

We decided to directly model the ETC system as a single functionality with (polynomially) many parties that reactively participate in (polynomially) many interactions. This leads to a clean interface but makes the security analysis highly non-trivial. At first sight, it seems tempting to follow a different approach: Consider the system as a composition of individual two-party protocols, analyze their security separately and argue about the security of the whole system using the UC composition theorem. We refrain from this approach, however, as it entails a slew of technical subtleties due to the shared state between the individual two-party protocols.

Moreover, although our system uses cryptographic building blocks for which UC formalizations exist (commitments, signatures, NIZK), these abstractions cannot be used. For example, UC-signatures are just random strings that are information-theoretic independent of the message they sign. Thus it is impossible to prove in zero-knowledge any statement about message-signature-pairs. Hence, our security proof has to start almost from scratch. Although parts of it are inspired by proofs from the literature, it is very complex and technically demanding in its entirety.

*1.6.3 Implementation.* In addition to our theoretical framework, we also evaluate the real-world practicality of P4TC by means of an implementation. We specifically benchmarked our protocols on an embedded processor which is known to be used in currently available OBUs such as the Savari MobiWAVE [64]. The major advantage for real-world deployment originates in the use of non-interactive zero-knowledge proofs, where major parts of the proofs can be precomputed and verification equations can be batched efficiently. This effectively minimizes the computations which have to be performed by the OBU and the RSU during an actual protocol run. Our implementation suggests that provably-secure ETC at full speed can indeed be realized using present-day hardware.

## 2 SECURITY MODEL

Our toll collection system  $\mathcal{F}_{\text{P4TC}}$  is defined within the (G)UC-framework by Canetti [23], which is a simulation-based security notion. We briefly explain the UC-framework in this section; for a more detailed introduction to UC and its variant GUC see [20, 23].

### 2.1 Introduction to UC

In the UC-framework security is defined by indistinguishability of two experiments: the *ideal experiment* and the *real experiment*. In the ideal experiment the task at hand is carried out by dummy parties with the help of an ideal incorruptible entity—called the ideal functionality  $\mathcal{F}$ —which plainly solves the problem at hand in a secure and privacy preserving manner. In the real experiment the parties execute a protocol  $\pi$  in order to solve the prescribed tasks themselves. A protocol  $\pi$  is said to be a (secure) *realization*  $\mathcal{F}$  if no PPT-machine  $\mathcal{Z}$ , called the *environment*, can distinguish between two experiments.

*2.1.1 The Basic Model of Computation.* The basic model of computation consists of a set of instances (ITIs) of interactive Turing machines (ITMs). An ITM is the description of a Turing machine with additional tapes for its identity, for subroutine input and output and for incoming and outgoing network messages. An ITI is a tangible instantiation of an ITM and is identified by the content of its identity tape. The order of activation of the ITIs is message-driven. If the ITI provides subroutine output or writes an outgoing message, the activation of the ITI completes and the ITI to whom the message has been delivered to gets activated next. Each experiment has two special ITIs: the environment  $\mathcal{Z}$  and the adversary  $\mathcal{A}$  (in the real experiment) or the simulator  $\mathcal{S}$  (in the ideal experiment). The environment  $\mathcal{Z}$  is the ITI that is initially activated. If the environment  $\mathcal{Z}$  provides subroutine output, the whole experiments stops. The output of the experiment is the output of  $\mathcal{Z}$ .

*2.1.2 The (Dummy) Adversary.* The adversary  $\mathcal{A}$  is instructed by  $\mathcal{Z}$  and represents  $\mathcal{Z}$ 's interface to the network, e.g.,  $\mathcal{A}$  reports all network messages generated by any party to  $\mathcal{Z}$  and can manipulate, reroute, inject and/or suppress messages on  $\mathcal{Z}$ 's order. This modeling reflects the idea of an unreliable and untrusted network. Please note: Only incoming/outgoing messages are under the control of  $\mathcal{A}$ . Subroutine input/output is authenticated, confidential and of integrity. Moreover,  $\mathcal{Z}$  may instruct  $\mathcal{A}$  to corrupt a party. In this case,  $\mathcal{A}$  takes over the role of the corrupted party, reports its internal state to  $\mathcal{Z}$  and from then on may arbitrarily deviate from the protocol  $\pi$  in the name of the corrupted party as requested by  $\mathcal{Z}$ .

*2.1.3 The Real Experiment.* In the real experiment, denoted by  $\text{EXEC}_{\pi, \mathcal{A}, \mathcal{Z}}(1^n)$ , the environment  $\mathcal{Z}$  interacts with parties running the actual protocol  $\pi$  and is supported by a real adversary  $\mathcal{A}$ . The environment  $\mathcal{Z}$  specifies the input of the honest parties, receives their output and determines the overall course of action.

*2.1.4 The Ideal Experiment.* In the ideal experiment, denoted by  $\text{EXEC}_{\mathcal{F}, \mathcal{S}, \mathcal{Z}}(1^n)$ , the protocol parties are mere dummies that pass their input to a trusted third party  $\mathcal{F}$  and hand over  $\mathcal{F}$ 's output as their own output. The ideal functionality  $\mathcal{F}$  executes the task at hand in a trustworthy manner and is incorruptible. The real adversary  $\mathcal{A}$

is replaced by a simulator  $\mathcal{S}$ . The simulator must mimic the behavior of  $\mathcal{A}$ , e.g., simulate appropriate network messages (there are no network messages in the ideal experiment), and come up with a convincing internal state for corrupted parties (dummy parties do not have an internal state).

**2.1.5 Definition of Security.** A protocol  $\pi$  is said to securely UC-realize an ideal functionality  $\mathcal{F}$ , denoted by  $\pi \geq_{\text{UC}} \mathcal{F}$ , if and only if

$$\exists \mathcal{S} \forall \mathcal{Z} : \text{EXEC}_{\pi, \mathcal{A}, \mathcal{Z}}(1^n) \stackrel{c}{\equiv} \text{EXEC}_{\mathcal{F}, \mathcal{S}, \mathcal{Z}}(1^n)$$

holds whereby the randomness is taken over the initial input of  $\mathcal{Z}$  and the random tapes of all ITIs.<sup>5</sup> If no environment  $\mathcal{Z}$  can tell executions of the real and the ideal experiment apart, then any successful attack existing in the real experiment would also exist in the ideal experiment. Therefore the real protocol  $\pi$  guarantees the same level of security as the (inherently secure) ideal functionality  $\mathcal{F}$ .

## 2.2 Remarks on Privacy

Regarding privacy, please note that all parties (including the simulator) use the ideal functionality as a black-box and only know what the ideal functionality explicitly allows them to know as part of their prescribed output. The output to the simulator is also called leakage. This makes UC suitable to reason about privacy in a very nice way. As no additional information is unveiled, the achieved level of privacy can directly be deduced from the defined output of the ideal functionality. In other words, the privacy assessment can be conducted onto the ideal functionality and is completely decoupled from the analysis of the protocol implementation. The proof of indistinguishability asserts that any secure realization of the functionality provides the same level of privacy.

## 2.3 UC Model Conventions

The bare UC model does not specify many important aspects. For example, it leaves open which ITIs are allowed to communicate with each other, how parties are corrupted and which ITI is allowed to invoke what kind of new ITIs. In this section we clarify these aspects. Our conventions are probably the mostly used one and quite natural.

- Each party is identified by its party identifier (PID)  $pid$  which is unique to the party and is the UC equivalent of the physical identity of this party. A party runs a protocol  $\pi$  by means of an ITI which is called the main party of this instance of  $\pi$ . An ITI can invoke subsidiary ITIs to execute sub-protocols. A subsidiary and its parent use their *subroutine input/output tapes* to communicate with each other. The set of ITIs taking part in the same protocol but for different parties communicate through their *message tapes*. An instance of a protocol is identified by its session identifier (SID)  $sid$ . All ITIs taking part in the same protocol instance share the same SID. A specific ITI is identified by its ID  $id = (pid, sid)$ .
- Ideal functionalities are an exception to this rule. An instance of an ideal functionality  $\mathcal{F}$  owns a SID but no PID. Input to and output from  $\mathcal{F}$  is performed through dummy parties which share the same SID as  $\mathcal{F}$ , but additionally have the PIDs of the parties they belong to.
- As already stated in [subsection 2.1.2](#) subroutine input/output is identifying. This implies in particular that an ideal functionality knows the  $pid$  of the party to/from whom it writes/reads input/output, because an ideal functionality communicates with its calling ITI by means of intermediary dummy ITIs and subroutine input/output.
- We assume PID-wise corruption: either all ITIs of party are corrupted or none.

All these conventions capture the natural intuition that ITIs with the same PID represent processes within the same physical computer and therefore can communicate with each other directly. But communication between

<sup>5</sup>N.b.: To streamline this short introduction we directly defined the so-called dummy adversary. The original definition all-quantifies over all adversaries in the first step. In the second step it is shown that the dummy adversary—as defined here—is the most severe one and complete.

different physical computers (i.e., parties) is only possible through an unreliable network under the control of an adversary.

#### 2.4 Setup Assumptions — CRS $\mathcal{F}_{\text{CRS}}$ and Bulletin-Board $\overline{\mathcal{G}}_{\text{bb}}$

As commonly found in the UC setting, we also draw from constitutive trust assumptions, or setup assumptions in the UC terminology. Setup assumptions are ideal functionalities that remain ideal in the real experiment, e.g., their secure realization is left unspecified. In our scenario, we assume a common reference string (CRS), denoted  $\mathcal{F}_{\text{CRS}}$ , and a globally available bulletin board,  $\overline{\mathcal{G}}_{\text{bb}}$  [25, Fig. 3], which is sometimes also referred to as key registration service in the literature [8, 22].

A CRS is a short piece of information that has been honestly generated and is shared between all parties. A bulletin board can be depicted as the abstract formalization of a globally available key registration service which associates (physical) party identifiers (PIDs) with (cryptographic) public keys. Every party has the option to register a key for itself once and  $\overline{\mathcal{G}}_{\text{bb}}$  ensures that the registering party cannot lie about its (physical) identity *pid*. In our scenario the PID of users could be a passport number or SSN. For RSUs the geo-location could be used as a PID. The key can be any value  $v$ , i.e.,  $\overline{\mathcal{G}}_{\text{bb}}$  is totally agnostic to the data. Moreover, every party can obtain keys that have been registered by any other party in a trustworthy way. Our augmentations to  $\overline{\mathcal{G}}_{\text{bb}}$  are only syntactical bagatelles. Parties can also perform a reverse lookup, i.e., lookup the PID of the party that registered a value, and parties can not only register an opaque string of bits but an ordered tuple of strings. The latter is required to support reverse searches on substrings.

#### 2.5 Generalized UC

In the plain UC model the environment  $\mathcal{Z}$  is only allowed to invoke a single session of the challenge protocol and is only allowed to exchange messages with the respective main parties. This implies that many real-world scenarios cannot be appropriately captured with plain UC. For example, a globally available PKI that also provides its services to other protocols cannot be modeled in UC. Generalized UC [23] and its simplification Externalized UC (EUC) relaxes this restriction such that a particular, designated ITI—the global functionality—can be accessed by any other ITI including the environment. In our setting, the bulletin-board  $\overline{\mathcal{G}}_{\text{bb}}$  is allowed to be global.

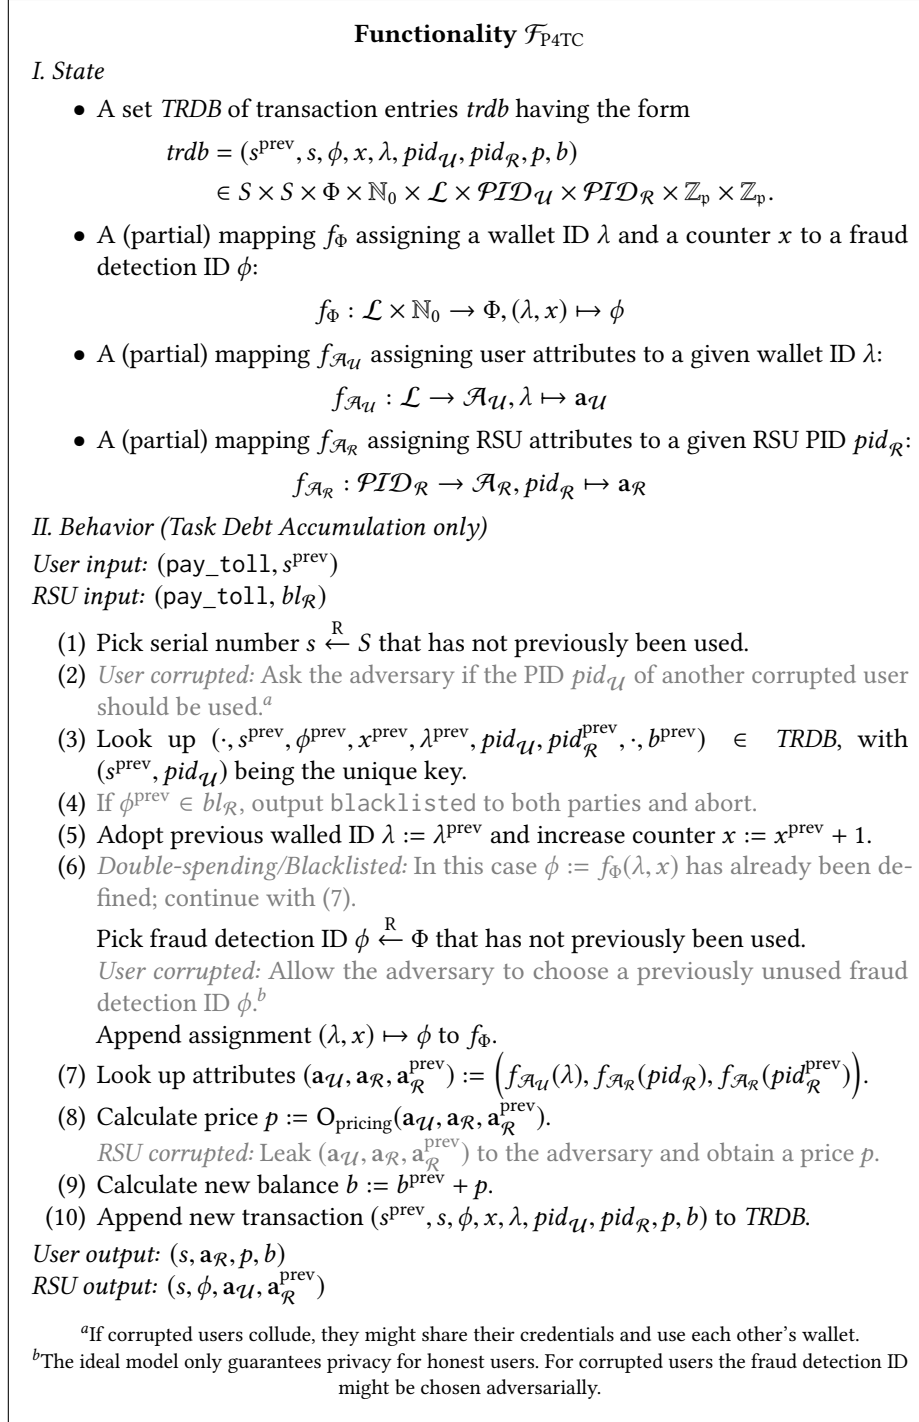
### 3 SYSTEM DEFINITION

We now give a condensed description of our ideal functionality  $\mathcal{F}_{\text{P4TC}}$  and point out how it ensures the desired properties given in [subsection 1.5.3](#). We do not formalize each task (e.g., Wallet Issuing, Debt Accumulation, ...) as an individual ideal functionality, but the whole system as a monolithic, highly reactive functionality  $\mathcal{F}_{\text{P4TC}}$  with polynomially many parties as users and RSUs. This allows for a shared state between the individual interactions and to define correctness and security more easily.

An excerpt of  $\mathcal{F}_{\text{P4TC}}$  is depicted in [Fig. 2](#). For better understanding we refrain from giving a complete description of all the tasks  $\mathcal{F}_{\text{P4TC}}$  provides at this point. Instead we focus on the most important task, namely Debt Accumulation, and only sketch the remaining tasks. We also state why the ideal model reflects the security and privacy level we wish to achieve. Again, we only concentrate on Debt Accumulation. The full-fledged ideal functionality can be found in [Appendix B](#).

For the ease of presentation, all instructions in [Fig. 2](#) that are typically executed are printed in normal font, while some conditional side tracks are grayed out. The normal execution path defines the level of security and privacy achieved for honest, well-behaving parties. The conditional branches deal with corrupted, misbehaving,<sup>6</sup> or blacklisted parties. The key idea behind  $\mathcal{F}_{\text{P4TC}}$  is to keep track of all conducted transactions in a pervasive

<sup>6</sup>Please note, that users do not need to be formally corrupted in order to commit double-spending. We call these users honest, but misbehaving.

Fig. 2. An excerpt of the ideal functionality  $\mathcal{F}_{\text{P4TC}}$ .

transaction database  $TRDB$  (cp. Fig. 2). Each transaction entry  $trdb \in TRDB$  is of the form

$$trdb = (s^{\text{prev}}, s, \phi, x, \lambda, pid_{\mathcal{U}}, pid_{\mathcal{R}}, p, b).$$

It contains the identities  $pid_{\mathcal{U}}$  and  $pid_{\mathcal{R}}$  of the involved user and RSU (or TSP) respectively,<sup>7</sup> the price  $p$  (toll) associated with this particular transaction and the total balance  $b$  of the user’s wallet, i.e., the accumulated sum of all prices so far including this transaction. In other words,  $\mathcal{F}_{\text{P4TC}}$  implements a trustworthy global bookkeeping service that manages the wallets of all users. Each transaction entry is uniquely identified by a serial number  $s$ . Additionally, each entry contains a pointer  $s^{\text{prev}}$  to the logically previous transaction, a unique wallet ID  $\lambda$ , a counter  $x$  indicating the number of previous transactions with this particular wallet, and a fraud detection ID  $\phi$ .

Some explanations are in order with respect to the different IDs. In a truly ideal world  $\mathcal{F}_{\text{P4TC}}$  would use the user identity  $pid_{\mathcal{U}}$  to look up its most recent entry in the database and append a new entry. Such a scheme, however, could only be implemented by an online system. Since we require offline capabilities—allowing a user and RSU to interact without the help of other parties and without permanent access to a global database—the inherent restrictions of such a setting must be reflected in the ideal model:

- (Even formally honest) users can misbehave, reuse old wallet states and commit double-spending without being noticed *instantly*.
- Double-spending is eventually detected after-the-fact.

In order to accurately define security, these technicalities have to be incorporated into the ideal functionality, which causes the bookkeeping to be more involved. The whole transaction database is best depicted as a directed graph. Nodes are identified by serial numbers  $s$  and additionally labeled with  $(\phi, x, \lambda, pid_{\mathcal{U}}, b)$ . Edges are given by  $(s^{\text{prev}}, s)$  and additionally labeled with  $(pid_{\mathcal{R}}, p)$ . A user’s wallet is represented by the subgraph of all nodes with the same wallet ID  $\lambda$  and forms a connected component. As long as a user does not commit double-spending the particular subgraph is a linked, linear list. In this case, each transaction entry has a globally unique fraud detection ID  $\phi$ . If a user misbehaves and reuses an old wallet state (i.e., there are edges  $(s^{\text{prev}}, s)$  and  $(s^{\text{prev}}, s')$ ), the corresponding subgraph degenerates into a directed tree. In this case, all transaction entries that constitute a double-spending, i.e., all nodes with the same predecessor, should share the same fraud detection ID  $\phi$ . To this end, the counter  $x$  and the injective map  $f_{\phi} : \mathcal{L} \times \mathbb{N}_0 \rightarrow \Phi$  have been introduced in order to manage fraud detection IDs consistently: For any newly issued wallet with ID  $\lambda$ , the counter  $x$  starts at zero and  $x = (x^{\text{prev}} + 1)$  always holds. It counts the number of subsequent transactions of a wallet since its generation, i.e.,  $x$  equals the depth of a node. The function  $f_{\phi}$  maps a transaction to its fraud detection ID  $\phi$ , given its wallet (aka tree)  $\lambda$  and its depth  $x$ .

Besides storing transaction data,  $\mathcal{F}_{\text{P4TC}}$  also keeps track of parties’ attributes by internally storing RSU attributes  $a_{\mathcal{R}}$  upon certification and user (or rather wallet) attributes  $a_{\mathcal{U}}$  when the wallet is issued. To give a better understanding of how  $\mathcal{F}_{\text{P4TC}}$  works, we explain the task Debt Accumulation in more detail.

### 3.1 Task “Debt Accumulation”

In this task, the user provides a serial number  $s^{\text{prev}}$  as input to  $\mathcal{F}_{\text{P4TC}}$ , indicating which past wallet state he wishes to use for this transaction. The participating RSU inputs a list of fraud detection IDs that are blacklisted.

Firstly,  $\mathcal{F}_{\text{P4TC}}$  randomly picks a fresh serial number  $s$  for the upcoming transaction. If (and only if) the user is corrupted,  $\mathcal{F}_{\text{P4TC}}$  allows the simulator to provide a different value for  $pid_{\mathcal{U}}$  that belongs to a another corrupted user. This is a required technicality as corrupted users might share their credentials and thus might use each

<sup>7</sup>The party identifier (PID) can best be depicted as the model’s counterpart of a party’s physical identity in the real world. E.g., the user’s physical identity could be a passport number or SSN; the “identity” of an RSU could be its geo-location. Generally, there is no necessary one-to-one correspondence between a PID and a cryptographic key. Also, the ideal functionality always knows the PID of the party it interacts with by the definition of the UC framework.

other's wallet. Please note that this does not affect honest users.  $\mathcal{F}_{P4TC}$  looks up the previous wallet state  $trdb^{prev}$  in  $TRDB$  and asserts that its fraud detection ID is not blacklisted; otherwise it aborts. The ideal functionality extracts the wallet ID from the previous record ( $\lambda := \lambda^{prev}$ ) and increases the counter for this particular wallet ( $x := x^{prev} + 1$ ). Then, it checks if a fraud detection ID  $\phi$  has already been defined for the wallet ID  $\lambda$  and counter  $x$  (n.b.: tree and depth of node). If so, the current transaction record will be assigned the same fraud detection ID  $\phi$ . Otherwise,  $\mathcal{F}_{P4TC}$  ties a fresh, uniformly and independently drawn fraud detection ID ( $(\lambda, x) \mapsto \phi$ ) to the  $x$ 'th transaction of the wallet  $\lambda$ . If and only if the user is corrupted and  $\phi$  has not previously been defined, the adversary is allowed to overwrite the fraud detection ID with another value.<sup>8</sup> Moreover,  $\mathcal{F}_{P4TC}$  looks up the user's attributes bound to this particular wallet ( $\mathbf{a}_U := f_{\mathcal{A}_U}(\lambda)$ ) and the attributes of the current and previous RSU ( $\mathbf{a}_R := f_{\mathcal{A}_R}(pid_R)$ ,  $\mathbf{a}_R^{prev} := f_{\mathcal{A}_R}(pid_R^{prev})$ ). Finally, the ideal functionality queries the pricing oracle  $O_{pricing}$  for the price  $p$  of this transaction, calculates the new balance of the wallet ( $b := b^{prev} + p$ ) and appends a new record to the transaction database.

If and only if the RSU is corrupted, the adversary learns the involved attributes and is allowed to override the price. Please note that leaking the user/RSU attributes to the adversary does not weaken the privacy guarantees as the (corrupted) RSU learns the attributes as an output anyway. The option to manipulate the price on the other hand was a design decision. It was made to enable implementations in which the pricing function is unilaterally evaluated by the RSU and the user initially just accepts the price. It is assumed that a user usually collects any toll willingly in order to proceed and (in case of a dispute) files an out-of-band claim later (before making the actual payment).

The user's output are the serial number  $s$  of the current transaction, the current RSU's attributes  $\mathbf{a}_R$ , the price  $p$  to pay and the updated balance  $b$  of his wallet. The RSU's output are the serial number  $s$  of the current transaction, the fraud detection ID  $\phi$  and the attributes  $\mathbf{a}_U$  and  $\mathbf{a}_R^{prev}$  of the user and the previous RSU, respectively. Learning their mutual attributes is necessary, because the RSU and user must evaluate the pricing function themselves in the real implementation without the help of a third party.<sup>9</sup>

### 3.2 Correctness and Operator Security

Both are immediately asserted by the ideal functionality. As  $\mathcal{F}_{P4TC}$  represents an incorruptible bookkeeper, the user has no chance to cheat. The only input of a (possibly malicious) user is his choice of which previous wallet state should be used. Hence, the user has no chance to lie about his attributes  $\mathbf{a}_U$ , the balance  $b^{prev}$  associated with this state, the calculated new balance  $b$  nor anything else. If the user is malicious, the adversary has the additional options to choose an alternative (malicious) user that is being charged and to choose a different fraud detection ID. Both do not effect operator security. The former does not change the total amount due to the operator. It only changes the party liable which is unavoidable, if a group of corrupted users collude. The latter does not affect the operator at all.

### 3.3 User Security and Privacy Leakage

User security follows by the same arguments as for operator security.

The information leakage that needs to be considered for an assessment of user privacy directly follows from the in- and output of the ideal functionality. At this point, we exemplarily consider the leakage of Debt Accumulation only. We refer to [Appendix A](#) for an overview of the leakage of the remaining tasks.

<sup>8</sup>Again, this is a technical concession to the security proof. Corrupted users are not obliged to use "good" randomness. This might affect untrackability, but we do not aim to provide this guarantee for corrupted users.

<sup>9</sup>At least, it is unavoidable for any *practical* implementation given our scenario. Due to the offline capabilities there is no third party available. Hence, the only alternative approach would be to use some generic SFE/2PC techniques which are prohibitively inefficient for an open-road ETC system using low-end hardware. Moreover, the legal framework of most countries requires that the value being charged can be checked manually.

We stress that we only care about privacy for honest, well-behaving, non-blacklisted<sup>10</sup> users. Hence, the grayed out steps can be ignored. First note that the serial number of the previous transaction  $s^{\text{prev}}$  is a private input of the user and never output to any party. The RSU only learns the serial number  $s$  and the fraud detection ID  $\phi$  of the current transaction which are both freshly, uniformly and independently drawn by  $\mathcal{F}_{\text{P4TC}}$ . Hence, it is *information-theoretically impossible* to track the honest and well-behaving user across any pair of transactions using any of these numbers. The only “real” information leakage (to the RSU) in this task is determined by the user’s and the previous RSU’s attributes  $\mathbf{a}_U$  and  $\mathbf{a}_R^{\text{prev}}$ .

Independence of the fraud detection ID might be lost for two cases:

- The user is corrupted and the adversary playing a corrupted user knows all his secrets anyway. In this case, the adversary is allowed to pick the fraud detection ID if not previously defined.
- The mapping  $(\lambda, x) \mapsto \phi$  is already defined, because the user repeatedly runs Debt Accumulation with the same inputs and thus tries to attempt double-spending or the user has been blacklisted. (Our blacklisting task—sketched later—prefills  $f_\phi$  and creates a pool of fraud detection IDs to be used.)

In neither of these cases we aim to provide untrackability.

### 3.4 Remaining Tasks

We briefly sketch the remaining tasks of the ideal functionality omitting some details regarding corrupted parties. Full details are given in [Appendix B](#).

*Registration and Certification:* There are auxiliary tasks for registration/certification of parties. They are modeled in the obvious way and append the appropriate mappings (e.g.,  $f_{\mathcal{A}_R}$ ) with the given attributes for the PIDs.

*Wallet Issuing:* In Wallet Issuing a new wallet ID  $\lambda$  is freshly, uniformly and independently drawn. A new transaction entry for the particular user is inserted into the database using this  $\lambda$  and a zero balance. This entry can be depicted as the root node of a new wallet tree.

*Debt Clearance:* The task Debt Clearance is very similar to Debt Accumulation described above except that the TSP additionally learns the user’s party ID  $pid_U$  and the wallet balance  $b$ . Debt Clearance is identifying for the user to allow the operator to invoice him and check if he (physically) pays the correct amount. Also, the user does not obtain a new serial number such that this the transaction entry becomes a leaf node of the wallet tree.

*User Blacklisting:* The task User Blacklisting is run between the DR and TSP. The TSP inputs the PID of a user it wishes to blacklist and obtains the debt the user owes and a list of upcoming serial numbers. For the latter,  $\mathcal{F}_{\text{P4TC}}$  draws a sequence of fresh, uniform and independent fraud detection IDs  $\phi_i$  and prefills the mapping  $f_\phi$  for all wallets the user owns. This ensures that upcoming transactions use predetermined fraud detection IDs that are actually blacklisted.

*Prove Participation:* The task Prove Participation simply checks if a record exists in the database for the particular user and serial number.

*Double-Spending Detection:* The task Double-Spending Detection checks if the given fraud detection ID exists multiple times in the database, i.e. if double-spending has occurred with this ID. If so,  $\mathcal{F}_{\text{P4TC}}$  leaks the identity of the corresponding user to the adversary and asks the adversary to provide an arbitrary bit string that serves as a proof of guilt.  $\mathcal{F}_{\text{P4TC}}$  outputs both—the user’s identity and the proof—to the TSP and records the proof as valid for the user. Please note, correctness, completeness and soundness is guaranteed by the ideal bookkeeping service.

*Guilt Verification:* The task Guilt Verification checks if the given proof is internally recorded as being issued for the particular user.

<sup>10</sup>Note that the TSP may only blacklist users with the help of the incorruptible DR who only cooperates if the user agreed or misbehaved.



### 3.5 Mapping of the Desired Properties to the Ideal Model

In this section, we illustrate how the definition of our ideal model  $\mathcal{F}_{\text{P4TC}}$  ensures the desired properties of a toll collection scheme (cp. [subsection 1.5.3](#), (P1)-(P10)).

- (P1) to (P3): In the tasks of Debt Accumulation and Debt Clearance, the user only ever inputs the serial number of the wallet state he wants to use. All relevant information is then looked up internally by  $\mathcal{F}_{\text{P4TC}}$ . Therefore, the user is not able to pretend his wallet state contains any other information than it actually does (like different attributes in Debt Accumulation or a lower balance in Debt Clearance).  $\mathcal{F}_{\text{P4TC}}$  also checks that the wallet actually belongs to this user. Only corrupted users are able to use wallets of one another, but not wallets of honest users.
- (P4): As explained above, the same fraud detection ID  $\phi$  occurs in multiple transaction entries if and only if double-spending was committed. The TSP can check this using the task Double-Spending Detection. If double-spending is detected,  $\mathcal{F}_{\text{P4TC}}$  provides the TSP with the identity  $pid_u$  of the respective user and a publicly verifiable proof  $\pi$  that this user has actually committed double-spending. This prevents the TSP from falsely accusing a user of double-spending.
- (P5) and (P7): As discussed in [Section 1.5](#), we assume that if a user does not properly participate in Debt Accumulation, he is physically identified outside the scope of  $\mathcal{F}_{\text{P4TC}}$ . The important feature  $\mathcal{F}_{\text{P4TC}}$  provides for this property is the task Prove Participation. In this task the SA inputs a set  $S_{\mathcal{R}}^{\text{pp}}$  of serial numbers and the ID of the user in question. With the users consent  $\mathcal{F}_{\text{P4TC}}$  then checks whether the user successfully participated in any of the transactions from this list. If so, it just responds with OK to the SA who does not learn anything about any of the user's transactions beyond that.
- (P6): After an instance of Debt Accumulation, an RSU gets the information  $(s, \phi, a_u, a_{\mathcal{R}}^{\text{prev}})$  as output. As already discussed in [Section 1.5](#), we assume user and RSU attributes to be sufficiently indistinct that they do not enable tracking of the user. This is not ensured within the scope of  $\mathcal{F}_{\text{P4TC}}$ , apart from attribute outputs to the user which enable him to check they are not identifying. Unless double-spending is committed, the serial numbers  $s$  and fraud detection IDs  $\phi$  are fresh and randomly drawn for every transaction and can therefore not compromise unlinkability. In case of double-spending, the TSP is able to directly link transactions with the same fraud detection ID and also find out the ID of the user responsible, but is still only able to link any other transactions of this user with help from the DR.
- (P8): Given a user ID  $pid_u$ , the task User Blacklisting provides the TSP a set of fraud detection IDs  $\phi$  of all wallets  $\lambda$  of this user. In Debt Accumulation, the RSU inputs a blacklist  $bl_{\mathcal{R}}$  (containing fraud detection IDs from the TSP). Then,  $\mathcal{F}_{\text{P4TC}}$  checks whether the fraud detection ID of the wallet state the user wishes to use for the transaction is contained in this list.
- (P9): Within the task of User Blacklisting,  $\mathcal{F}_{\text{P4TC}}$  sums up all prices of all past transactions of all wallets of the user in question and outputs the resulting amount  $b^{\text{bill}}$  to the TSP. As each instance of Debt Clearance results in a transaction entry  $trdb^*$  with  $p^* = -b^{\text{bill}*}$  as the price, correctly cleared and paid wallets cancel out in this sum and the result accurately gives the amount of debt still owed by the user.
- (P10): This is handled outside the scope of  $\mathcal{F}_{\text{P4TC}}$  by encoding a limited time of validity into the RSU's attributes (cp. [subsection 1.5.2](#)).

## 4 PROTOCOL OVERVIEW

For the sake of brevity, this section gives only a *simplified* description of the main protocols illustrating the ideas behind P4TC. For the full protocols see [Appendix D](#).

Before describing our protocols, some remarks about the used crypto building blocks, secure channels and the structure of user wallets are in order.

#### 4.1 Crypto Building Blocks and Algebraic Setting

Our construction makes use of ( $F_{gp}$ -extractable) non-interactive zero-knowledge (NIZK) proofs, equivocal and extractable homomorphic commitments, digital signatures, public-key encryption, and pseudo-random functions (PRFs). The latter building blocks need to be efficiently and securely combinable with the chosen NIZK (which is Groth-Sahai [40] in our case). For readers not familiar with these building blocks, we refer to [Appendix C](#) for a detailed description.

Our system instantiation is based on an asymmetric bilinear group setting  $(G_1, G_2, G_T, e, p, g_1, g_2)$ . Here,  $G_1, G_2, G_T$  are cyclic groups of prime order  $p$  (where no efficiently computable homomorphisms between  $G_1$  and  $G_2$  are known);  $g_1, g_2$  are generators of  $G_1, G_2$ , respectively, and  $e: G_1 \times G_2 \rightarrow G_T$  is a (non-degenerated) bilinear map. We rely on the co-CDH assumption ([Definition C.3](#)) as well as on the security of the building blocks in this setting.

In this section we use a simpler notation than the one introduced in [Appendix C](#) in order to improve the comprehension. Here, a commitment scheme consists of the PPT algorithms  $\text{Com}$  and  $\text{Open}$ , where  $\text{Com}$  returns a commitment  $c$  together with an opening information  $d$  for a given message  $m$  and  $\text{Open}$  verifies if a commitment  $c$  opens to a particular message  $m$  given the opening information  $d$ . A digital signature scheme consists of the PPT algorithms  $\text{Sgn}$  and  $\text{Vfy}$ , where  $\text{Sgn}$  takes as input the secret key  $sk$  and a message  $m \in \mathcal{M}$ , and outputs a signature  $\sigma$  and  $\text{Vfy}$  takes as input the public key  $pk$ , a message  $m \in \mathcal{M}$ , and a purported signature  $\sigma$ , and outputs a bit indicating whether or not the signature is valid.

#### 4.2 Secure (Authenticated) Channels

In our system, all protocol messages are encrypted using CCA-secure encryption. For this purpose, a new session key chosen by the user is encrypted under the public key of an RSU/TSP for each interaction. We omit these encryptions when describing the protocols. Apart from the task Debt Accumulation we furthermore assume all channels to be authenticated.

#### 4.3 Wallets

A user wallet essentially consists of two signed commitments  $(c_{\mathcal{T}}, c_{\mathcal{R}})$ , where  $c_{\mathcal{T}}$  represents the fixed part of a wallet and  $c_{\mathcal{R}}$  the updatable part. Accordingly, the fixed part is signed along with the user's attributes  $a_{\mathcal{U}}$  by the TSP using  $sk_{\mathcal{T}}^{\mathcal{T}}$  during wallet issuing. Every time  $c_{\mathcal{R}}$  is updated, it is signed along with the serial number  $s$  of the transaction by the RSU using  $sk_{\mathcal{R}}$ . The fixed part  $c_{\mathcal{T}} = \text{Com}(\lambda, sk_{\mathcal{U}})^{11}$  is a commitment on the PRF key  $\lambda$  (which is used as wallet ID) and the user's secret key  $sk_{\mathcal{U}}$ . The updatable part  $c_{\mathcal{R}} = \text{Com}(\lambda, b, u_1, x)$  also contains  $\lambda$  (to link both parts), the current balance  $b$  (debt), some user-chosen randomness  $u_1$  to generate double-spending tags for the current state of the wallet, and a counter value  $x$  being the input to the PRF. The value  $\text{PRF}(\lambda, x)$  serves as the wallet's fraud detection ID  $\phi$ . This choice of the fraud detection ID has the advantage that the different versions of a wallet are traceable given  $\lambda$  but untraceable if  $\lambda$  is unknown.

#### 4.4 The P4TC Protocols

We stress again that the following description of the P4TC protocols is a simplified version. The actual protocols, including protocol interaction diagrams, can be found in [Appendix D](#).

**4.4.1 System Setup.** In the setup phase, a trusted third party generates the bilinear group and a public common reference string for the system. The latter contains parameters for some of the building blocks, namely the NIZK and commitment scheme(s) we use.

<sup>11</sup>Note that by abuse of notation, we sometimes ignore the opening or decommitment value which is also an output of  $\text{Com}(\cdot)$ .

**4.4.2 DR/TSP/RSU Key Generation.** The DR generates a key pair  $(pk_{DR}, sk_{DR})$  for an IND-CCA secure encryption scheme, where  $pk_{DR}$  is used to deposit a user-specific trapdoor (a PRF-key) which allows to link this user's transactions in case of a dispute. An individual signature key pair  $(pk_{\mathcal{R}}, sk_{\mathcal{R}})$  is generated for each RSU to sign the updatable part of a user's wallet. Moreover, the TSP generates several key pairs  $(pk_{\mathcal{T}}, sk_{\mathcal{T}}^{\mathcal{T}})$ ,  $(pk_{\mathcal{T}}^{\text{cert}}, sk_{\mathcal{T}}^{\text{cert}})$ ,  $(pk_{\mathcal{T}}^{\mathcal{R}}, sk_{\mathcal{T}}^{\mathcal{R}})$  for an EUF-CMA secure signature scheme. The key  $sk_{\mathcal{T}}^{\mathcal{T}}$  is used to sign fixed user-specific information when a new wallet is issued. Using  $sk_{\mathcal{T}}^{\mathcal{R}}$ , the TSP can play the role of the initial RSU to sign the updatable part of a new wallet.

**4.4.3 RSU Certification.** An RSU engages with the TSP in this protocol to get its certificate  $\text{cert}_{\mathcal{R}}$ . It contains the RSU's public key  $pk_{\mathcal{R}}$ , its attributes  $\mathbf{a}_{\mathcal{R}}$  (that are assigned in this protocol by the TSP), and a signature on both, generated by the TSP using  $sk_{\mathcal{T}}^{\text{cert}}$ .

**4.4.4 User Registration.** Before a user can participate in the system, he needs to generate a key pair  $(pk_{\mathcal{U}} = g_1^{\text{sk}_{\mathcal{U}}}, sk_{\mathcal{U}}) \in G_1 \times \mathbb{Z}_p$ . We assume that binding  $pk_{\mathcal{U}}$  to a verified physical user ID, in order to hold the user liable in case of a misuse, is done out-of-band. For instance, this could be done by some trusted (external) certification authority. Throughout the protocols the key pair  $(pk_{\mathcal{U}}, sk_{\mathcal{U}})$  is used to bind a wallet to a user.

**4.4.5 Wallet Issuing.** This protocol is executed between a user and the TSP to create a new wallet with a fresh wallet ID  $\lambda$  and balance 0. Additionally, the PRF key  $\lambda$  is deposited under the DR's public encryption key. To generate the wallet, the user encodes  $\lambda$  together with his other secrets into the wallet. The PRF key  $\lambda$  needs to be chosen jointly and remain secret to the TSP. If only the user chose the key, an adversary could tamper with recalculations and blacklisting, as well as with double-spending detection (e.g., by choosing the same key for two different users). If only the TSP chose it, a user's transactions would be linkable.

To this end, the user and the TSP engage in the first two messages of a Blum coin toss. After the second message  $\lambda$  is fixed and the user knows his own share  $\lambda'$  as well as the TSP's share  $\lambda''$ . Then the user computes the commitments  $c_{\mathcal{T}} = \text{Com}(\lambda, sk_{\mathcal{U}})$  and  $c_{\mathcal{R}} := \text{Com}(\lambda, b := 0, u_1, x := 0)$ . Additionally, he prepares the deposit of  $\lambda$ . This is a bit tricky, as the user needs to prove that the ciphertext he gives to the TSP is actually an encryption of  $\lambda$  under  $pk_{DR}$ . For practical reasons, we use Groth-Sahai NIZKs and the Dodis-Yampolski PRF. Ideally, one would want an encryption scheme that is compatible with Groth-Sahai and whose message space equals the key space of the PRF, i.e.,  $\mathbb{Z}_p$ . Unfortunately, we are not aware of any such scheme.<sup>12</sup> Instead, we use a CCA-secure structure-preserving encryption scheme for vectors of  $G_1$ -elements and the following workaround: The user splits up its share  $\lambda'$  into small chunks  $\lambda'_i < \mathcal{B}$  (e.g.,  $\mathcal{B} = 2^{32}$ ) such that recovering the discrete logarithm of  $\Lambda'_i := g_1^{\lambda'_i}$  becomes feasible. All chunks  $\Lambda_{i \in \{0, \dots, \ell-1\}}$ , the TSP's share  $\Lambda'' := g_1^{\lambda''}$ , and the user's public key  $pk_{\mathcal{U}}$  are jointly encrypted as  $e^*$  under  $pk_{DR}$ . The CCA-secure ciphertext  $e^*$  unambiguously binds the PRF key  $\lambda$  to the user's key  $pk_{\mathcal{U}}$  and rules out malleability attacks. Otherwise, a malicious TSP could potentially trick the DR into recovering the trapdoor for a different (innocent) user. It is not necessary, to split the TSP share  $\lambda''$ , because it is known to the TSP anyway and thus can additionally be stored in the clear alongside  $e^*$ . The user then sends over  $e^*$ ,  $c_{\mathcal{T}}$ ,  $c_{\mathcal{R}}$  along with a NIZK  $\pi$  proving that everything has been generated honestly and the wallet is bound to the user owning  $sk_{\mathcal{U}}$ . When the TSP receives this data, it verifies the NIZK first. If the check passed, the TSP signs  $c_{\mathcal{T}}$  along with attributes  $\mathbf{a}_{\mathcal{U}}$  and  $c_{\mathcal{R}}$  along with  $s$  using  $sk_{\mathcal{T}}^{\mathcal{R}}$ .<sup>13</sup> The resulting signatures  $\sigma_{\mathcal{T}}$ ,  $\sigma_{\mathcal{R}}$  are sent to the user, who checks their correctness. The user finally stores his freshly generated state token

$$\tau := (\mathbf{a}_{\mathcal{U}}, c_{\mathcal{R}}, d_{\mathcal{R}}, \sigma_{\mathcal{R}}, c_{\mathcal{T}}, d_{\mathcal{T}}, \sigma_{\mathcal{T}}, \lambda := \lambda' + \lambda'', b := 0, u_1, x := 1, s),$$

<sup>12</sup>Note that Paillier encryption works in a different algebraic setting and cannot easily be combined with Groth-Sahai proofs.

<sup>13</sup>During the protocol run, a uniformly random serial number  $s$  for this transaction was jointly generated by user and TSP by means of a Blum-like coin toss (see Debt Accumulation for a description).

where  $d_{\mathcal{R}}$  and  $d_{\mathcal{T}}$  are the decommitment values required to open  $c_{\mathcal{R}}$  and  $c_{\mathcal{T}}$ , respectively. The TSP stores  $htd := (\text{pk}_{\mathcal{U}}, s, \lambda'', e^*)$  as hidden user trapdoor to recover  $\lambda$  with help of DR.

**4.4.6 User Blacklisting.** This is a protocol executed by the DR upon request of the TSP. We assume that DR and TSP agreed out-of-band that the user with public key  $\text{pk}_{\mathcal{U}}^{\text{DR}}$  should be blacklisted before the protocol starts. For example, the TSP might have presented a proof of guilt to the DR or the user affirmed to be voluntarily blacklisted (for example due to a stolen car). Given  $e^*$  and  $\lambda''$ , the DR recovers the contained PRF key but only if it is bound to  $\text{pk}_{\mathcal{U}}^{\text{DR}}$ . To this end, DR decrypts  $e^*$  using  $\text{sk}_{\text{DR}}$  to obtain  $\Lambda'_i, \Lambda''$ , and  $\text{pk}_{\mathcal{U}}$ . If  $\text{pk}_{\mathcal{U}} \neq \text{pk}_{\mathcal{U}}^{\text{DR}}$  or  $\Lambda'' \neq g_1^{\lambda''}$  it aborts. Otherwise, it computes the (small) discrete logarithms of the  $\Lambda'_i$  to the base  $g_1$  to recover the chunks  $\lambda'_i$  of the user's share of the PRF key. This is feasible as the DLOGs are very small. The key is computed as  $\lambda := \lambda'' + \sum_{i=0}^{\ell-1} \lambda'_i \cdot \mathcal{B}^i$ . The DR also checks if the provided share of the TSP  $\lambda''$  matches the decrypted version  $\Lambda''$  and if the decrypted user key  $\text{pk}_{\mathcal{U}}^{\mathcal{T}}$  equals the expected key  $\text{pk}_{\mathcal{U}}^{\text{DR}}$ . If the check fails, the DR aborts.

Using  $\lambda$ , the fraud detection IDs belonging to the previous and upcoming versions of a user's wallet can be computed. Thus, all interactions of the user (including double-spending) in the TSP's database can be linked and the legitimate debt recalculated. Also, the fraud detection IDs for upcoming transactions with this wallet can be precomputed and blacklisted.

**4.4.7 Debt Accumulation.** In this protocol (Fig. 3) executed between an anonymous user and an RSU, the toll  $p$  is determined and a new state of the user's wallet with a debt increased by  $p$  is created. The toll  $p$  may depend on the user's attributes, the current and the previous reader's attributes, as well as other factors like the current time, etc. The user stays anonymous, i.e., the RSU does not receive the user's key or see the previous balance. Only  $\mathbf{a}_{\mathcal{U}}$  and  $\mathbf{a}_{\mathcal{R}}^{\text{prev}}$  (the attributes of the RSU who signed the previous state of the wallet) are revealed to the RSU, which are assumed not to contain identifying information.

The user's main input is the state token

$$\tau^{\text{prev}} := (\mathbf{a}_{\mathcal{U}}, c_{\mathcal{R}}^{\text{prev}}, d_{\mathcal{R}}^{\text{prev}}, \sigma_{\mathcal{R}}^{\text{prev}}, c_{\mathcal{T}}, d_{\mathcal{T}}, \sigma_{\mathcal{T}}, \lambda, b^{\text{prev}}, u_1^{\text{prev}}, x, s^{\text{prev}})$$

containing the state of his previous protocol run and wallet. To update his wallet state, the user computes a fresh commitment  $c'_{\mathcal{R}}$  on  $(\lambda, b^{\text{prev}}, u_1^{\text{next}}, x)$ , where except for  $u_1^{\text{next}}$  the same values are contained in the previous wallet state. The randomness  $u_1^{\text{next}}$  is freshly chosen by the user and is used to generate a double-spending tag for the new wallet state. In order to get his new wallet state certified by the RSU, the user needs to invalidate his previous state. To do so, he needs to reveal the fraud detection ID and double-spending tag of his previous state. For the latter, the RSU sends a random challenge  $u_2$  along with a commitment  $c''_{\text{ser}} = \text{Com}(s'')$  on his share of the serial number of this transaction (which is part of the Blum coin toss). Upon receiving these values, the user computes the double-spending tag  $t := u_2 \cdot \text{sk}_{\mathcal{U}} + u_1^{\text{next}} \pmod{\mathfrak{p}}$  (a linear equation in the unknowns  $u_1$  and  $\text{sk}_{\mathcal{U}}$ ), the fraud detection ID  $\phi := \text{PRF}(\lambda, x)$ , and a hidden user ID  $c_{\text{hid}} := \text{Com}(\text{pk}_{\mathcal{U}})$ . The latter is used in Prove Participation to associate this interaction with the user.<sup>14</sup> As response, the user sends over  $c_{\text{hid}}, c'_{\mathcal{R}}, \phi, t, \mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}}, \pi$ , and  $s'$ , where  $\pi$  is a NIZK proving that everything has been computed honestly, and  $s'$  is the user's share of the serial number. In particular,  $\pi$  shows that the user knows a certified wallet state involving commitments  $c_{\mathcal{T}}$  and  $c_{\mathcal{R}}^{\text{prev}}$  such that  $c_{\mathcal{R}}^{\text{prev}}$  and  $c'_{\mathcal{R}}$  are commitments on the same messages except for the double-spending randomness, that the (hidden) signature on  $c_{\mathcal{R}}^{\text{prev}}$  verifies under some (hidden) RSU key  $\text{pk}_{\mathcal{R}}^{\text{prev}}$  certified by the TSP, and that  $t, \phi$ , and  $c_{\text{hid}}$  have been computed using the values contained in  $c_{\mathcal{T}}$  and  $c_{\mathcal{R}}^{\text{prev}}$ .

When receiving this data, the RSU first checks that  $\phi$  is not on the blacklist and  $\pi$  is correct. Then it calculates the price  $p$ , adds it to the user's balance  $b^{\text{prev}}$  and increases the counter  $x$  by 1 using the homomorphic property of  $c'_{\mathcal{R}}$ . The resulting commitment  $c_{\mathcal{R}}$  is signed along with the serial number  $s := s' \cdot s''$  using  $\text{sk}_{\mathcal{R}}$ . Then the

<sup>14</sup>We are aware of alternatives realizing this mechanism without introducing  $c_{\text{hid}}$ , which are, however, less efficient.




Fig. 3. A Simplified Version of Debt Accumulation

RSU sends  $c_{\mathcal{R}}$  to the user, along with its decommitment information  $d_{\mathcal{R}}$ , its signature  $\sigma_{\mathcal{R}}$ , the added price  $p$ , the certificate for  $\text{pk}_{\mathcal{R}}$ , RSU's share  $s''$  of the serial number and the decommitment value  $d''_{\text{ser}}$  for  $c''_{\text{ser}}$ .

The user checks if the received data is valid and ends up with an updated state token

$$\tau := (\text{a}_{\mathcal{U}}, c_{\mathcal{R}}, d_{\mathcal{R}}, \sigma_{\mathcal{R}}, c_{\mathcal{T}}, d_{\mathcal{T}}, \sigma_{\mathcal{T}}, \lambda, b^{\text{prev}} + p, u_1^{\text{next}}, x + 1, s)$$

containing his new wallet state  $c_{\mathcal{T}}, c_{\mathcal{R}}, \sigma_{\mathcal{T}}, \sigma_{\mathcal{R}}$ . He additionally stores  $(c_{\text{hid}}, d_{\text{hid}})$ , where  $d_{\text{hid}}$  allows to open the hidden ID. The RSU stores  $(\phi, p, t, u_2, c_{\text{hid}}, s)$  as transaction information.

**4.4.8 Prove Participation.** In this protocol executed between a user and the SA, the user proves that he participated in one of the Debt Accumulation transactions under audit by the SA. This protocol is identifying, as the SA retrieved the user's physical ID from his license plate number and, thus, is aware of his public key  $\text{pk}_{\mathcal{U}}$ . First, it

sends the list  $\Omega_{\mathcal{R}}^{\text{pp}}$  of hidden ID commitments observed by the RSU during the considered interactions. If the user owns some  $c_{\text{hid}}$  contained in  $\Omega_{\mathcal{R}}^{\text{pp}}$ , he simply sends over  $c_{\text{hid}}$  along with the corresponding opening  $d_{\text{hid}}$  he stored and the serial number  $s$  of the transaction.<sup>15</sup> The SA accepts if the commitment  $d_{\text{hid}}$  indeed opens to  $\text{pk}_{\mathcal{U}}$  and  $s$  is part of the corresponding transaction information the RSU stored. No other user may claim to have sent  $c_{\text{hid}}$ , as this would imply to open  $c_{\text{hid}}$  with a different public key.

**4.4.9 Debt Clearance.** In this protocol executed between a user and the TSP, the user identifies himself using  $\text{sk}_{\mathcal{U}}$  and reveals the balance for his current wallet state. The wallet is invalidated by not creating a new state. Upon receiving a challenge  $u_2$  from the TSP, the user computes the double-spending tag  $t$  and fraud detection ID  $\phi$  for his wallet. This is the same as in Debt Accumulation. Then the user sends over  $\text{pk}_{\mathcal{U}}$ ,  $b^{\text{prev}}$ ,  $\phi$ ,  $t$  and a NIZK  $\pi$ . This NIZK essentially shows that the user knows a certified wallet state with balance  $b^{\text{prev}}$ , fraud detection ID  $\phi$ , and double-spending tag  $t$  which is bound to  $\text{pk}_{\mathcal{U}}$ . Note that if the user does not make use of his latest wallet state, double-spending detection will reveal this. If the proof verifies, the balance and the double-spending information, i.e.,  $(b^{\text{prev}}, \phi, t, u_2)$ , is stored by the TSP.

**4.4.10 Double-Spending Detection.** This algorithm is applied by the TSP to its database containing all double-spending tags collected by the RSUs. The fraud detection ID  $\phi$  is used as the database index associated with the double-spending tag for a wallet state. If the same index appears twice in the TSP's database, a double-spending occurred and the cheater's key pair can be reconstructed from the two double-spending tags as follows: Let us assume there exist two records  $(\phi, t, u_2), (\phi, t', u'_2)$ . In this case,  $\text{sk}_{\mathcal{U}}$  can be recovered as  $\text{sk}_{\mathcal{U}} = (t - t')(u_2 - u'_2)^{-1} \bmod p$  with overwhelming probability. The cheater's public key  $\text{pk}_{\mathcal{U}}$  can be computed from  $\text{sk}_{\mathcal{U}}$ . As a consequence, the wallet bound to  $\text{pk}_{\mathcal{U}}$  could be blacklisted. The output of the protocol is the fraudulent user's public key  $\text{pk}_{\mathcal{U}}$  along with a proof-of-guilt  $\pi := \text{sk}_{\mathcal{U}}$ .

**4.4.11 Guilt Verification.** This algorithm can be executed by any party to verify the guilt of an accused double-spender. Given a public key  $\text{pk}_{\mathcal{U}}$  and a proof-of-guilt  $\pi$ , it checks if  $g_1^\pi = \text{pk}_{\mathcal{U}}$ .

## 5 SECURITY THEOREM

Assuming co-CDH is hard and our building blocks are secure, we prove that the protocols from [Section 4](#) (combined as one comprehensive toll collection protocol  $\pi_{\text{P4TC}}$ ) GUC-realize the ideal model  $\mathcal{F}_{\text{P4TC}}$  from [Section 3](#) (in the  $(\mathcal{F}_{\text{CRS}}, \overline{\mathcal{G}}_{\text{bb}})$ -hybrid model), i.e.,

$$\pi_{\text{P4TC}}^{\mathcal{F}_{\text{CRS}}, \overline{\mathcal{G}}_{\text{bb}}} \geq_{\text{UC}} \mathcal{F}_{\overline{\mathcal{G}}_{\text{bb}}}.$$

Informally, this means the ideal model and our protocol are indistinguishable and therefore provide the same guarantees regarding security and privacy. The statement holds given static corruption of either

- (1) A subset of users.
- (2) All users and a subset of RSUs, TSP and SA.
- (3) A subset of RSUs, TSP and SA.
- (4) All RSUs, TSP and SA as well as a subset of users.

Note that we assume the DR to be honest. In [Appendix E.1](#), more details on corruption possibilities are presented. Other details of our adversarial model, i.e. the underlying channel model and handling of aborts, can be found in [Appendix E](#) as well. [Section 5](#) gives an outline of the proof, the full proof is given in [Appendix F](#).

<sup>15</sup>By additionally providing a time interval out-of-band this search can be accelerated.

## Proof Outline

As mentioned above, we separately prove correctness, system security and user security and privacy. Technically, the first step is not necessary, as the case in which all parties are honest is a special case of a scenario in which a subset of parties is corrupted. Therefore both [Theorems F.22](#) and [F.36](#) individually imply correctness. Showing correctness separately, however, serves two purposes: For didactic reasons a restricted setting (i.e., all parties are assumed to be honest) is convenient to introduce notation and to present the underlying idea that is common to all our proofs. Secondly, having already shown correctness as a special case helps during the course of the proof of system security to keep matters simple, because we can revert to the special case.

To show correctness of our protocol, we start by asserting several structural properties of the ideal transaction database  $TRDB$ . First, note that the entries  $trdb = (s^{\text{prev}}, s, \phi, \lambda, pid_{\mathcal{U}}, pid_{\mathcal{R}}, p, b)$  of the ideal transaction database define a graph structure where serial numbers  $s$  are considered vertices and predecessors  $s^{\text{prev}}$  define edges  $(s^{\text{prev}}, s)$ . Assigning a label  $(pid_{\mathcal{R}}, p)$  to this edge and  $(\phi, \lambda, pid_{\mathcal{U}}, b)$  to the vertex  $s$  results in a graph where every vertex represents the state of a wallet and the incoming edge represents the transaction that led to this wallet state. We call this perception of  $TRDB$  the *Ideal Transaction Graph* and give graph-theoretic proofs of its structural properties. These properties include that the graph as a whole is a directed forest where each tree corresponds to a wallet ID  $\lambda$ , double-spending corresponds to branching and different wallet states have the same fraud detection ID  $\phi$  if and only if they have the same depth in the same tree.

Secondly we add in- and out-commitments  $(c_{\mathcal{R}}^{\text{in}}, c_{\mathcal{T}}^{\text{in}})$  and  $(c_{\mathcal{R}}^{\text{out}}, c_{\mathcal{T}}^{\text{out}})$  from the real protocols to each transaction node in the Ideal Transaction Graph. These commitments are the fixed and updateable part of the wallet before and after the transaction (cp. [Section 4](#)). This information gives a second set of edges where two transactions  $trdb$  and  $trdb^*$  are connected if  $(c_{\mathcal{R}}^{\text{out}}, c_{\mathcal{T}}^{\text{out}})$  corresponds to  $(c_{\mathcal{R}}^{\text{in}*}, c_{\mathcal{T}}^{\text{in}*})$ . We call the resulting graph the *Augmented Transaction Graph*. Showing that in case of an honest execution of the toll collection protocol both of those graph structures coincide with overwhelming probability yields correctness.

The proofs of system security and user security and privacy are conducted by explicitly specifying a simulator and reducing the indistinguishability of real and ideal world to the security of our building blocks. During an execution of the protocol our simulator generates the Augmented Transaction Graph as explained above. After showing that all messages sent by the simulator are statistically close to the real messages, i.e., simulated perfectly, that leaves only two kinds of reasons the environment  $\mathcal{Z}$  could be able to distinguish both worlds. The first kind are *failure events* where the two graph structures in the augmented transaction graph diverge and the second are *discrepancy events* where some party's outputs could be distinguished. We show that both of those cases only occur with negligible probability by various reductions to our cryptographic building blocks and hardness assumptions. The details of the proofs can be found in [Appendix F](#).

## 6 PERFORMANCE EVALUATION

In order to evaluate the practicality of P4TC, we implemented our system for a realistic target platform. We performed our measurements for the user side on an i.MX6 Dual-Core processor running at 800MHz with 1GB DDR3 RAM and 4GB eMMC Flash, the same processor as used in the Savari MobiWAVE-1000 OBU [\[64\]](#). The processor runs an embedded Linux, is ARM Cortex-A9 based (32-bit), and also exists in a more powerful Quad-Core variant. For the RSU hardware we take the ECONOLITE Connected Vehicle CoProcessor Module as a reference system, which was specifically “designed to enable third-party-developed or processor-intensive applications” [\[41\]](#) and measured on comparable hardware. For simplicity, we measured TSP runtime on a standard laptop featuring an i7-6600U processor, although we can expect the TSP to be equipped with much more powerful hardware.

As the CPU of the OBU and the RSU are comparable, we measured both sides using the OBU processor.

Table 1. Averaged performance results for execution time  $t$  and transmitted data  $n$ . Byte values are rounded.

Protocol	$ \mathbf{a}_U  =  \mathbf{a}_R  = 1$				$ \mathbf{a}_U  =  \mathbf{a}_R  = 4$			
	$t_{\text{user}}$ [ms]	$t_{\text{RSU/TSP}}$ [ms]	$n_{\text{user}}$ [byte]	$n_{\text{RSU/TSP}}$ [byte]	$t_{\text{user}}$ [ms]	$t_{\text{RSU/TSP}}$ [ms]	$n_{\text{user}}$ [byte]	$n_{\text{RSU/TSP}}$ [byte]
Session Key Generation	15.01	0.03	131	–	15.00	0.01	131	–
Wallet Issuing	27 063.61	(8 490.00)	87 951	944	27 182.82	(8 544.74)	88 107	1 152
Debt Accumulation								
– Precomp (offline)	2 715.03	–	–	–	2 715.66	–	–	–
– Online	348.14	474.64	8 128	976	455.52	525.83	8 336	1 088
– Online (cached certificate)	41.01	474.64	8 128	976	40.23	525.83	8 336	1 088
– Postprocessing (offline)	33.84	–	–	–	33.78	–	–	–
Debt Clearance	2 435.47	(745.30)	7 071	96	2 435.78	(766.90)	7 280	96

### 6.1 Building Block Instantiation

We implemented P4TC in C++14 using the RELIC toolkit v.0.4.1, an open source cryptography and arithmetic library written in C, with support for pairing-friendly elliptic curves [5]. We developed our own library for Groth-Sahai NIZK proofs [37, 40] and employed the method in [15] to realize the range proofs required in Wallet Issuing. We make use of two types of commitments: the shrinking commitment scheme from [3], as well as a (dual-mode) extractable commitment scheme from [40]. Moreover, we implemented the structure-preserving signature scheme from [1]. We use an adopted variant of the IND-CCA-secure encryption scheme from [16] (cp. Appendix C.2.5) for encrypting PRF key shares in the Wallet Issuing protocol, as well as the IND-CCA-secure encryption scheme from [26] (in combination with AES-CBC and HMAC-SHA256) to establish secure channels for exchanging protocol messages. The PRF is instantiated with the Dodis-Yampolskiy construction introduced in [36].

### 6.2 Parameter Choice

As for the bilinear group setting, we use the Barreto-Naehrig curves Fp254BNb and Fp254n2BNb presented by Aranha et al. [11, 51]. For the pairing function, we use the optimal Ate pairing since it results in the shortest execution times [60]. This yields a security level of about 100 bit [9].

We evaluated P4TC for two sizes of attribute vectors:  $|\mathbf{a}_U| = |\mathbf{a}_R| = 1$  and  $|\mathbf{a}_U| = |\mathbf{a}_R| = 4$ . With curves of 254-bit order, each vector component can encode up to 253 bits of arbitrary information. In practice, it should be possible to encode multiple attributes into one such component. Since the actual encoding of data depends on the concrete scenario, we only focus here on evaluating the performance penalties when increasing the size of the vectors.

### 6.3 Implementation Results

Table 1 shows the results of our measurements on the OBU and on the RSU/TSP side in terms of execution time and transmitted data. All values are averaged over 1000 independent protocol executions. Note that the processor is running an embedded Linux, hence execution times can vary by tens of milliseconds due to internal processes and scheduling. We only considered the main protocols which include the expensive NIZKs. The row entitled “Session Key Generation” in Table 1 includes the execution time and size of data to setup a session key for the



secure channel which is established prior to any protocol run.<sup>16</sup> In order to utilize the capabilities of our OBU hardware, the user side algorithms were optimized for two CPU cores. Note that the given TSP times (the gray entries in Table 1) are likely an upper bound, since the TSP is equipped with more powerful hardware than a standard laptop.

*Debt Accumulation.* While the protocols Wallet Issuing and Debt Clearance can be regarded as non-time-critical, Debt Accumulation is performed while driving (possibly at high speed). Thus, execution has to be as efficient as possible. Fortunately, all parts of the expensive NIZK proof which do not involve the challenge value  $u_2$  (provided by the RSU) can be precomputed (*offline* phase) at the OBU which takes approximately 2700ms. Regarding  $|a_U| = |a_R| = 1$ , computations in this online phase take only approximately 350ms on the OBU, mostly due to the verification of the RSU certificate. When caching valid certificates, the runtime can be reduced to approximately 40ms. After the OBU received a response from the RSU, which computes for about 475ms, the internal wallet has to be updated. Since this step can be done offline again, we measure the execution time separately and call the phase *postprocessing*. We also optimized the computations performed by the RSU, taking advantage of the 4 CPU cores and the batching techniques for Groth-Sahai verification by Herold et al. [44]. In summary, all computations in the online phase of Debt Accumulation can be performed in about 825ms or just 515ms when the certificate has already been cached.

The WAVE data transmission standard on DSRC guarantees a transmission rate of 24 Mbit/s [57]. At this rate, all data of the Debt Accumulation protocol are transmitted in approximately 27ms. While the standard claims communication ranges of up to 1km, we assume a toll collection zone of 50m. Moreover, we may assume one RSU per lane [54] such that the workload can be easily spread among the available units. Going at 120km/h, it takes a car 1.5s to travel this distance, leaving us with a time buffer of about 700ms for uncached certificates or 1s for cached certificates. Considering the mandated safety distance at this speed, there should only be a single car inside the 50m zone on each lane. However, since computations take less time than it takes a car to cross the toll collection zone, the system could theoretically handle cars with a distance of only 28m (uncached certificate) or 17m (cached certificate) at 120km/h. We therefore conclude that the performance is sufficient for real-world scenarios.

*Storage Requirements.* During Debt Accumulation, the RSU and the OBU collect data in order to, e.g., prevent double-spending or to prove participation in a protocol run. In Debt Accumulation, the OBU has to store 137 bytes of transaction information and (optionally) 268 bytes to cache the RSU certificate. Assuming that in one billing period 10000 transactions are performed by the OBU, it only has to store 1.37MB of transaction information and, even if all visited RSUs were different, 2.68MB of cached certificates. The wallet itself consumes 1kB of memory and is fixed in size. The RSU has to store 246 bytes of transaction information after each run of Debt Accumulation (for 32-bit toll values). All this information is eventually aggregated at the TSP's database. The US-based toll collection system E-ZPass reported about 252.4 million transactions per month in 2016 [42], which would result in a database of size 62GB. In case a wallet is blacklisted, an RSU is updated with a list of future fraud detection IDs. Each detection ID consumes 35 bytes of memory. Using appropriate data structures such as hashsets or hashmaps, a lookup is performed in less than 1ms in sets of size  $10^6$ . Hence, blacklisting does not affect the execution time on the RSU. Such a set consumes 35MB of memory, neglecting the overhead induced by the data structure.

*Computing DLOGs.* To blacklist a user, the DR has to compute a number of discrete logarithms to recover  $\lambda$ . With our choice of parameters,  $\lambda$  is split into 32-bit values, thus resulting in the computation of eight 32-bit DLOGs. While DLOGs of this size can be brute-forced naively, the technique of Bernstein et al. [14] can be used

<sup>16</sup>All communication between the participants is encrypted with AES-CBC and HMAC-SHA256 to realize a secure channel. Therefore, each individual protocol is preceded by a symmetric key exchange via the KEM by Cash et al. [26].

to speed up this process. Using their algorithm, computing a discrete logarithm in an interval of order  $2^{32}$  takes around 1.5 seconds on a single core of a standard desktop using a 55kB table of precomputed elements. These precomputations need to be done only *once* by the DR when setting up the system and take one hour on a desktop computer. Thus, the required DLOGs can be computed in reasonable time by the DR.

## REFERENCES

- [1] Masayuki Abe, Jens Groth, Kristiyan Haralambiev, and Miyako Ohkubo. 2011. Optimal Structure-Preserving Signatures in Asymmetric Bilinear Groups. In *Advances in Cryptology – CRYPTO 2011 (Lecture Notes in Computer Science)*, Phillip Rogaway (Ed.), Vol. 6841. Springer, Heidelberg, 649–666.
- [2] Masayuki Abe, Jens Groth, Miyako Ohkubo, and Takeya Tango. 2014. Converting Cryptographic Schemes from Symmetric to Asymmetric Bilinear Groups. In *Advances in Cryptology – CRYPTO 2014, Part I (Lecture Notes in Computer Science)*, Juan A. Garay and Rosario Gennaro (Eds.), Vol. 8616. Springer, Heidelberg, 241–260. [https://doi.org/10.1007/978-3-662-44371-2\\_14](https://doi.org/10.1007/978-3-662-44371-2_14)
- [3] Masayuki Abe, Markulf Kohlweiss, Miyako Ohkubo, and Mehdi Tibouchi. 2015. Fully Structure-Preserving Signatures and Shrinking Commitments. In *Advances in Cryptology – EUROCRYPT 2015, Part II (Lecture Notes in Computer Science)*, Elisabeth Oswald and Marc Fischlin (Eds.), Vol. 9057. Springer, Heidelberg, 35–65. [https://doi.org/10.1007/978-3-662-46803-6\\_2](https://doi.org/10.1007/978-3-662-46803-6_2)
- [4] Carolin Albrecht, Frank Gurski, Jochen Reithmann, and Eda Yilmaz. 2018. Knapsack Problems: A Parameterized Point of View. *Theoretical Computer Science* (2018).
- [5] D. F. Aranha and C. P. L. Gouvêa. 2016. RELIC is an Efficient Library for Cryptography. Online Resource. (2016). <https://github.com/relic-toolkit/relic>.
- [6] Josep Balasch, Alfredo Rial, Carmela Troncoso, Bart Preneel, Ingrid Verbauwhede, and Christophe Geuens. 2010. PrETP: Privacy-Preserving Electronic Toll Pricing (extended version). In *Proceedings of the 19th USENIX Security Symposium*. 63–78.
- [7] Foteini Baldimtsi, Melissa Chase, Georg Fuchsbauer, and Markulf Kohlweiss. 2015. Anonymous Transferable E-Cash. In *Public-Key Cryptography - PKC 2015 - 18th IACR International Conference on Practice and Theory in Public-Key Cryptography, Gaithersburg, MD, USA, March 30 - April 1, 2015, Proceedings (Lecture Notes in Computer Science)*, Jonathan Katz (Ed.), Vol. 9020. Springer, 101–124. [https://doi.org/10.1007/978-3-662-46447-2\\_5](https://doi.org/10.1007/978-3-662-46447-2_5)
- [8] Boaz Barak, Ran Canetti, Jesper Buus Nielsen, and Rafael Pass. 2004. Universally Composable Protocols with Relaxed Set-Up Assumptions. In *45th Annual Symposium on Foundations of Computer Science*. IEEE Computer Society Press, 186–195.
- [9] Razvan Barbulescu and Sylvain Duquesne. 2017. Updating key size estimations for pairings. Cryptology ePrint Archive, Report 2017/334. (2017). <http://eprint.iacr.org/2017/334>.
- [10] Amira Barki, Solenn Brunet, Nicolas Desmoulins, Sébastien Gambs, Saïd Gharout, and Jacques Traoré. 2016. Private eCash in Practice (Short Paper). In *International Conference on Financial Cryptography and Data Security (FC 2016)*. Springer, 99–109.
- [11] Paulo S. L. M. Barreto and Michael Naehrig. 2006. Pairing-Friendly Elliptic Curves of Prime Order. In *SAC 2005: 12th Annual International Workshop on Selected Areas in Cryptography (Lecture Notes in Computer Science)*, Bart Preneel and Stafford Tavares (Eds.), Vol. 3897. Springer, Heidelberg, 319–331.
- [12] Mihir Bellare. 2015. New Proofs for NMAC and HMAC: Security Without Collision-Resistance. *Journal of Cryptology* 28, 4 (2015), 844–878.
- [13] Mihir Bellare, Anand Desai, David Pointcheval, and Phillip Rogaway. 1998. Relations Among Notions of Security for Public-Key Encryption Schemes. In *Advances in Cryptology – CRYPTO’98 (Lecture Notes in Computer Science)*, Hugo Krawczyk (Ed.), Vol. 1462. Springer, Heidelberg, 26–45.
- [14] Daniel J. Bernstein and Tanja Lange. 2012. Computing small discrete logarithms faster. Cryptology ePrint Archive, Report 2012/458. (2012). <http://eprint.iacr.org/2012/458>.
- [15] Jan Camenisch, Rafik Chaabouni, and abhi shelat. 2008. Efficient Protocols for Set Membership and Range Proofs. In *Advances in Cryptology – ASIACRYPT 2008 (Lecture Notes in Computer Science)*, Josef Pieprzyk (Ed.), Vol. 5350. Springer, Heidelberg, 234–252.
- [16] Jan Camenisch, Kristiyan Haralambiev, Markulf Kohlweiss, Jorn Lapon, and Vincent Naessens. 2011. Structure Preserving CCA Secure Encryption and Applications. In *Advances in Cryptology – ASIACRYPT 2011 (Lecture Notes in Computer Science)*, Dong Hoon Lee and Xiaoyun Wang (Eds.), Vol. 7073. Springer, Heidelberg, 89–106.
- [17] Jan Camenisch, Susan Hohenberger, and Anna Lysyanskaya. 2005. Compact E-Cash. In *Advances in Cryptology - EUROCRYPT 2005, 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Aarhus, Denmark, May 22-26, 2005, Proceedings (Lecture Notes in Computer Science)*, Ronald Cramer (Ed.), Vol. 3494. Springer, 302–321. [https://doi.org/10.1007/11426639\\_18](https://doi.org/10.1007/11426639_18)
- [18] Jan Camenisch, Susan Hohenberger, and Anna Lysyanskaya. 2005. Compact E-Cash. In *Advances in Cryptology – EUROCRYPT 2005 (Lecture Notes in Computer Science)*, Ronald Cramer (Ed.), Vol. 3494. Springer, Heidelberg, 302–321.
- [19] Sébastien Canard and Aline Gouget. 2008. Anonymity in Transferable E-cash. In *Applied Cryptography and Network Security, 6th International Conference, ACNS 2008, New York, NY, USA, June 3-6, 2008. Proceedings (Lecture Notes in Computer Science)*, Steven M. Bellovin,

- Rosario Gennaro, Angelos D. Keromytis, and Moti Yung (Eds.), Vol. 5037. 207–223. [https://doi.org/10.1007/978-3-540-68914-0\\_13](https://doi.org/10.1007/978-3-540-68914-0_13)
- [20] Ran Canetti. 2001. Universally Composable Security: A New Paradigm for Cryptographic Protocols. In *42nd Annual Symposium on Foundations of Computer Science*. IEEE Computer Society Press, 136–145.
- [21] Ran Canetti. 2006. Security and Composition of Cryptographic Protocols: A Tutorial. Cryptology ePrint Archive, Report 2006/465. (2006). <http://eprint.iacr.org/2006/465>.
- [22] Ran Canetti. 2007. Obtaining Universally Composable Security: Towards the Bare Bones of Trust. Cryptology ePrint Archive, Report 2007/475. (2007). <http://eprint.iacr.org/2007/475>.
- [23] Ran Canetti, Yevgeniy Dodis, Rafael Pass, and Shabsi Walfish. 2007. Universally Composable Security with Global Setup. In *TCC 2007: 4th Theory of Cryptography Conference (Lecture Notes in Computer Science)*, Salil P. Vadhan (Ed.), Vol. 4392. Springer, Heidelberg, 61–85.
- [24] Ran Canetti, Yehuda Lindell, Rafail Ostrovsky, and Amit Sahai. 2002. Universally composable two-party and multi-party secure computation. In *34th Annual ACM Symposium on Theory of Computing*. ACM Press, 494–503.
- [25] Ran Canetti, Daniel Shahaf, and Margarita Vald. 2016. Universally Composable Authentication and Key-Exchange with Global PKI. In *PKC 2016: 19th International Conference on Theory and Practice of Public Key Cryptography, Part II (Lecture Notes in Computer Science)*, Chen-Mou Cheng, Kai-Min Chung, Giuseppe Persiano, and Bo-Yin Yang (Eds.), Vol. 9615. Springer, Heidelberg, 265–296. [https://doi.org/10.1007/978-3-662-49387-8\\_11](https://doi.org/10.1007/978-3-662-49387-8_11)
- [26] David Cash, Eike Kiltz, and Victor Shoup. 2008. The Twin Diffie-Hellman Problem and Applications. In *Advances in Cryptology – EUROCRYPT 2008 (Lecture Notes in Computer Science)*, Nigel P. Smart (Ed.), Vol. 4965. Springer, Heidelberg, 127–145.
- [27] Xihui Chen, Gabriele Lenzini, Souke Mauw, and Jun Pang. 2012. A Group Signature Based Electronic Toll Pricing System. In *2012 Seventh International Conference on Availability, Reliability and Security (ARES 2012)*. 85–93.
- [28] Xihui Chen, Gabriele Lenzini, Sjouke Mauw, and Jun Pang. 2013. Design and Formal Analysis of A Group Signature Based Electronic Toll Pricing System. *JoWUA* 4, 1 (2013), 55–75. <http://isyou.info/jowua/papers/jowua-v4n1-3.pdf>
- [29] European Commission. 2015. Study on State of the Art of Electronic Road Tolling. [https://ec.europa.eu/transport/sites/transport/files/modes/road/road\\_charging/doc/study-electronic-road-tolling.pdf](https://ec.europa.eu/transport/sites/transport/files/modes/road/road_charging/doc/study-electronic-road-tolling.pdf). (2015). [Online; accessed April-19-2018].
- [30] European Commission. 2017. Proposal for a Directive of the European Parliament and of the Council on the Interoperability of Electronic Road Toll Systems and Facilitating Crossborder Exchange of Information on the Failure to Pay Road Fees in the Union (recast). <https://ec.europa.eu/transport/sites/transport/files/com20170280-eets-directive.pdf>. (2017). [Online; accessed April-19-2018].
- [31] Morten Dahl, Stéphanie Delaune, and Graham Steel. 2012. Formal Analysis of Privacy for Anonymous Location Based Services. *Theory of Security and Applications* (2012), 98–112.
- [32] Datatilsynet. 2007. Statens Vegvesen Holdt Tilbake Viktig AutoPASS-Infomasjon (Press release). <http://www.datatilsynet.no/>. (2007). [Online; accessed April-19-2018].
- [33] Jeremy Day, Yizhou Huang, Edward Knapp, and Ian Goldberg. 2011. SPecTRE: Spot-checked Private Ecash Tolling at Roadside. In *Proceedings of the 10th Annual ACM Workshop on Privacy in the Electronic Society (WPES '11)*. 61–68.
- [34] Wiebren de Jonge and Bart Jacobs. 2008. Privacy-Friendly Electronic Traffic Pricing via Commits. In *Formal Aspects in Security and Trust – FAST 2008 (Lecture Notes in Computer Science)*, Vol. 5491. 143–161.
- [35] Deutsches Bundesamt für Güterverkehr. 2018. Monatliche Mautstatistik für Januar 2018. [https://www.bag.bund.de/SharedDocs/Downloads/DE/Statistik/Lkw-Maut/18\\_Monatstab\\_01.html](https://www.bag.bund.de/SharedDocs/Downloads/DE/Statistik/Lkw-Maut/18_Monatstab_01.html) (in German). (2018). [Online; accessed January-09-2019].
- [36] Yevgeniy Dodis and Aleksandr Yampolskiy. 2004. A Verifiable Random Function With Short Proofs and Keys. Cryptology ePrint Archive, Report 2004/310. (2004). <http://eprint.iacr.org/2004/310>.
- [37] Alex Escala and Jens Groth. 2014. Fine-Tuning Groth-Sahai Proofs. In *PKC 2014: 17th International Conference on Theory and Practice of Public Key Cryptography (Lecture Notes in Computer Science)*, Hugo Krawczyk (Ed.), Vol. 8383. Springer, Heidelberg, 630–649. [https://doi.org/10.1007/978-3-642-54631-0\\_36](https://doi.org/10.1007/978-3-642-54631-0_36)
- [38] eugdpr.org. 2018. The EU General Data Protection Regulation (GDPR). <https://www.eugdpr.org/>. (2018). [Online; accessed April-19-2018].
- [39] Flavio D Garcia, Eric R Verheul, and Bart Jacobs. 2011. Cell-Based Roadpricing. In *European Public Key Infrastructure Workshop – EuroPKI 2011 (Lecture Notes in Computer Science)*, Vol. 7163. 106–122.
- [40] Jens Groth and Amit Sahai. 2008. Efficient Non-interactive Proof Systems for Bilinear Groups. In *Advances in Cryptology – EUROCRYPT 2008 (Lecture Notes in Computer Science)*, Nigel P. Smart (Ed.), Vol. 4965. Springer, Heidelberg, 415–432.
- [41] ECONOLITE Group. 2018. Connected Vehicle CoProcessor Module. <http://www.econolitegroup.com/wp-content/uploads/2017/05/controllers-connectedvehicle-datashet.pdf>. (2018). [Online; accessed 07-April-2018].
- [42] E-ZPass Group. 2017. E-ZPass Statistics: 2005 - 2016. <https://e-zpassag.com/about-us/statistics>. (2017). [Online; accessed May-08-2018].
- [43] Gunnar Hartung, Max Hoffmann, Matthias Nagel, and Andy Rupp. 2017. BBA+: Improving the Security and Applicability of Privacy-Preserving Point Collection. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS 2017, Dallas, TX, USA, October 30 - November 03, 2017*, Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu (Eds.). ACM, 1925–1942. <https://doi.org/10.1145/3133956.3134071>
- [44] Gottfried Herold, Max Hoffmann, Michael Kloof, Carla Ràfols, and Andy Rupp. 2017. New Techniques for Structural Batch Verification in Bilinear Groups with Applications to Groth-Sahai Proofs. In *ACM CCS 17: 24th Conference on Computer and Communications Security*,

- Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu (Eds.). ACM Press, 1547–1564.
- [45] Jaap-Henk Hoepman and George Huitema. 2010. Privacy Enhanced Fraud Resistant Road Pricing. *What Kind of Information Society? Governance, Virtuality, Surveillance, Sustainability, Resilience* (2010), 202–213.
- [46] Kapsch (Toll Collection System Integrator and Supplier). 2018. Personal Communication. (2018).
- [47] Yuval Ishai, Manoj Prabhakaran, and Amit Sahai. 2008. Founding Cryptography on Oblivious Transfer - Efficiently. In *Advances in Cryptology – CRYPTO 2008 (Lecture Notes in Computer Science)*, David Wagner (Ed.), Vol. 5157. Springer, Heidelberg, 572–591.
- [48] Roger Jardí-Cedó, Jordi Castellà-Roca, and Alexandre Viejo. 2014. Privacy-Preserving Electronic Toll System with Dynamic Pricing for Low Emission Zones. In *Data Privacy Management, Autonomous Spontaneous Security and Security Assurance – DPM 2014, SETOP 2014 and QASA 2014. Revised Selected Papers (Lecture Notes in Computer Science)*, Vol. 8872. 327–334.
- [49] Roger Jardí-Cedó, Macià Mut-Puigserver, M. Magdalena Payeras-Capellà, Jordi Castellà-Roca, and Alexandre Viejo. 2014. Electronic Road Pricing System for Low Emission Zones to Preserve Driver Privacy. In *Modeling Decisions for Artificial Intelligence – MDAI 2014. Proceedings*. 1–13.
- [50] Roger Jardí-Cedó, Macià Mut-Puigserver, M. Magdalena Payeras-Capellà, Jordi Castellà-Roca, and Alexandre Viejo. 2016. Privacy-preserving Electronic Road Pricing System for Multifare Low Emission Zones. In *Proceedings of the 9th International Conference on Security of Information and Networks (SIN 2016)*. 158–165.
- [51] Yuto Kawahara, Tetsutaro Kobayashi, Michael Scott, and Akihiro Kato. 2016. *Barreto-Naehrig Curves*. Internet Draft. Internet Engineering Task Force. Work in Progress.
- [52] Florian Kerschbaum and Hoon Wei Lim. 2015. Privacy-Preserving Observation in Public Spaces. In *ESORICS 2015: 20th European Symposium on Research in Computer Security, Part II (Lecture Notes in Computer Science)*, Günther Pernul, Peter Y. A. Ryan, and Edgar R. Weippl (Eds.), Vol. 9327. Springer, Heidelberg, 81–100. [https://doi.org/10.1007/978-3-319-24177-7\\_5](https://doi.org/10.1007/978-3-319-24177-7_5)
- [53] Eike Kiltz, Jiaxin Pan, and Hoeteck Wee. 2015. Structure-Preserving Signatures from Standard Assumptions, Revisited. In *Advances in Cryptology – CRYPTO 2015, Part II (Lecture Notes in Computer Science)*, Rosario Gennaro and Matthew J. B. Robshaw (Eds.), Vol. 9216. Springer, Heidelberg, 275–295. [https://doi.org/10.1007/978-3-662-48000-7\\_14](https://doi.org/10.1007/978-3-662-48000-7_14)
- [54] Hamish Koelmeyer and Sithampanathan Kandeepan. 2017. Tagless Tolling using DSRC for Intelligent Transport System: An Interference Study. In *Asia Modelling Symposium*. IEEE.
- [55] Sebastian Kummer, Maria Dieplinger, and Mario Dobrovnik. 2015. Endbericht der Studie “Flächendeckende Schwerverkehrs-Maut in Österreich”. [https://www.wko.at/branchen/stmk/transport-verkehr/gueterbefoerderungsgewerbe/Endbericht\\_FlaechendeckendeMaut\\_final.pdf](https://www.wko.at/branchen/stmk/transport-verkehr/gueterbefoerderungsgewerbe/Endbericht_FlaechendeckendeMaut_final.pdf) (in German). (2015). [Online; accessed January-09-2019].
- [56] Jorn Lapon, Markulf Kohlweiss, Bart De Decker, and Vincent Naessens. 2010. Performance Analysis of Accumulator-Based Revocation Mechanisms. In *Security and Privacy - Silver Linings in the Cloud - 25th IFIP TC-11 International Information Security Conference, SEC 2010, Held as Part of WCC 2010, Brisbane, Australia, September 20-23, 2010. Proceedings (IFIP Advances in Information and Communication Technology)*, Kai Rannenberg, Vijay Varadarajan, and Christian Weber (Eds.), Vol. 330. Springer, 289–301. [https://doi.org/10.1007/978-3-642-15257-3\\_26](https://doi.org/10.1007/978-3-642-15257-3_26)
- [57] Yunxin Jeff Li. 2010. An overview of the DSRC/WAVE technology. In *International Conference on Heterogeneous Networking for Quality, Reliability, Security and Robustness*. Springer, 544–558.
- [58] Markets and Markets. 2017. Electronic Toll Collection Market Study. <https://www.marketsandmarkets.com/Market-Reports/electronic-toll-collection-system-market-224492059.html>. (2017). [Online; accessed April-19-2018].
- [59] Sarah Meiklejohn, Keaton Mowery, Stephen Checkoway, and Hovav Shacham. 2011. The Phantom Tollbooth: PrivacyPreserving Electronic Toll Collection in the Presence of Driver Collusion. In *Proceedings of the 20th USENIX Security Symposium*.
- [60] Dustin Moody, Rene C. Peralta, Ray A. Perlner, Andrew R. Regenscheid, Allen L. Roginsky, and Lidong Chen. 2015. Report on Pairing-based Cryptography. In *Journal of Research of the National Institute of Standards and Technology*, Vol. 120. National Insitute of Standards and Technology, Gaithersburg, MD, USA, 11–27.
- [61] ABC News. 2008. Toll Records Catch Unfaithful Spouses. <https://abcnews.go.com/Technology/story?id=3468712&page=1>. (2008). [Online; accessed April-19-2018].
- [62] European Parliament. 2014. Technology Options for the European Electronic Toll Service. [http://www.europarl.europa.eu/RegData/etudes/STUD/2014/529058/IPOL\\_STUD\(2014\)529058\\_EN.pdf](http://www.europarl.europa.eu/RegData/etudes/STUD/2014/529058/IPOL_STUD(2014)529058_EN.pdf). (2014). [Online; accessed April-19-2018].
- [63] Raluca Ada Popa, Hari Balakrishnan, and Andrew J. Blumberg. 2009. VPriv: Protecting Privacy in Location-Based Vehicular Services. In *Proceedings of the 18th USENIX Security Symposium*. 335–350.
- [64] Savari.net. 2017. MobiWAVE On-Board-Unit (OBU). [http://savari.net/wp-content/uploads/2017/05/MW-1000\\_April2017.pdf](http://savari.net/wp-content/uploads/2017/05/MW-1000_April2017.pdf). (2017). [Online; accessed 05-February-2018].
- [65] Latanya Sweeney. 2002. k-Anonymity: A Model for Protecting Privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 10, 5 (2002), 557–570. <https://doi.org/10.1142/S0218488502001648>
- [66] American Civil Liberties Union. 2015. Toll Records Catch Unfaithful Spouses. <https://www.aclu.org/blog/privacy-technology/location-tracking/newly-obtained-records-reveal-extensive-monitoring-e-zpass>. (2015). [Online; accessed April-19-2018].

Table 2. Information an adversary learns about honest users.

Protocol	Leakage					
	$pk_{\mathcal{U}}$	$a_{\mathcal{U}}$	$a_{\mathcal{R}}/a_{\mathcal{T}}$	$a_{\mathcal{R}}^{\text{prev}}$	$p$	$b^{\text{bill}}$
User Registration	•					
Wallet Issuing	•	•	•			
Debt Accumulation		•	•	•	•	
Debt Clearance	•	(•)		•		•
Prove Participation	•					

## A INFORMATION LEAKAGE AND DISCUSSION ON PRIVACY IMPLICATIONS

As briefly discussed in [subsection 1.5.2](#), possibly known background information and the leakage of the ideal functionality determines the level of user privacy in P4TC. Since we prove our real protocol  $\pi_{\text{P4TC}}$  to be indistinguishable from the ideal functionality  $\mathcal{F}_{\text{P4TC}}$ , it is ensured that an adversary attacking  $\pi_{\text{P4TC}}$  in the real world can only learn as much about a user as an adversary in the ideal model. [Table 2](#) summarizes what an adversary learns about the users in each task. We omitted the serial number  $s$  and the fraud detection ID  $\phi$  in the table as these are independently and uniformly drawn randomness and thus cannot be exploited. In all tasks except Debt Accumulation the user’s public key  $pk_{\mathcal{U}}$  is leaked. The variables  $a_{\mathcal{U}}$ ,  $a_{\mathcal{R}}$ ,  $a_{\mathcal{R}}^{\text{prev}}$  and  $a_{\mathcal{T}}$  refer to attributes of the participating parties. The variable  $p$  denotes the price of a Debt Accumulation transaction, and  $b^{\text{bill}}$  is the total debt the user owes at the end of the task Debt Clearance.

For every billing period, the TSP collects all transaction information from every RSU. Hence, the TSP eventually possesses two datasets:

- (1) A database of users that are identified by their public key  $pk_{\mathcal{U}}$  together with their attributes and total debt. This dataset comprises all information from every conducted task but Debt Accumulation.
- (2) A database of anonymous transactions. This dataset stems from the Debt Accumulation tasks (cp. [Table 2](#)).

With respect to practical privacy considerations one can naturally pose several questions: Can a single transaction be linked to a specific user? Has a user passed by a particular RSU? Can a user be mapped to a complete track, i.e., a sequence of consecutive transactions? A final answer to these questions crucially depends on the concrete instantiation of the attributes  $a_{\mathcal{U}}$ ,  $a_{\mathcal{R}}$  and the pricing function but also on “environmental” parameters that cannot be chosen by the system designer such as the total number of registered users, the average length of a trip, etc. An in-depth analysis would require plausible and justifiable assumptions about probability distributions for these parameters, and would constitute a separate line of research in its own right.

In the following, however, we would like to elaborate a bit on the general aspects of the question, how a user can be linked to a full track. This problem can be depicted as a graph-theoretical problem of finding a path in a directed, layered graph. The graph consists of initial nodes, inner nodes that are ordered in layers and terminal nodes. Initial nodes represent Wallet Issuing transactions and are linked to a users. Terminal nodes represent Debt Clearance transactions and are also linked to users and final balances  $b^{\text{bill}}$ .

Inner nodes represent the (anonymous) transactions in between. Assuming that transactions can only occur at discrete points in time, the inner nodes can be ordered in layers. A directed edge connects two nodes if the target node is a plausible successor of the source node. As a bare minimum, this requires that the represented transactions have equal user attributes  $a_{\mathcal{U}}$ , the attribute  $a_{\mathcal{R}}^{\text{prev}}$  of the target node equal  $a_{\mathcal{R}}$  of the source node and the target node is in a later layer than the source node (because time can only increase). Additionally, background knowledge such as the geo-position of the RSUs, the given road infrastructure, etc. can be utilized to only insert an edge between two nodes if, e.g., the corresponding RSUs are within a certain distance bound. Obviously, the

average in- and out-degree of a node heavily depends on the distribution of distinct values for  $a_U$  and  $a_R$ . Given this graph, the task is to find a path from the initial node to the terminal node for a particular user such that the sum of the prices of the transactions on this path equals the total balance. Moreover, if the transactions of all users are taken into account, there must be a path for each user such that all paths are pairwise disjoint and every node (i.e., transaction) lies on exactly one path.

For privacy, two characteristics are important: How many solutions do exist and what is the computational complexity to find one (or all) solutions? This results in a trade-off between two borderline cases:

- (1) There is exactly one unique solution. At first glance, this contradicts privacy. However, the mere existence of a unique solution is worthless, if it is computationally infeasible to find it.
- (2) Finding a solution is easy but there are many equally valid solutions. In this case privacy is preserved as well.

If additional background information is omitted, the problem can be cast as a specialized instance of various NP-complete problems, e.g., the parallel-version of the KNAPSACK problem. Parallel KNAPSACK is defined as follows: Let  $\{u_i\}_{i \in \{1, \dots, n\}}$  be a finite set of knapsacks (users) with volume  $b_i := b(u_i)$  (total balance). Further let  $\{t_j\}_{j \in \{1, \dots, m\}}$  be a finite set of items (transactions) with volumes  $p_j := p(t_j)$  (price) respectively. Let  $x_{ij} \in \{0, 1\}$  be variables that indicate if knapsack  $i$  contains item  $j$ . The task is to maximize the objective function

$$\sum_{j=1}^m \sum_{i=1}^n x_{ij} v(t_j) \quad (1)$$

under the constraints

$$\sum_{j=1}^m x_{ij} p_j \leq b_i \text{ for } i \in \{1, \dots, m\} \quad (2)$$

and

$$\sum_{i=1}^n x_{ij} \leq 1 \text{ for } j \in \{1, \dots, m\} \quad (3)$$

Here,  $v$  denotes a value function that assigns a benefit to each item. In our case, we set  $v = \text{const}$ , i.e., each item (transaction) is equally valuable. Informally, this means to pack as many items (transactions) into knapsacks (invoices) without exceeding the volume (balance) of each knapsack (invoice) or assigning an item to more than one knapsack. Since we know that all knapsacks can be filled exactly without any item being left over, we can conclude that any solution that does not completely use a knapsack up to its limit must omit an item. Hence, such a solution is never optimal and equality holds for all optimal solutions in Eq. (2).

Two remarks are in order. With the above explanation, we map an instance of our track-finding problem onto an instance of the KNAPSACK problem, i.e., this only shows that our problem at hand is not harder than KNAPSACK. To be fully correct, we would have to show the inverse direction, namely that a *restricted set* of instances of the parallel KNAPSACK problem can be cast into our track-finding problem (without background knowledge). The *general* KNAPSACK problem is NP-complete. This is beneficial as it implies that finding a solution is generally believed to be intractable. However, there might be good heuristics for all “natural” instances. Especially since we only need to consider those instances for which the optimal solution is a perfect solution (i.e., each knapsack is completely filled). Nonetheless, this restricted class of KNAPSACK problems is still NP-complete and can indeed be mapped onto our track-finding problem. Moreover, depending on the concrete parameters (e.g., an upper bound on the maximum price  $p$  or the balance  $b^{\text{bill}}$ ) the problem might become fixed-parameter tractable [4]. In other words, although solving the general problem has super-polynomial runtime in the instance

size, it might still be practically solvable for “real world” instances. We stress again, that an in-depth analysis requires to look at concrete distributions of these parameters which may be the basis for an independent work.

Nonetheless, there are indicators that—if finding one solution is easy—there might be a myriad of solutions, which again yields privacy. We take the German Toll Collect<sup>17</sup> for trucks and analyze the statistics from January, 2018 [35], as a more concrete example. This system uses 24 distinct<sup>18</sup> values for the user attributes. We assume that transaction times are recorded with 5 min resolution which is a little bit larger than the timespan a truck needs to travel between two subsequent RSUs.<sup>19</sup> Within this timeslot each RSU is passed by 28.4 other trucks. Picking a fixed RSU and a fixed timeslot, this implies that the probability that at least one other truck with identical attributes passes by equals

$$p = 1 - \left(\frac{23}{24}\right)^{28.4} \approx 0.7 \quad (4)$$

This means, for a randomly chosen transaction in the graph the probability to find at least one other transaction in the same layer (i.e., timeslot) for the same RSU and for identical user attributes equals 0.7. In other words, for a randomly chosen node in the transaction graph the in- and out-degree is at least two with probability 0.7.

The average length of trips per month of a single vehicle equals 3 398 km and the vehicle passes 613 RSUs. This means the average trip consists of 613 transactions. For approximately  $0.7 \cdot 613 = 429$  of these transactions there are at least two identically looking transactions that are a plausible predecessor or successor. However, it is completely wrong to assume this would yield  $2^{429}$  identically looking paths that could be assigned to a particular user. This kind of analysis would assume an independence of the transactions which does not hold. Again, let us randomly pick a particular RSU and timeslot and consider the geographically next RSU in the following time slot. If two identical transactions are observed at the first RSU, the *conditional probability* to observe two identical transactions at the next RSU again is assumedly much higher than 0.7, given that there are no entry points, exit points or intersections in between. Vice versa, if an transaction is unique at the first RSU, the *conditional probability* to observe a second, identical transaction at the next RSU is assumedly much lower. On the contrary, we expect to have some transaction nodes with a high in-/out-degree, namely those before an exit point or after an entry point. If a truck leaves the tolling system, it might either promptly re-enter the system at some nearby entry point or at some distant entry point in the remote future.

In summary, the analysis above is overly simplified as we assume uniform distributions of all values and an independence of random variables that is likely not given in reality. Nonetheless, this indicates that the solution space for mapping a particular user to a specific trip might be vast and it points out why an in-depth analysis would justify an independent work on its own. Please also note that this analysis is only based on trucks not passenger cars where we expect much higher probabilities for joint occurrence of identical attributes due to the higher number of participants and higher travel speeds.

*Remark A.1.* In practice, several privacy notions like  $k$ -anonymity are established. For several reasons these notions are not directly applicable here. First of all, these notions evaluate the privacy level of a concrete dataset and we stress again that this is out of the scope of this work. While at first glance the calculations above might suggest that our system features  $k$ -anonymity [65] for some yet to be determined  $k$ , the notion of  $k$ -anonymity is actually not applicable due to formal reasons. The definition of  $k$ -anonymity requires the database to have exactly one entry for each individual, but our transaction database features several entries per user. Therefore, the notion of  $k$ -anonymity is syntactically not applicable to the users of our system. While we could still discuss

<sup>17</sup><https://www.toll-collect.de/en/>

<sup>18</sup>The actual system uses more attribute values (i.e., 720) for statistical purposes. However, we only counted attribute values that the price depends on.

<sup>19</sup>Since the German toll collection system is GNSS-based and has no RSUs, the average distance between two RSUs is assumed to be 5 km as in the Austrian Toll Collection system [55]. Additionally, we assumed  $80 \frac{\text{km}}{\text{h}}$  travel speed for trucks.

### Functionality $\mathcal{F}_{\text{P4TC}}$

#### I. State

- Set  $TRDB = \{trdb\}$  of transactions

$$trdb = (s^{\text{prev}}, s, \phi, x, \lambda, pid_{\mathcal{U}}, pid_{\mathcal{R}}, p, b) \\ \in S \times S \times \Phi \times \mathbb{N}_0 \times \mathcal{L} \times \mathcal{PID}_{\mathcal{U}} \times \mathcal{PID}_{\mathcal{R}} \times \mathbb{Z}_p \times \mathbb{Z}_p.$$

- A (partial) mapping  $f_{\Phi}$  giving the fraud detection ID  $\phi$  corresponding to given wallet ID  $\lambda$  and counter  $x$ :

$$f_{\Phi} : \mathcal{L} \times \mathbb{N}_0 \rightarrow \Phi, (\lambda, x) \mapsto \phi$$

- A (partial) mapping  $f_{\mathcal{A}_{\mathcal{U}}}$  assigning user attributes to a given wallet ID  $\lambda$ :

$$f_{\mathcal{A}_{\mathcal{U}}} : \mathcal{L} \rightarrow \mathcal{A}_{\mathcal{U}}, \lambda \mapsto \mathbf{a}_{\mathcal{U}}$$

- A (partial) mapping  $f_{\mathcal{A}_{\mathcal{R}}}$  assigning RSU attributes to a given RSU PID  $pid_{\mathcal{R}}$ :

$$f_{\mathcal{A}_{\mathcal{R}}} : \mathcal{PID}_{\mathcal{R}} \rightarrow \mathcal{A}_{\mathcal{R}}, pid_{\mathcal{R}} \mapsto \mathbf{a}_{\mathcal{R}}$$

- A (partial) mapping  $f_{\Pi}$  assigning a user PID  $pid_{\mathcal{U}}$  and a proof of guilt  $\pi$  to a validity bit:

$$f_{\Pi} : \mathcal{PID}_{\mathcal{U}} \times \Pi \rightarrow \{\text{OK}, \text{NOK}\}$$

#### II. Behavior

- |                              |                               |                                       |
|------------------------------|-------------------------------|---------------------------------------|
| • DR Registration (Fig. 5)   | • Wallet Issuing (Fig. 10)    | • Prove Participation (Fig. 13)       |
| • TSP Registration (Fig. 6)  | • Debt Accumulation (Fig. 11) | • Double-Spending Detection (Fig. 14) |
| • RSU Registration (Fig. 7)  | • Debt Clearance (Fig. 12)    | • Guilt Verification (Fig. 15)        |
| • User Registration (Fig. 8) |                               | • User Blacklisting (Fig. 16)         |
| • RSU Certification (Fig. 9) |                               |                                       |

Fig. 4. The functionality  $\mathcal{F}_{\text{P4TC}}$

$k$ -anonymity in this setting if the TSP combined all entries that pertain to the same user into one single entry, privacy of our system largely stems from the TSP not being able to link transactions of the same user in this way and hence such a discussion would largely undervalue the privacy protection P4TC provides.

## B FULL SYSTEM DEFINITION

In this appendix we give a detailed description and explanation of our ideal privacy-preserving electronic toll collection functionality  $\mathcal{F}_{\text{P4TC}}$ . As explained before we define this as a monolithic, reactive functionality with polynomially many parties. This is mainly due to a shared state that the system requires. We will therefore first explain how this state is recorded by  $\mathcal{F}_{\text{P4TC}}$  before we go on to describe its behavior in a modular way by explaining each task<sup>20</sup> it provides.

The main feature of  $\mathcal{F}_{\text{P4TC}}$  is that it keeps track of all conducted transactions in a global transaction database  $TRDB$  (see Fig. 4). Note that in this case by “transaction” we mean every instance of the tasks Wallet Issuing, Debt Accumulation or Debt Clearance, not just Debt Accumulation. Each transaction entry  $trdb \in TRDB$  is of the form

$$trdb = (s^{\text{prev}}, s, \phi, x, \lambda, pid_{\mathcal{U}}, pid_{\mathcal{R}}, p, b).$$

<sup>20</sup>Note that we are intentionally avoiding the word “phase”—which is commonly used in other composite functionalities—as it suggests a predefined order/number of executions.



Table 3. Notation that only occurs in the ideal functionality

Identifier	Description
$PID_{\text{corrupt}}$	set of corrupted party identifiers
$f_\phi$	(partial) mapping giving the fraud detection ID $\phi$ corresponding to given wallet ID $\lambda$ and counter $x$
$f_{\mathcal{A}_U}$	(partial) mapping assigning user attributes to a given wallet ID $\lambda$
$f_{\mathcal{A}_R}$	(partial) mapping $f_{\mathcal{A}_R}$ assigning RSU attributes to a given RSU PID $pid_{\mathcal{R}}$

It contains the identities  $pid_{\mathcal{U}}$  and  $pid_{\mathcal{R}}$  of the involved user and RSU (or TSP in the case of Wallet Issuing and Debt Clearance) respectively, the ID  $\lambda$  of the wallet that was used as well as the price  $p$  and total balance  $b$  of the wallet state after this transaction. Furthermore, each transaction entry is identified by a unique serial number  $s$  and links via  $s^{\text{prev}}$  to the previous transaction  $trdb^{\text{prev}}$  (which corresponds to the wallet state before  $trdb$ ). Lastly, a fraud detection ID  $\phi$  and a counter  $x$  are part of the transaction entry. The counter starts at zero for any newly registered wallet and  $x = (x^{\text{prev}} + 1)$  always holds. Hence, it is unique across all wallet states belonging to the same wallet  $\lambda$  if and only if no double-spending has been committed with this wallet. The fraud detection ID is constant for each pair  $(\lambda, x)$  of wallet ID and counter instead of being unique for each transaction, but unique for different pairs of wallet ID and counter. Therefore, fraud detection IDs are stored in a partially defined, but one-to-one, mapping  $f_\phi : (\mathcal{L} \times \mathbb{N}_0) \rightarrow \Phi$  within  $\mathcal{F}_{\text{P4TC}}$ . Full transaction entries  $trdb$  are only created by instances of Debt Accumulation. Both Wallet Issuing and Debt Clearance create stubs of the form

$$(\perp, s, \phi, 0, \lambda, pid_{\mathcal{U}}, pid_{\mathcal{T}}, 0, 0) \quad \text{and} \\ (s^{\text{prev}}, \perp, \phi, x, \lambda, pid_{\mathcal{U}}, pid_{\mathcal{T}}, -b^{\text{bill}}, 0)$$

respectively. Every other task does not alter  $TRDB$  but only queries it. Although this database and the mapping to fraud detection IDs contains most of the information our toll collection scheme needs,  $\mathcal{F}_{\text{P4TC}}$  stores three more partially defined mappings:  $f_{\mathcal{A}_U} : \mathcal{L} \rightarrow \mathcal{A}_U$  and  $f_{\mathcal{A}_R} : PID_{\mathcal{R}} \rightarrow \mathcal{A}_R$  of user and RSU attribute vectors as well as  $f_\Pi$  of proofs of guilt that have been issued or queried in the context of double-spending detection.

The ideal function  $\mathcal{F}_{\text{P4TC}}$  provides twelve different tasks in total which we divide up into three categories: “System Setup Tasks” (comprising all Registrations and RSU Certification), “Basic Tasks” (Wallet Issuing, Debt Accumulation and Debt Clearance), and “Feature Tasks” (Prove Participation, Double-Spending Detection, Guilt Verification and User Blacklisting).

For better clarity of the following task descriptions, an overview of the variables used can be found in [Tables 3 and 4](#).

## B.1 System Setup Tasks

To set up the system two things are required: All parties—the DR, TSP, RSUs and users—have to register public keys with the bulletin board  $\overline{\mathcal{G}}_{\text{bb}}$  to be able to participate in the toll collection system. As all of these registration tasks are similar, we will not describe them separately. In the special case of RSUs a certification conducted with the TSP also needs to take place.

**B.1.1 Registrations.** The tasks of DR, RSU and User Registration (cp. [Figs. 5 to 8](#)) are straightforward and analogous. They do not take any input apart from “register”, but in the case of the user we assume the physical identity of the party has been verified out-of-band before this task is conducted. In each case a check is performed first whether the task has been run before for this party. If this does not lead to an abort, the adversary is asked to

Table 4. Notation that occurs in the ideal functionality and in the real protocol

Identifier	Abstract Domain	Instantiation	Description
$pid_{DR}$	$\mathcal{PID}_{DR}$	$\{0, 1\}^*$	party identifier of the DR
$pid_{\mathcal{T}}$	$\mathcal{PID}_{\mathcal{T}}$	$\{0, 1\}^*$	party identifier of the TSP
$pid_{\mathcal{R}}$	$\mathcal{PID}_{\mathcal{R}}$	$\{0, 1\}^*$	party identifier of a RSU
$pid_{\mathcal{U}}$	$\mathcal{PID}_{\mathcal{U}}$	$\{0, 1\}^*$	party identifier of a user
$pk_{DR}$	$\mathcal{PK}_{DR}$	$G_1^3 \times G_2^3 \times (G_1^2)^{\ell+2} \times (G_2^2)^4 \times (G_2^2)^{\ell+2}$	public DR key
$pk_{\mathcal{T}}$	$\mathcal{PK}_{\mathcal{T}}$	$G_1^{j+3} \times (G_1^3 \times G_2^{y+3}) \times (G_1^3 \times G_2)$	public TSP key
$pk_{\mathcal{R}}$	$\mathcal{PK}_{\mathcal{R}}$	$G_1^3 \times G_2$	public RSU key
$pk_{\mathcal{U}}$	$\mathcal{PK}_{\mathcal{U}}$	$G_1$	public user key
$a_{\mathcal{U}}$	$\mathcal{A}_{\mathcal{U}}$	$G_2^j$	user attributes
$a_{\mathcal{R}}$	$\mathcal{A}_{\mathcal{R}}$	$G_1^y$	RSU attributes
$a_{\mathcal{T}}$	$\mathcal{A}_{\mathcal{R}}$	$G_1^y$	TSP attributes
$b$	$\mathbb{Z}_p$	$\mathbb{Z}_p$	balance
$p$	$\mathbb{Z}$	$\mathbb{Z}$	price to pay at an RSU
$\phi$	$\Phi$	$G_1$	fraud detection ID
$s$	$S$	$G_1$	serial number
$\lambda$	$\mathcal{L}$	$\mathbb{Z}_p$	wallet ID; is used as PRF seed
$x$	$\mathbb{N}_0$	$\{0, \dots, n_{\text{PRF}}\}$	(PRF) counter
$bl_{\mathcal{R}}$	list of $\Phi$ elements	list of $G_1$ elements	RSU blacklist
$x_{bl_{\mathcal{R}}}$	$\mathbb{N}$	$\mathbb{N}$	RSU blacklist parameter
$bl_{\mathcal{T}}$	list of $\mathcal{PK}_{\mathcal{U}}$ elements	list of $G_1$ elements	TSP blacklist

### Functionality $\mathcal{F}_{P4TC}$ (cont.) – Task DR Registration

DR input: (register)

- (1) If this task has been run before, output  $\perp$  and abort.
- (2) Send (registering\_dr,  $pid_{DR}$ ) to the adversary and obtain the key ( $pk_{DR}$ ).<sup>a</sup>
- (3) Call  $\overline{\mathcal{G}}_{bb}$  with input (register,  $pk_{DR}$ ).

DR output: ( $pk_{DR}$ )

<sup>a</sup>Giving the adversary the power to control the key generation serves two purposes: i) It gives a stronger security guarantee, i.e., security for a honest TSP is retained, even if its keys are maliciously generated (due to bad random number generator). ii) It gives the simulator the lever to simulate faithfully.

Fig. 5. The functionality  $\mathcal{F}_{P4TC}$  (cont. from Fig. 4)

**Functionality  $\mathcal{F}_{P4TC}$  (cont.) – Task TSP Registration***TSP input:* (register,  $a_{\mathcal{T}}$ )

- (1) If this task has been run before, output  $\perp$  and abort.
- (2)  $TRDB := \emptyset$
- (3) Send (registering\_tsp,  $pid_{\mathcal{T}}$ ,  $a_{\mathcal{T}}$ ) to the adversary and obtain the key  $(pk_{\mathcal{T}})$ .<sup>a</sup>
- (4) Call  $\overline{\mathcal{G}}_{bb}$  with input (register,  $pk_{\mathcal{T}}$ ).

*TSP output:* ( $pk_{\mathcal{T}}$ )

<sup>a</sup>Giving the adversary the power to control the key generation serves two purposes: i) It gives a stronger security guarantee, i.e., security for a honest TSP is retained, even if its keys are maliciously generated (due to bad random number generator). ii) It gives the simulator the lever to simulate faithfully.

Fig. 6. The functionality  $\mathcal{F}_{P4TC}$  (cont. from Fig. 4)**Functionality  $\mathcal{F}_{P4TC}$  (cont.) – Task RSU Registration***RSU input:* (register)

- (1) If this task has been run before, output  $\perp$  and abort.
- (2) Send (registering\_rsu,  $pid_{\mathcal{R}}$ ) to the adversary and obtain the key  $(pk_{\mathcal{R}})$ .<sup>a</sup>
- (3) Call  $\overline{\mathcal{G}}_{bb}$  with input (register,  $pk_{\mathcal{R}}$ ).

*RSU output:* ( $pk_{\mathcal{R}}$ )

<sup>a</sup>Giving the adversary the power to control the key generation serves two purposes: i) It gives a stronger security guarantee, i.e., security for honest parties is retained, even if their keys are maliciously generated (due to bad random number generator). ii) It gives the simulator the lever to simulate faithfully.

Fig. 7. The functionality  $\mathcal{F}_{P4TC}$  (cont. from Fig. 4)**Functionality  $\mathcal{F}_{P4TC}$  (cont.) – Task User Registration***User input:* (register)

- (1) If this task has been run before, output  $\perp$  and abort.
- (2) Send (registering\_user,  $pid_{\mathcal{U}}$ ) to the adversary and obtain the key  $(pk_{\mathcal{U}})$ .<sup>a</sup>
- (3) Call  $\overline{\mathcal{G}}_{bb}$  with input (register,  $pk_{\mathcal{U}}$ ).

*User output:* ( $pk_{\mathcal{U}}$ )

<sup>a</sup>Giving the adversary the power to control the key generation serves two purposes: i) It gives a stronger security guarantee, i.e., security for honest parties is retained, even if their keys are maliciously generated (due to bad random number generator). ii) It gives the simulator the lever to simulate faithfully.

Fig. 8. The functionality  $\mathcal{F}_{P4TC}$  (cont. from Fig. 4)

provide a public key  $pk$  for the respective party which is then registered with the bulletin board  $\overline{\mathcal{G}}_{bb}$  and output to the newly registered party.

The registration of the TSP is slightly different (cp. Fig. 6). In addition to “register” it takes an attribute vector  $a_{\mathcal{T}}$  as input, which—after a check if this task has been run before—is leaked to the adversary together with  $pid_{\mathcal{T}}$

**Functionality  $\mathcal{F}_{\text{P4TC}}$  (cont.) – Task *RSU Certification***

*RSU input:* (certify)

*TSP input:* (certify,  $\mathbf{a}_{\mathcal{R}}$ )

- (1) If  $f_{\mathcal{A}_{\mathcal{R}}}(pid_{\mathcal{R}})$  is already defined, output  $\perp$  to both parties and abort; else append  $f_{\mathcal{A}_{\mathcal{R}}}(pid_{\mathcal{R}}) := \mathbf{a}_{\mathcal{R}}$  to  $f_{\mathcal{A}_{\mathcal{R}}}$ .
- (2) Leak (certifying\_rsu,  $pid_{\mathcal{R}}$ ,  $\mathbf{a}_{\mathcal{R}}$ ) to the adversary.

*RSU output:* ( $\mathbf{a}_{\mathcal{R}}$ )

*TSP output:* (OK)

Fig. 9. The functionality  $\mathcal{F}_{\text{P4TC}}$  (cont. from Fig. 4)

when the public key  $pk_{\mathcal{T}}$  is obtained. In addition, all data structures of  $\mathcal{F}_{\text{P4TC}}$  are initialized as empty sets and empty (partial) mappings respectively. Again, the public key  $pk_{\mathcal{T}}$  is output to the TSP.

**B.1.2 RSU Certification.** RSU certification (cp. Fig. 9) is a two-party task between the TSP and an RSU in which the RSU is assigned an attribute vector  $\mathbf{a}_{\mathcal{R}}$ .<sup>21</sup> The content of attribute vectors is in no way restricted by  $\mathcal{F}_{\text{P4TC}}$  and can be used to implement different scenarios like location-based toll collection or entry-exit toll collection, but could also be maliciously used to void unlinkability. This  $\mathbf{a}_{\mathcal{R}}$  is input by the TSP, while the RSU only inputs its desire to be certified.  $\mathcal{F}_{\text{P4TC}}$  checks if there have already been attributes assigned to the RSU previously (in which case it aborts) and otherwise appends  $f_{\mathcal{A}_{\mathcal{R}}}(pid_{\mathcal{R}}) := \mathbf{a}_{\mathcal{R}}$  to the partial mapping  $f_{\mathcal{A}_{\mathcal{R}}}$  which internally stores all RSU attributes already assigned. The identity  $pid_{\mathcal{R}}$  and attributes  $\mathbf{a}_{\mathcal{R}}$  are leaked to the adversary before the attributes are output to the RSU.

## B.2 Basic Tasks

In this section we describe the basic tasks you would expect from any toll collection scheme. These tasks are Wallet Issuing, Debt Accumulation and Debt Clearance. As mentioned before, those are the only tasks in which transaction entries are created.

**B.2.1 Wallet Issuing.** Wallet Issuing (cp. Fig. 10) is a two-party task between a user and the TSP in which a new and empty wallet is created for the user. The TSP inputs an attribute vector  $\mathbf{a}_{\mathcal{U}}$  and a blacklist  $bl_{\mathcal{T}}$  of user public keys that are not allowed to obtain any new wallets. First,  $\mathcal{F}_{\text{P4TC}}$  randomly picks a (previously unused) serial number  $s$  for the new transaction entry  $trdb$ . If the user is corrupted, the adversary may at this point choose another corrupted user's identity  $pid_{\mathcal{U}}$  that is to be used for this wallet. Multiple corrupted users are allowed to have wallets issued for one another but are not able to request a new wallet for an honest user. The corresponding public key for the user ID  $pid_{\mathcal{U}}$  is obtained from the bulletin board  $\overline{\mathcal{G}}_{\text{bb}}$  and checked against the TSP's blacklist  $bl_{\mathcal{T}}$ . If this does not lead to an abort, a new wallet ID  $\lambda$  and fraud detection ID  $\phi$  are uniquely and randomly picked, unless the user is corrupted in which case the adversary chooses  $\phi$ . This may infringe upon the unlinkability of the user's transactions and we do not give any privacy guarantees for corrupted users. Finally, a transaction entry

$$trdb := (\perp, s, \phi, 0, \lambda, pid_{\mathcal{U}}, pid_{\mathcal{T}}, 0, 0)$$

corresponding to the new and empty wallet is stored in  $TRDB$  and the wallet's attributes  $f_{\mathcal{A}_{\mathcal{U}}}(\lambda) := \mathbf{a}_{\mathcal{U}}$  are appended to the partial mapping  $f_{\mathcal{A}_{\mathcal{U}}}$ . Both parties get the serial number  $s$  as output; the user also receives

<sup>21</sup>Although only attributes are set in this task, we will later see in the task of Debt Accumulation that these also serve as a kind of certificate, as RSUs are only able to successfully participate in Debt Accumulation if they have been assigned attributes by the TSP.

**Functionality  $\mathcal{F}_{\text{P4TC}}$  (cont.) – Task *Wallet Issuing****User input:* (issue)*TSP input:* (issue,  $\mathbf{a}_U, bl_{\mathcal{T}}$ )

- (1) Pick serial number  $s \xleftarrow{R} S$  that has not previously been used.
- (2) If  $pid_U \in \mathcal{PID}_{\text{corrupt}}$  leak  $(s, \mathbf{a}_U)$  to the adversary, and ask if another PID  $pid_U \in \mathcal{PID}_{\text{corrupt}}$  should be used instead.<sup>a</sup>
- (3) Receive  $pk_U$  from the bulletin-board  $\overline{\mathcal{G}}_{\text{bb}}$  for PID  $pid_U$ .<sup>†</sup>
- (4) If  $pk_U \in bl_{\mathcal{T}}$ , output blacklisted to both parties and abort.
- (5) Pick wallet ID  $\lambda \xleftarrow{R} \mathcal{L}$  that has not previously been used.
- (6) If  $pid_U \notin \mathcal{PID}_{\text{corrupt}}$  pick  $\phi \xleftarrow{R} \Phi$  that has not previously been used, otherwise ask the adversary for a fraud detection ID  $\phi$  that has not previously been used.<sup>b</sup> Append  $f_{\Phi}(\lambda, 0) := \phi$  to  $f_{\Phi}$ .
- (7) Append  $trdb := (\perp, s, \phi, 0, \lambda, pid_U, pid_{\mathcal{T}}, 0, 0)$  to  $TRDB$
- (8) Append  $f_{\mathcal{A}_U}(\lambda) := \mathbf{a}_U$  to  $f_{\mathcal{A}_U}$ .

*User output:* ( $s, \mathbf{a}_U$ )*TSP output:* ( $s$ )<sup>†</sup>If this does not exist, output  $\perp$  and abort.<sup>a</sup>If a group of corrupted users collude, a correct mapping to a specific users cannot be guaranteed, because corrupted users might share their credentials.<sup>b</sup>Picking the upcoming fraud detection ID randomly asserts untrackability for honest users. For corrupted user, we do not (and cannot) provide such a guarantee.Fig. 10. The functionality  $\mathcal{F}_{\text{P4TC}}$  (cont. from Fig. 4)

the attribute vector  $\mathbf{a}_U$  to check this has been assigned correctly and more importantly does not contain any identifying information.

**B.2.2 Debt Accumulation.** This two-party task (cp. Fig. 11) is conducted whenever a registered user passes an RSU and it serves the main purpose of adding toll to a previous wallet state of the user. In this task the user only inputs a serial number  $s^{\text{prev}}$ , indicating which past wallet state he wishes to use for this transaction. The participating RSU in turn inputs a blacklist  $bl_{\mathcal{R}}$  of fraud detection IDs. First,  $\mathcal{F}_{\text{P4TC}}$  randomly picks a (previously unused) serial number  $s$  for the new transaction entry  $trdb$ . If the user is corrupted, the adversary may at this point choose another corrupted user's identity  $pid_U$  that is to be used for this transaction.  $\mathcal{F}_{\text{P4TC}}$  looks up if a wallet state  $trdb^{\text{prev}}$  in  $TRDB$  corresponds to the user input  $s^{\text{prev}}$  and belongs to the users  $pid_U$ . This guarantees that each user can only accumulate debt on a wallet that was legitimately issued to him. Multiple corrupted users may choose to swap wallets between them but are not able to use an honest user's wallet. From the previous wallet state

$$trdb^{\text{prev}} = (\cdot, s^{\text{prev}}, \phi^{\text{prev}}, x^{\text{prev}}, \lambda^{\text{prev}}, pid_U, pid_{\mathcal{R}}^{\text{prev}}, \cdot, b^{\text{prev}})$$

the ideal function gets a fraud detection ID  $\phi^{\text{prev}}$  that is checked against  $bl_{\mathcal{R}}$ . The other information in  $trdb^{\text{prev}}$  is then used to determine the content of the new transaction entry  $trdb$ . The user ID  $pid_U$  and wallet ID  $\lambda$  stay the same,  $pid_{\mathcal{R}}$  is set to the identity of the participating RSU, and the counter  $x^{\text{prev}}$  is increased by one to obtain  $x$ .  $\mathcal{F}_{\text{P4TC}}$  checks if there is already a fraud detection ID  $\phi := f_{\Phi}(\lambda, x)$  assigned to the pair  $(\lambda, x)$  (either because the user committed double-spending or because it has been precalculated for blacklisting purposes). If not and the user is honest, it picks a new  $\phi$  uniquely at random. If the user is corrupted, the fraud detection ID is not randomly drawn but picked by the adversary. This may infringe upon the unlinkability of the user's transactions, but, as mentioned before, we do not give any privacy guarantees for corrupted users. The attributes  $\mathbf{a}_U$ ,  $\mathbf{a}_{\mathcal{R}}$  and

**Functionality  $\mathcal{F}_{P4TC}$  (cont.) – Task *Debt Accumulation***

*User input:* (pay\_toll,  $s^{\text{prev}}$ )

*RSU input:* (pay\_toll,  $bl_{\mathcal{R}}$ )

- (1) Pick serial number  $s \xleftarrow{R} S$  that has not previously been used.
- (2) If  $pid_{\mathcal{U}} \in \mathcal{PID}_{\text{corrupt}}$  ask the adversary, if another PID  $pid_{\mathcal{U}} \in \mathcal{PID}_{\text{corrupt}}$  should be used instead.<sup>a</sup>
- (3) Select  $(\cdot, s^{\text{prev}}, \phi^{\text{prev}}, x^{\text{prev}}, \lambda^{\text{prev}}, pid_{\mathcal{U}}, pid_{\mathcal{R}}^{\text{prev}}, \cdot, b^{\text{prev}}) \in TRDB$  (with  $s^{\text{prev}}, pid_{\mathcal{U}}$  being the unique key)<sup>†</sup>.
- (4) If  $\phi^{\text{prev}} \in bl_{\mathcal{R}}$ , output blacklisted to both parties and abort.
- (5) Set  $\lambda := \lambda^{\text{prev}}$  and  $x := x^{\text{prev}} + 1$ .
- (6) If  $f_{\Phi}(\lambda, x)$  is already defined, set  $\phi := f_{\Phi}(\lambda, x)$ .  
Else, if  $pid_{\mathcal{U}} \notin \mathcal{PID}_{\text{corrupt}}$  pick  $\phi \xleftarrow{R} \Phi$  that has not previously been used, otherwise ask the adversary for a fraud detection ID  $\phi$  that has not previously been used.<sup>b</sup> Append  $f_{\Phi}(\lambda, x) := \phi$  to  $f_{\Phi}$ .
- (7) Set  $\mathbf{a}_{\mathcal{U}} := f_{\mathcal{A}_{\mathcal{U}}}(\lambda)$ ,  $\mathbf{a}_{\mathcal{R}} := f_{\mathcal{A}_{\mathcal{R}}}(pid_{\mathcal{R}})$ , and  $\mathbf{a}_{\mathcal{R}}^{\text{prev}} := f_{\mathcal{A}_{\mathcal{R}}}(pid_{\mathcal{R}}^{\text{prev}})$ .<sup>‡</sup>
- (8) Calculate price  $p := \text{O}_{\text{pricing}}(\mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}})$ . If  $pid_{\mathcal{R}} \in \mathcal{PID}_{\text{corrupt}}$ , then leak  $(\mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}})$  to the adversary and obtain a price  $p$ .
- (9)  $b := b^{\text{prev}} + p$ .
- (10) Append  $(s^{\text{prev}}, s, \phi, x, \lambda, pid_{\mathcal{U}}, pid_{\mathcal{R}}, p, b)$  to  $TRDB$ .

*User output:* ( $s, \mathbf{a}_{\mathcal{R}}, p, b$ )

*RSU output:* ( $s, \phi, \mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}}$ )

<sup>†</sup>If this does not exist, output  $\perp$  and abort.

<sup>a</sup>If a group of corrupted users collude, a correct mapping to a specific users cannot be guaranteed, because corrupted users might share their credentials.

<sup>b</sup>The ideal model only guarantees privacy for honest users. For corrupted users the fraud detection ID might be chosen adversarially (cp. text body).

Fig. 11. The functionality  $\mathcal{F}_{P4TC}$  (cont. from Fig. 4)

$\mathbf{a}_{\mathcal{R}}^{\text{prev}}$  are again looked up internally and leaked to the adversary who chooses the price  $p$  of this transaction. Having the price determined in this way makes it clear that  $\mathcal{F}_{P4TC}$  does not give any guarantees on the “right” amount of debt being added at this point. Instead, it gives the user enough information about the transaction to appeal out-of-band afterwards if the wrong amount of debt is added. We assume this detectability will keep RSUs in the real world from adding too much debt. Finally, the new balance  $b$  is calculated from the price and old balance before  $trdb$  is stored in  $TRDB$ . Note that all information leading to the new wallet state came from data internally stored in  $\mathcal{F}_{P4TC}$  itself, not from an input by the user or RSU, and can therefore not be compromised. The serial number, RSU attributes, price and balance are output to the user so he may check he only paid the amount he expected. The RSU gets the serial number as well but also the fraud detection ID to enable double-spending detection and the attributes of the user and previous RSU.

**B.2.3 Debt Clearance.** As Debt Clearance (cp. Fig. 12) is very similar to the task of Debt Accumulation, we will refrain from describing it again in full detail but rather just highlight the differences to Debt Accumulation. The first difference is that it is conducted with the TSP rather than an RSU and no blacklist is taken as input as we do not want to prevent anyone from paying their debt. Although this task results in a transaction entry  $trdb$  as well, no new serial number  $s$  is picked. This emphasizes that the new wallet state is final and can not be updated again by using its serial number as input for another transaction. Instead of obtaining a price from the adversary, the attributes  $\mathbf{a}_{\mathcal{R}}^{\text{prev}}$  of the previous RSU  $pid_{\mathcal{R}}^{\text{prev}}$  are leaked to the adversary in case the TSP is corrupted. The (negative) price for the transaction entry is set to the billing amount  $b^{\text{bill}}$  which in turn is taken to be the previous

**Functionality  $\mathcal{F}_{\text{P4TC}}$  (cont.) – Task *Debt Clearance****User input:* (clear\_debt,  $s^{\text{prev}}$ )*TSP input:* (clear\_debt)

- (1) Pick serial number  $s \xleftarrow{R} S$  that has not previously been used.
- (2) If  $pid_{\mathcal{U}} \in \mathcal{PID}_{\text{corrupt}}$  ask the adversary, if another PID  $pid_{\mathcal{U}} \in \mathcal{PID}_{\text{corrupt}}$  should be used instead.<sup>a</sup>
- (3) Receive  $pk_{\mathcal{U}}$  from the bulletin-board  $\overline{\mathcal{G}}_{\text{bb}}$  for PID  $pid_{\mathcal{U}}$ .<sup>†</sup>
- (4) Select  $(\cdot, s^{\text{prev}}, \phi^{\text{prev}}, x^{\text{prev}}, \lambda^{\text{prev}}, pid_{\mathcal{U}}, pid_{\mathcal{R}}^{\text{prev}}, \cdot, b^{\text{prev}}) \in TRDB$  (with  $s^{\text{prev}}, pid_{\mathcal{U}}$  being the unique key).<sup>†</sup>
- (5) Set  $\lambda := \lambda^{\text{prev}}$  and  $x := x^{\text{prev}} + 1$ .
- (6) If  $f_{\Phi}(\lambda, x)$  is already defined, set  $\phi := f_{\Phi}(\lambda, x)$ .  
Else, if  $pid_{\mathcal{U}} \notin \mathcal{PID}_{\text{corrupt}}$  pick  $\phi \xleftarrow{R} \Phi$  that has not previously been used, otherwise ask the adversary for a fraud detection ID  $\phi$  that has not previously been used.<sup>b</sup> Append  $f_{\Phi}(\lambda, x) := \phi$  to  $f_{\Phi}$ .
- (7) If  $pid_{\mathcal{T}} \in \mathcal{PID}_{\text{corrupt}}$ , set  $a_{\mathcal{R}}^{\text{prev}} := f_{\mathcal{A}_{\mathcal{R}}}(pid_{\mathcal{R}}^{\text{prev}})^{\perp}$ , and leak  $a_{\mathcal{R}}^{\text{prev}}$  to the adversary.
- (8)  $b^{\text{bill}} := b^{\text{prev}}$ .
- (9) Append  $(s^{\text{prev}}, s, \phi, x, \lambda, pid_{\mathcal{U}}, pid_{\mathcal{T}}, -b^{\text{bill}}, 0)$  to  $TRDB$ .

*User output:* ( $b^{\text{bill}}$ )*TSP output:* ( $pk_{\mathcal{U}}, \phi, b^{\text{bill}}$ )<sup>†</sup>If this does not exist, output  $\perp$  and abort.<sup>a</sup>If a group of corrupted users collude, a correct mapping to a specific users cannot be guaranteed, because corrupted users might share their credentials.<sup>b</sup>The ideal model only guarantees privacy for honest users. For corrupted users the fraud detection ID might be chosen adversarially (cp. text body).Fig. 12. The functionality  $\mathcal{F}_{\text{P4TC}}$  (cont. from Fig. 4)

balance  $b^{\text{prev}}$  of the wallet. As new transaction entry

$$trdb := (s^{\text{prev}}, \perp, \phi, x, \lambda, pid_{\mathcal{U}}, pid_{\mathcal{T}}, -b^{\text{bill}}, 0)$$

is added to  $TRDB$  and the bill  $b^{\text{bill}}$  output to both parties. Furthermore, the TSP gets the user's ID  $pid_{\mathcal{U}}$ , as we assume Debt Clearance to be identifying, as well as the fraud detection ID  $\phi$  to enable double-spending detection.

### B.3 Feature Tasks

To obtain a more secure toll collection system we also provide the feature tasks Prove Participation, Double-Spending Detection, Guilt Verification and User Blacklisting. All of those tasks deal with different aspects arising from fraudulent user behavior.

**B.3.1 Prove Participation.** This is a two-party task involving a user and the SA (cp. Fig. 13) and assumed to be conducted with every user that has been physically caught by one of the SA's cameras. It allows an honest user to prove his successful participation in a transaction with the RSU where the photo was taken, while the fraudulent user will not be able to do so. The SA inputs the public key  $pk_{\mathcal{U}}$  of the user which the SA wishes to prove its participation and a set  $S_{\mathcal{R}}^{\text{PP}}$  of serial numbers in question. The user has no explicit input but simply expresses its consent by running the protocol.

First note, there is no guarantee that the user which participates in the protocol (with PID  $pid_{\mathcal{U}}$ ) is the user the SA wants to prove its participation (with user key  $pk_{\mathcal{U}}$  and  $pid'_{\mathcal{U}}$ ). Nonetheless, the ideal functionality checks if  $TRDB$  contains a transaction for the requested PID  $pid'_{\mathcal{U}}$  and a serial number  $s$  in  $S_{\mathcal{R}}^{\text{PP}}$ . Hence, the ideal functionality ensures that either the SA receives the correct answer—even for corrupted users—or aborts.

**Functionality  $\mathcal{F}_{\text{P4TC}}$  (cont.) – Task *Prove Participation***

*User input:* (prove\_participation)

*SA input:* (prove\_participation,  $\text{pk}_{\mathcal{U}}, S_{\mathcal{R}}^{\text{PP}}$ )

- (1) Obtain  $\text{pid}'_{\mathcal{U}}$  for  $\text{pk}_{\mathcal{U}}$  from  $\overline{\mathcal{G}}_{\text{bb}}$ .<sup>a</sup>
- (2) If  $(\text{pid}'_{\mathcal{U}} \notin \text{PID}_{\text{corrupt}} \vee \text{pid}_{\mathcal{U}} \notin \text{PID}_{\text{corrupt}}) \wedge \text{pid}'_{\mathcal{U}} \neq \text{pid}_{\mathcal{U}}$ , abort.
- (3) If  $\text{pid}_{\mathcal{U}} \in \text{PID}_{\text{corrupt}}$  leak  $S_{\mathcal{R}}^{\text{PP}}$  to the adversary.
- (4) If  $\exists (\cdot, s, \cdot, \cdot, \cdot, \text{pid}'_{\mathcal{U}}, \cdot, \cdot, \cdot) \in \text{TRDB}$  such that  $s \in S_{\mathcal{R}}^{\text{PP}}$ ,  
then  $\text{out}_{\mathcal{U}} := \text{out}_{\text{SA}} := \text{OK}$   
else  $\text{out}_{\mathcal{U}} := \text{out}_{\text{SA}} := \text{NOK}$ .

*User output:* ( $\text{out}_{\mathcal{U}}$ )

*SA output:* ( $\text{out}_{\text{SA}}$ )

<sup>a</sup> $\text{pid}_{\mathcal{U}}$  is the implicit PID of the user participating in the protocol.  $\text{pid}_{\mathcal{U}}$  is not necessarily equal to  $\text{pid}'_{\mathcal{U}}$  which denotes the user that is expected by the SA.

Fig. 13. The functionality  $\mathcal{F}_{\text{P4TC}}$  (cont. from Fig. 4)

Some remarks are in order to the abort condition—which is only a technical concession to what can be practically realized. If  $\text{pid}'_{\mathcal{U}} = \text{pid}_{\mathcal{U}}$  holds (i.e., the SA communicates with the expected user), everything is fine. If  $\text{pid}'_{\mathcal{U}} \neq \text{pid}_{\mathcal{U}}$  holds (i.e., the SA communicates with the wrong user) and at least one of the users is honest, the ideal functionality aborts. This models the fact that a malicious user must neither be able to embody an honest user nor an honest user embodies a malicious user. But if  $\text{pid}'_{\mathcal{U}} \neq \text{pid}_{\mathcal{U}}$  holds (i.e., the SA communicates with the wrong user) and both users are corrupted, the ideal functionality proceeds normally, but guarantees to return the correct result for the user in question (i.e., with  $\text{pid}'_{\mathcal{U}}$ ). This models the limitation that two corrupted users can share their credentials. A corrupted user (with  $\text{pid}'_{\mathcal{U}}$ ) can pass its transaction information to another corrupted user (with  $\text{pid}_{\mathcal{U}}$ ) which then proves participation to the SA. However, the latter user can still only prove the participation of the original user and not misuse the information to falsely prove participation for itself.

If the user is corrupted, the set of serial numbers in question is leaked to the adversary. Note further that this task also deanonymizes the one transaction proven by the user and leaks the respective serial number. Although the SA only obtains a single bit of information, whether the user's serial number is a member of the set  $S_{\mathcal{R}}^{\text{PP}}$  or not, this single bit of information is sufficient to restore the complete serial number by means of a bi-sectional search. The SA could repeatedly run the task and summon the user to prove its participation for a descending sequence of bi-sected sets until the last set only contains a single serial number. Nonetheless, this does not effect the anonymity or unlinkability of any other transactions.

**B.3.2 Double-Spending Detection and Guilt Verification.** Due to our requirement to allow offline RSUs, a user is able to fraudulently collect debt on outdated states of his wallet. This double-spending can not be prevented but must be detected afterwards. To ensure this,  $\mathcal{F}_{\text{P4TC}}$  provides the tasks Double-Spending Detection (cp. Fig. 14) and Guilt Verification (cp. Fig. 15).

Double-Spending Detection is a one-party task performed by the TSP. It takes a fraud detection ID  $\phi$  as input and checks the transaction database  $\text{TRDB}$  for two distinct entries containing this same fraud detection ID. In case such entries are present the adversary is asked for a proof  $\pi$  to be issued for this instance of double-spending. The user ID and proof  $(\text{pid}_{\mathcal{U}}, \pi)$  are appended to  $f_{\Pi}$  and marked as valid. Additionally, both are output to the TSP.

Guilt Verification is a one-party task as well but can be performed by any party. It takes a user ID  $\text{pid}_{\mathcal{U}}$  and a double-spending proof  $\pi$  as input. First, it checks if this particular pair  $(\text{pid}_{\mathcal{U}}, \pi)$  has already been defined and



**Functionality  $\mathcal{F}_{\text{P4TC}}$  (cont.) – Task *Double-Spending Detection****TSP input:* (scan\_for\_fraud,  $\phi$ )

- (1) Pick  $trdb \neq trdb'$  in  $TRDB$  such that  $trdb = (\cdot, \cdot, \phi, \cdot, \cdot, pid_{\mathcal{U}}, \cdot, \cdot, \cdot)$  and  $trdb' = (\cdot, \cdot, \phi, \cdot, \cdot, pid_{\mathcal{U}}, \cdot, \cdot, \cdot)$ .<sup>‡</sup>
- (2) Ask the adversary for a proof  $\pi \in \Pi$  corresponding to  $pid_{\mathcal{U}}$  and append  $(pid_{\mathcal{U}}, \pi) \mapsto \text{OK}$  to  $f_{\Pi}$ .
- (3) Receive  $pk_{\mathcal{U}}$  from the bulletin-board  $\overline{\mathcal{G}}_{\text{bb}}$  for  $pid_{\mathcal{U}}$ .

*TSP output:* ( $pk_{\mathcal{U}}, \pi$ )<sup>‡</sup>If this does not exist, output  $\perp$  and abort.Fig. 14. The functionality  $\mathcal{F}_{\text{P4TC}}$  (cont. from Fig. 4)**Functionality  $\mathcal{F}_{\text{P4TC}}$  (cont.) – Task *Guilt Verification****Party input:* (verify\_guilt,  $pk_{\mathcal{U}}, \pi$ )

- (1) Receive  $pid_{\mathcal{U}}$  from the bulletin-board  $\overline{\mathcal{G}}_{\text{bb}}$  for key  $pk_{\mathcal{U}}$ .<sup>‡</sup>
- (2) If  $f_{\Pi}(pid_{\mathcal{U}}, \pi)$  is defined, then set  $\text{out} := f_{\Pi}(pid_{\mathcal{U}}, \pi)$  and output (out).
- (3) If  $pid_{\mathcal{U}} \in \mathcal{PID}_{\text{corrupt}}$ , then leak  $(pid_{\mathcal{U}}, \pi)$  to the adversary and obtain result out, else set  $\text{out} := \text{NOK}$ .
- (4) Append  $(pid_{\mathcal{U}}, \pi) \mapsto \text{out}$  to  $f_{\Pi}$ .

*Party output:* (out)<sup>‡</sup>If this does not exist, output  $\perp$  and abort.Fig. 15. The functionality  $\mathcal{F}_{\text{P4TC}}$  (cont. from Fig. 4)

outputs whatever has been output before. This is necessary to ensure consistency across different invocations. If  $(pid_{\mathcal{U}}, \pi)$  has neither been issued nor queried before *and* the affected user is corrupted, the adversary is allowed to decide if this proof should be accepted. This implies that we do not protect corrupted users from false accusations of guilt. If the user is honest and  $(pid_{\mathcal{U}}, \pi)$  has neither been issued nor queried before, then the proof is marked as invalid. This protects honest users from being accused by made-up proofs which have not been issued by the ideal functionality itself. Finally, the result is recorded for the future and output to the party. This possibility of public verification is vital to prevent the TSP from wrongly accusing any user of double-spending and should for instance be utilized by the DR before it agrees to blacklist and therefore deanonymize a user on the basis of double-spending.

**B.3.3 User Blacklisting.** User Blacklisting (cp. Fig. 16) is a two-party task between the DR and TSP and serves two purposes: First, the debt  $b^{\text{bill}}$  owed by the user that is to be blacklisted is calculated. Second, fraud detection IDs for all of the user's wallets are determined and handed to the TSP so it may add them to the RSU blacklist  $bl_{\mathcal{R}}$ . Note that the generation of the blacklist  $bl_{\mathcal{T}}$  of user public keys is handled internally by the TSP and not in the scope of this task or  $\mathcal{F}_{\text{P4TC}}$ .

Both parties—the TSP and the DR—input the public key ( $pk_{\mathcal{U}}^{\text{DR}}$  and  $pk_{\mathcal{U}}^{\mathcal{T}}$ , resp.) of the user that is going to be blacklisted. We assume both parties to agree on the same key out-of-band before the protocol starts. We first describe the “normal” case (cp. Fig. 16, Step 3a) for an honest TSP. (N.b.: The DR is assumed to be always honest.) To calculate the user's outstanding debt, all transaction entries in  $TRDB$  containing  $pid_{\mathcal{U}}$  are taken and their respective prices  $p$  summed up to obtain  $b^{\text{bill}}$ . Note that although this sum may contain the prices of transactions

**Functionality  $\mathcal{F}_{P4TC}$  (cont.) – Task User Blacklisting**

*DR input:* (blacklist\_user,  $pk_{\mathcal{U}}^{DR}$ )

*TSP input:* (blacklist\_user,  $pk_{\mathcal{U}}^T$ )

- (1) If  $pk_{\mathcal{U}}^{DR} \neq pk_{\mathcal{U}}^T$ , abort.
- (2) Receive  $pid_{\mathcal{U}}$  from the bulletin-board  $\overline{\mathcal{G}}_{bb}$  for  $pk_{\mathcal{U}}^T$ .<sup>‡</sup>
- (3) Distinguish 3 cases:
  - (a)  $pid_{\mathcal{T}} \notin \mathcal{PID}_{corrupt}$ : Set  $\mathcal{L}_{bl} := \{\lambda \mid (\cdot, \cdot, \cdot, \cdot, \lambda, pid_{\mathcal{U}}, \cdot, \cdot, \cdot) \in TRDB\}$ .
  - (b)  $pid_{\mathcal{T}} \in \mathcal{PID}_{corrupt}$ ,  $pid_{\mathcal{U}} \notin \mathcal{PID}_{corrupt}$ : Obtain a set of serial numbers  $S_{root}$  from the adversary and set  $\mathcal{L}_{bl} := \{\lambda \mid (\perp, s, \cdot, \cdot, \lambda, pid_{\mathcal{U}}, \cdot, \cdot, \cdot) \in TRDB \text{ with } s \in S_{root}\}$ .
  - (c)  $pid_{\mathcal{T}} \in \mathcal{PID}_{corrupt}$ ,  $pid_{\mathcal{U}} \in \mathcal{PID}_{corrupt}$ : Let the adversary decide on the output for DR and TSP and stop.
- (4)  $TRDB_{bl} := \{trdb \in TRDB \mid trdb = (\cdot, \cdot, \cdot, \cdot, \lambda, \cdot, \cdot, p, \cdot) \text{ s. t. } \lambda \in \mathcal{L}_{bl}\}$
- (5)  $b^{bill} := \sum_{trdb \in TRDB_{bl}} p$ .
- (6) For each  $\lambda \in \mathcal{L}_{bl}$ :
  - (a)  $x_{\lambda} := \max\{x \mid f_{\Phi}(\lambda, x) \text{ is already defined}\}$ .
  - (b) For  $x \in \{x_{\lambda} + 1, \dots, x_{bl_{\mathcal{R}}}\}$ :
    - (i) If  $pid_{\mathcal{U}} \notin \mathcal{PID}_{corrupt}$ , pick  $\phi \xleftarrow{R} \Phi$  that has not previously been used, otherwise leak  $(\lambda, x)$  to the adversary and obtain fraud detection ID  $\phi$  that has not previously been used.
    - (ii) Append  $(\lambda, x) \mapsto \phi$  to  $f_{\Phi}$ .
- (7)  $\Phi_{bl} := \{f_{\Phi}(\lambda, x) \mid \lambda \in \mathcal{L}_{bl}, 0 \leq x \leq x_{bl_{\mathcal{R}}}\}$ .

*DR output:* (OK)

*TSP output:* ( $b^{bill}$ ,  $\Phi_{bl}$ )

<sup>‡</sup>If this does not exist, output  $\perp$  and abort.

Fig. 16. The functionality  $\mathcal{F}_{P4TC}$  (cont. from Fig. 4)

and wallets that have already been cleared, this does not falsify the value of  $b^{bill}$  as every successful execution of Debt Clearance creates an entry with the amount that was cleared as negative price. For the actual blacklisting the set of all wallet IDs belonging to  $pid_{\mathcal{U}}$  is looked up and the remainder of the task is conducted for every wallet  $\lambda$  separately.  $\mathcal{F}_{P4TC}$  checks how many values of  $f_{\Phi}(\lambda, \cdot)$  are already defined and extends them to the first  $x_{bl_{\mathcal{R}}}$  fraud detection IDs, where  $x_{bl_{\mathcal{R}}}$  is a parameter we assume to be greater than the number of transactions a user would be involved in within one billing period. To that end, yet undefined fraud detection IDs  $f_{\Phi}(\lambda, x)$  with  $x \leq x_{bl_{\mathcal{R}}}$  are uniquely and randomly drawn or—in case of a corrupted user—obtained from the adversary. Finally, all fraud detection IDs  $\phi = f_{\Phi}(\lambda, x)$  for  $x \leq x_{bl_{\mathcal{R}}}$  and all wallets  $\lambda$  of the user are output to the TSP together with the outstanding debt  $b^{bill}$ .

It remains to describe the remaining two cases (cp. Fig. 16, Steps 3b and 3c). If the TSP is corrupted but the user in question honest (Step 3b), the TSP is free to drop some of the user’s wallets and only partially blacklist the user. This “attack”—a demented TSP—cannot be ruled out. In order to correctly map this in the ideal model, the adversary is asked to provide a set of associated serial numbers and only the wallets from the intersection are used for blacklisting. This ensures that a malicious TSP can only blacklist less wallets but not more wallets or even wallets of another user. If the TSP and the user in question are both corrupted (Step 3c), no guarantees are given. Please note that a corrupted TSP could even come up with an “imaginary” corrupted user (which only exists in the head of the TSP) and ask the DR to blacklist this user. Essentially, this is nothing else than a

cumbersome way to evaluate the PRF at inputs chosen by the TSP. However, the TSP can do this by itself anyway. We stress that this does not affect the security or privacy of honest parties in the system.

## C SETTING AND BUILDING BLOCKS

In this appendix we introduce the algebraic setting and building blocks we make use of. In particular, the latter includes non-interactive zero-knowledge proofs, commitments, signatures, encryption and pseudo-random functions. We also describe possible instantiations for these building blocks and explain how these primitives are used in our system.

### C.1 Algebraic Setting and Assumptions

Our protocol instantiations are based on an asymmetric bilinear group setting  $\text{gp} := (G_1, G_2, G_T, e, p, g_1, g_2)$ . We adopt the following definition from [43].

*Definition C.1 (Prime-order Bilinear Group Generator).* A *prime-order bilinear group generator* is a PPT algorithm  $\text{SetupGrp}$  that on input of a security parameter  $1^n$  outputs a tuple of the form

$$\text{gp} := (G_1, G_2, G_T, e, p, g_1, g_2) \leftarrow \text{SetupGrp}(1^n) \quad (5)$$

where  $G_1, G_2, G_T$  are descriptions of cyclic groups of prime order  $p$ ,  $\log p = \Theta(n)$ ,  $g_1$  is a generator of  $G_1$ ,  $g_2$  is a generator of  $G_2$ , and  $e: G_1 \times G_2 \rightarrow G_T$  is a map (aka pairing) which satisfies the following properties:

- *Efficiency:*  $e$  is efficiently computable.
- *Bilinearity:*  $\forall a \in G_1, b \in G_2, x, y \in \mathbb{Z}_p : e(a^x, b^y) = e(a, b)^{xy}$ .
- *Non-Degeneracy:*  $e(g_1, g_2)$  generates  $G_T$ .

The setting is called *asymmetric*, if no efficiently computable homomorphisms between  $G_1$  and  $G_2$  are known. In the remainder of this paper, we consider the asymmetric kind.

Our construction relies on the co-CDH assumption for identification, and the security of our building blocks (cp. [Appendix C.2](#)) in asymmetric bilinear groups. For our special instantiation of the building blocks (see there), security holds under the SXDH and co-DLIN assumption. The former implies the co-CDH assumption.

The SXDH assumption essentially asserts that the DDH assumption holds in both source groups  $G_1$  and  $G_2$  of the bilinear map and is formally defined as:

*Definition C.2 (DDH and SXDH Assumption).*

- (1) We say that the *DDH assumption* holds with respect to  $\text{SetupGrp}$  over  $G_i$  if the advantage  $\text{Adv}_{\text{SetupGrp}, i, \mathcal{A}}^{\text{DDH}}(1^n)$  defined by

$$\left| \Pr \left[ \begin{array}{l} b = b' \\ \text{gp} := (G_1, G_2, G_T, e, p, g_1, g_2) \leftarrow \text{SetupGrp}(1^n) \\ x, y, z \leftarrow \mathbb{Z}_p; h_0 := g_i^{xy}; h_1 := g_i^z \\ b \xleftarrow{R} \{0, 1\} \\ b' \leftarrow \mathcal{A}(1^n, \text{gp}, g_i^x, g_i^y, h_b) \end{array} \right] - \frac{1}{2} \right|$$

is a negligible function in  $n$  for all PPT algorithms  $\mathcal{A}$ .

- (2) We say that the *SXDH assumption* holds with respect to  $\text{SetupGrp}$  if the above holds for both  $i = 1$  and  $i = 2$ .

The co-CDH assumption is defined as follows:

*Definition C.3 (Co-CDH assumption).* We say that the *co-CDH assumption* holds with respect to  $\text{SetupGrp}$  if the advantage  $\text{Adv}_{\text{SetupGrp}, \mathcal{A}}^{\text{CO-CDH}}(1^n)$  defined by

$$\Pr \left[ a = g_2^x \mid \begin{array}{l} \text{gp} := (G_1, G_2, G_T, e, p, g_1, g_2) \leftarrow \text{SetupGrp}(1^n) \\ x \leftarrow \mathbb{Z}_p \\ a \leftarrow \mathcal{A}(1^n, \text{gp}, g_1^x) \end{array} \right]$$

is a negligible function in  $n$  for all PPT algorithms  $\mathcal{A}$ .

The co-DLIN assumption is defined as follows:

*Definition C.4.* We say that the *co-DLIN assumption* holds with respect to  $\text{SetupGrp}$  if the advantage  $\text{Adv}_{\text{SetupGrp}, \mathcal{A}}^{\text{CO-DLIN}}(1^n)$  defined by

$$\Pr \left[ b = b' \mid \begin{array}{l} \text{gp} := (G_1, G_2, G_T, e, p, g_1, g_2) \leftarrow \text{SetupGrp}(1^n) \\ \alpha, \beta, \gamma \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_p \\ b \stackrel{\mathbb{R}}{\leftarrow} \{0, 1\} \\ \check{h}_1 := g_1^\alpha, \check{h}_2 := g_1^\beta, \check{h}_3 := g_1^{\alpha+\beta+b\gamma} \\ \hat{h}_1 := g_2^\alpha, \hat{h}_2 := g_2^\beta, \hat{h}_3 := g_2^{\alpha+\beta+b\gamma} \\ b' \leftarrow \mathcal{A}(1^n, \text{gp}, \check{h}_1, \check{h}_2, \check{h}_3, \hat{h}_1, \hat{h}_2, \hat{h}_3) \end{array} \right]$$

is a negligible function in  $n$  for all PPT algorithms  $\mathcal{A}$ .

## C.2 Cryptographic Building Blocks

Our semi-generic construction makes use of various cryptographic primitives including ( $F_{\text{gp}}$ -extractable) NIZK proofs, equivocal and extractable homomorphic commitments, digital signatures, public-key encryption, symmetric encryption and pseudo-random functions. The latter building blocks need to be efficiently and securely combinable with the chosen NIZK proof system, which is Groth-Sahai (GS) in our case. In the following, we give an introduction to the formal definition of these building blocks.

*C.2.1 Group setup.* Let  $\text{SetupGrp}$  be a bilinear group generator (cp. [Definition C.1](#)) that outputs descriptions of asymmetric bilinear groups  $\text{gp} \leftarrow \text{SetupGrp}(1^n)$ . The following building blocks all make use of  $\text{SetupGrp}$  as their common group setup algorithm.

*C.2.2 NIZKs.* Let  $R$  be a witness relation for some NP language  $L = \{\text{stmt} \mid \exists \text{wit s.t. } (\text{stmt}, \text{wit}) \in R\}$ . A zero-knowledge proof system allows a prover  $P$  to convince a verifier  $V$  that some  $\text{stmt}$  is contained in  $L$  without  $V$  learning anything beyond that fact. In a non-interactive zero-knowledge (NIZK) proof, only one message, the proof  $\pi$ , is sent from  $P$  to  $V$  for that purpose.

More precisely, a (group-based) NIZK proof system is defined as:

*Definition C.5 (Group-based NIZK proof system).* Let  $R$  be an efficiently verifiable relation containing triples  $(\text{gp}, x, w)$ . We call  $\text{gp}$  the group setup,  $x$  the statement, and  $w$  the witness. Given some  $\text{gp}$ , let  $L_{\text{gp}}$  be the language containing all statements  $x$  such that  $(\text{gp}, x, w) \in R$ . Let  $\text{POK} := (\text{SetupGrp}, \text{SetupPoK}, \text{Prove}, \text{Vfy})$  be a tuple of PPT algorithms such that

- $\text{SetupGrp}$  takes as input a security parameter  $1^n$  and outputs public parameters  $\text{gp}$ . We assume that  $\text{gp}$  is given as implicit input to all algorithms.
- $\text{SetupPoK}$  takes as input  $\text{gp}$  and outputs a (public) common reference string  $\text{CRS}_{\text{pok}}$ .
- $\text{Prove}$  takes as input the common reference string  $\text{CRS}_{\text{pok}}$ , a statement  $x$ , and a witness  $w$  with  $(\text{gp}, x, w) \in R$  and outputs a proof  $\pi$ .
- $\text{Vfy}$  takes as input the common reference string  $\text{CRS}_{\text{pok}}$ , a statement  $x$ , and a proof  $\pi$  and outputs 1 or 0.

POK is called a non-interactive zero-knowledge proof system for  $R$  with  $F_{\text{gp}}$ -extractability, if the following properties are satisfied:

- (1) *Perfect completeness*: For all  $\text{gp} \leftarrow \text{SetupGrp}(1^n)$ ,  $\text{CRS}_{\text{pok}} \leftarrow \text{SetupPoK}(\text{gp})$ ,  $(\text{gp}, x, w) \in R$ , and  $\pi \leftarrow \text{Prove}(\text{CRS}_{\text{pok}}, x, w)$  we have that  $\text{Vfy}(\text{CRS}_{\text{pok}}, x, \pi) = 1$ .
- (2) *Perfect soundness*: For all (possibly unbounded) adversaries  $\mathcal{A}$  we have that

$$\Pr \left[ \text{Vfy}(\text{CRS}_{\text{pok}}, x, \pi) = 0 \mid \begin{array}{l} \text{gp} \leftarrow \text{SetupGrp}(1^n) \\ \text{CRS}_{\text{pok}} \leftarrow \text{SetupPoK}(\text{gp}) \\ (x, \pi) \leftarrow \mathcal{A}(\text{CRS}_{\text{pok}}) \\ x \notin L_{\text{gp}} \end{array} \right]$$

is 1.

- (3) *Perfect  $F_{\text{gp}}$ -extractability*: There exists a polynomial-time extractor ( $\text{SetupEPoK}, \text{ExtractW}$ ) such that for all (possibly unbounded) adversaries  $\mathcal{A}$

- (a) we have that the advantage  $\text{Adv}_{\text{POK}, \mathcal{A}}^{\text{pok-ext-setup}}(n)$  defined by

$$\left| \Pr \left[ 1 \leftarrow \mathcal{A}(\text{CRS}_{\text{pok}}) \mid \begin{array}{l} \text{gp} \leftarrow \text{SetupGrp}(1^n), \\ \text{CRS}_{\text{pok}} \leftarrow \text{SetupPoK}(\text{gp}) \end{array} \right] - \Pr \left[ 1 \leftarrow \mathcal{A}(\text{CRS}'_{\text{pok}}) \mid \begin{array}{l} \text{gp} \leftarrow \text{SetupGrp}(1^n), \\ (\text{CRS}'_{\text{pok}}, \text{td}_{\text{epok}}) \leftarrow \text{SetupEPoK}(\text{gp}) \end{array} \right] \right|$$

is zero.

- (b) we have that the advantage  $\text{Adv}_{\text{POK}, \mathcal{A}}^{\text{pok-ext}}(n)$  defined by

$$\Pr \left[ \exists w : F_{\text{gp}}(w) = W \wedge (\text{gp}, x, w) \in R \mid \begin{array}{l} \text{gp} \leftarrow \text{SetupGrp}(1^n) \\ (\text{CRS}'_{\text{pok}}, \text{td}_{\text{epok}}) \leftarrow \text{SetupEPoK}(\text{gp}) \\ (x, \pi) \leftarrow \mathcal{A}(\text{CRS}'_{\text{pok}}) \\ 1 \leftarrow \text{Vfy}(\text{CRS}'_{\text{pok}}, x, \pi) \\ W \leftarrow \text{ExtractW}(\text{CRS}'_{\text{pok}}, \text{td}_{\text{epok}}, x, \pi) \end{array} \right]$$

is 1.

- (4) *Composable Zero-knowledge*: There exists a polynomial-time simulator ( $\text{SetupSPoK}, \text{SimProof}$ ) and hint generator  $\text{GenHint}$  such that for all PPT adversaries  $\mathcal{A}$

- (a) we have that the advantage  $\text{Adv}_{\text{POK}, \mathcal{A}}^{\text{pok-zk-setup}}(n)$  defined by

$$\left| \Pr \left[ 1 \leftarrow \mathcal{A}(\text{CRS}_{\text{pok}}) \mid \begin{array}{l} \text{gp} \leftarrow \text{SetupGrp}(1^n), \\ \text{CRS}_{\text{pok}} \leftarrow \text{SetupPoK}(\text{gp}) \end{array} \right] - \Pr \left[ 1 \leftarrow \mathcal{A}(\text{CRS}'_{\text{pok}}) \mid \begin{array}{l} \text{gp} \leftarrow \text{SetupGrp}(1^n), \\ \text{hint} \leftarrow \text{GenHint}(\text{gp}), \\ (\text{CRS}'_{\text{pok}}, \text{td}_{\text{spok}}) \leftarrow \text{SetupSPoK}(\text{gp}, \text{hint}), \end{array} \right] \right|$$

is negligible in  $n$ .

- (b) we have that the advantage  $\text{Adv}_{\text{POK}, \mathcal{A}}^{\text{pok-zk}}(n)$  defined by

$$\left| \Pr \left[ 1 \leftarrow \mathcal{A}^{\text{SimProof}'(\text{CRS}'_{\text{pok}}, \text{td}_{\text{spok}}, \cdot, \cdot)}(1^n, \text{CRS}'_{\text{pok}}, \text{td}_{\text{spok}}) \right] - \Pr \left[ 1 \leftarrow \mathcal{A}^{\text{Prove}(\text{CRS}'_{\text{pok}}, \cdot, \cdot)}(1^n, \text{CRS}'_{\text{pok}}, \text{td}_{\text{spok}}) \right] \right|$$

is negligible in  $n$ , where  $\text{gp} \leftarrow \text{SetupGrp}(1^n)$ ,  $(\text{CRS}'_{\text{pok}}, \text{td}_{\text{spok}}) \leftarrow \text{SetupSPoK}(\text{gp})$ , and  $\text{SimProof}'(\text{CRS}'_{\text{pok}}, \text{td}_{\text{spok}}, \cdot, \cdot)$  is an oracle which on input  $(x, z) \in R$  returns  $\text{SimProof}(\text{CRS}'_{\text{pok}}, \text{td}_{\text{spok}}, x)$ . Both  $\text{SimProof}'$  and  $\text{Prove}$  return  $\perp$  on input  $(x, z) \notin R$ .

We wish to point out some remarks.

- Remark C.6.* (1) The considered language  $L_{\text{gp}}$  may depend on  $\text{gp}$ .  
 (2)  $F_{\text{gp}}$ -extractability actually implies soundness independent of  $F_{\text{gp}}$ : If there was a false statement  $x$  which verifies, violating soundness, then obviously there is no witness  $w$  for  $x$ , which violates extractability.  
 (3) Extractability essentially means that  $\text{ExtractW}$ —given a trapdoor  $\text{td}_{\text{epok}}$ —is able to extract  $F_{\text{gp}}(\text{wit})$  for an NP-witness  $\text{wit}$  for  $\text{stmnt} \in L_{\text{gp}}$  from any valid proof  $\pi$ . If  $F_{\text{gp}}$  is the identity function, then the actual witness is extracted and the system is called a *proof of knowledge*.

*Our Instantiation.* We choose the SXDH-based Groth-Sahai proof system [37, 40] as our NIZK, as it allows for very efficient proofs (under standard assumptions). On the other hand, GS comes with some drawbacks, which makes applying it sometimes pretty tricky: It only works for algebraic languages containing certain types of equations, it is not always zero-knowledge, and  $F_{\text{gp}}$  is not always the identity function. For the sake of completeness [Appendix C.3](#) contains a description what types of equations are supported by GS. When choosing our remaining building blocks and forming equations we ensured that they fit into this framework. Likewise, we ensured that the ZK-property holds for the languages we consider.

For proving correctness of the computations taking place on the user's side we need three different instantiations of the GS proof system, denoted by P1, P2 and P3, respectively. The corresponding functions  $F_{\text{gp}}^{(1)}$ ,  $F_{\text{gp}}^{(2)}$  and  $F_{\text{gp}}^{(3)}$  depend on the considered languages  $L_1$ ,  $L_2$  and  $L_3$  (defined in [Appendices D.6 to D.8](#)) but they have the following in common: They behave as the identity function with respect to group elements and map elements from  $\mathbb{Z}_p$  either to  $G_1$  or  $G_2$  (by exponentiation with basis  $g_1$  or  $g_2$ ) depending on whether these are used as exponents of a  $G_1$  or  $G_2$  element in the language.

All proof systems will share a common reference string. More precisely, we demand that there is a shared extraction setup algorithm which generates the CRS and also a single extraction trapdoor for P1, P2 and P3. Let us denote this algorithm by  $\text{SetupEPoK}$  and its output by  $(\text{CRS}_{\text{pok}}, \text{td}_{\text{epok}}) \leftarrow \text{SetupEPoK}(\text{gp})$  in the following. Furthermore, let us denote the prove and verify algorithms of these proof systems by  $\text{PX.Prove}$  and  $\text{PX.Vfy}$ , for  $1 \leq X \leq 3$ .

*Range Proofs.* For one particular task<sup>22</sup> we need range proofs in order to show that some  $\mathbb{Z}_p$ -element  $\lambda'_i$  is “smaller” than some fixed system parameter  $\mathcal{B}$  with both elements being regarded as elements from  $\{0, \dots, p-1\}$  and the normal  $\leq$ -relation from the integers. We realize these range proofs using Groth-Sahai by applying the signature-based technique in [15]. Here, the verifier initially chooses parameters  $q$  and  $t$  such that every possible  $\lambda'_i$  can be represented as  $\lambda'_i = \sum_{j=0}^t d_j q^j$  with  $0 \leq d_j < q$ . He also generates a signature on every possible value of a digit, i.e.,  $0, \dots, q-1$ . The prover then shows using a Groth-Sahai NIZK that each  $\lambda'_i$  can be indeed represented in this way and that he knows a signature by the verifier for each of its digits. Clearly, a structure-preserving signature scheme is needed for this purpose and we use the one in [1].

**C.2.3 Commitments.** A commitment scheme allows a user to commit to a message  $m$  and publish the result, called commitment  $c$ , in a way that  $m$  is hidden from others, but also the user cannot claim a different  $m$  afterwards when he opens  $c$ . A commitment scheme is called an  $F_{\text{gp}}$ -binding commitment scheme for a bijective function  $F_{\text{gp}}$  on the message space, if one commits to a message  $m$  but opens the commitment using  $F_{\text{gp}}(m)$ . We call the codomain of  $F_{\text{gp}}$  the implicit message space.

<sup>22</sup>More precisely: The task Wallet Issuing

*Definition C.7.* A commitment scheme  $\text{COM} := (\text{SetupGrp}, \text{Gen}, \text{Com}, \text{Open})$  consists of four algorithms:

- $\text{SetupGrp}$  takes as input a security parameter  $1^n$  and outputs public parameters  $\text{gp}$ . These parameters also define a message space  $\mathcal{M}$ , an implicit message space  $\mathcal{M}'$  and a function  $F_{\text{gp}} : \mathcal{M} \rightarrow \mathcal{M}'$  mapping a message to its implicit representation. We assume that  $\text{gp}$  is given as implicit input to all algorithms.
- $\text{Gen}$  is a PPT algorithm, which takes  $\text{gp}$  as input and outputs public parameters  $\text{CRS}_{\text{com}}$ .
- $\text{Com}$  is a PPT algorithm, which takes as input parameters  $\text{CRS}_{\text{com}}$  and a message  $m \in \mathcal{M}$  and outputs a commitment  $c$  to  $m$  and some decommitment value  $d$ .
- $\text{Open}$  is a deterministic polynomial-time algorithm, which takes as input parameters  $\text{CRS}_{\text{com}}$ , commitment  $c$ , an implicit message  $M \in \mathcal{M}'$ , and opening  $d$ . It returns either 0 or 1.

$\text{COM}$  is *correct* if for all  $\text{gp} \leftarrow \text{SetupGrp}(1^n)$ ,  $\text{CRS}_{\text{com}} \leftarrow \text{Gen}(\text{gp})$ ,  $m \in \mathcal{M}$ , and  $(c, d) \leftarrow \text{Com}(\text{CRS}_{\text{com}}, m)$  it holds that  $1 = \text{Open}(\text{CRS}_{\text{com}}, F_{\text{gp}}(m), c, d)$ .

We say that  $\text{COM}$  is a (computationally) *hiding*,  $F_{\text{gp}}$ -*binding*, *equivocal*, *extractable* commitment scheme if it has the following properties:

- (1) *Hiding*: For all PPT adversaries  $\mathcal{A}$  it holds that the advantage  $\text{Adv}_{\text{COM}, \mathcal{A}}^{\text{Hiding}}(1^n)$  defined by

$$\Pr \left[ b = b' \mid \begin{array}{l} \text{gp} \leftarrow \text{SetupGrp}(1^n) \\ \text{CRS}_{\text{com}} \leftarrow \text{Gen}(\text{gp}) \\ (m_0, m_1, \text{state}) \leftarrow \mathcal{A}(1^n, \text{CRS}_{\text{com}}) \\ b \xleftarrow{\mathbb{R}} \{0, 1\} \\ (c, d) \leftarrow \text{Com}(\text{CRS}_{\text{com}}, m_b) \\ b' \leftarrow \mathcal{A}(c, \text{state}) \end{array} \right] = \frac{1}{2}$$

is negligible in  $n$ . The scheme is called *statistically hiding* if  $\text{Adv}_{\text{COM}, \mathcal{A}}^{\text{Hiding}}(1^n)$  is negligible even for an unbounded adversary  $\mathcal{A}$ .

- (2)  $F_{\text{gp}}$ -*Binding*: For all PPT adversaries  $\mathcal{A}$  it holds that the advantage  $\text{Adv}_{\mathcal{A}}^{F_{\text{gp}}\text{-Binding}}(1^n)$  defined by

$$\Pr \left[ \begin{array}{l} \text{Open}(\text{CRS}_{\text{com}}, M, c, d) = 1 \\ \wedge \\ \text{Open}(\text{CRS}_{\text{com}}, M', c, d') = 1 \end{array} \mid \begin{array}{l} \text{gp} \leftarrow \text{SetupGrp}(1^n) \\ \text{CRS}_{\text{com}} \leftarrow \text{Gen}(\text{gp}) \\ (c, M, d, M', d') \leftarrow \mathcal{A}(1^n, \text{CRS}_{\text{com}}) \\ M \neq M' \end{array} \right]$$

is negligible in  $n$ .

- (3) *Equivocal*: There exist PPT algorithms  $\text{SimGen}$ ,  $\text{SimCom}$  and  $\text{Equiv}$  such that for all PPT adversaries  $\mathcal{A}$

- (a) we have that the advantage  $\text{Adv}_{\text{COM}, \mathcal{A}}^{\text{SimGen}}(n)$  defined by

$$\left| \Pr \left[ 1 \leftarrow \mathcal{A}(\text{CRS}_{\text{com}}) \mid \begin{array}{l} \text{gp} \leftarrow \text{SetupGrp}(1^n), \\ \text{CRS}_{\text{com}} \leftarrow \text{Gen}(\text{gp}) \end{array} \right] - \Pr \left[ 1 \leftarrow \mathcal{A}(\text{CRS}'_{\text{com}}) \mid \begin{array}{l} \text{gp} \leftarrow \text{SetupGrp}(1^n), \\ (\text{CRS}'_{\text{com}}, \text{td}_{\text{eqcom}}) \leftarrow \text{SimGen}(\text{gp}) \end{array} \right] \right|$$

is negligible in  $n$ .

(b) we have that the advantage  $\text{Adv}_{\text{COM}, \mathcal{A}}^{\text{Equiv}}(n)$  defined by

$$\left| \begin{array}{l} \Pr \left[ 1 \leftarrow \mathcal{A}(\text{CRS}'_{\text{com}}, \text{td}_{\text{eqcom}}, m, c, d) \right] \\ -\Pr \left[ 1 \leftarrow \mathcal{A}(\text{CRS}'_{\text{com}}, \text{td}_{\text{eqcom}}, m, c', d') \right] \end{array} \middle| \begin{array}{l} \text{gp} \leftarrow \text{SetupGrp}(1^n), \\ (\text{CRS}'_{\text{com}}, \text{td}_{\text{eqcom}}) \leftarrow \text{SimGen}(\text{gp}), \\ m \leftarrow \mathcal{M}, \\ (c, d) \leftarrow \text{Com}(\text{CRS}'_{\text{com}}, m) \\ \text{gp} \leftarrow \text{SetupGrp}(1^n), \\ (\text{CRS}'_{\text{com}}, \text{td}_{\text{eqcom}}) \leftarrow \text{SimGen}(\text{gp}), \\ (c', r) \leftarrow \text{SimCom}(\text{gp}), \\ m \leftarrow \mathcal{M}, \\ d' \leftarrow \text{Equiv}(\text{CRS}'_{\text{com}}, \text{td}_{\text{eqcom}}, m, r) \end{array} \right|$$

is zero.

(4) *Extractable*: There exist PPT algorithms  $\text{ExtGen}$  and  $\text{Extract}$  such that for all PPT adversaries  $\mathcal{A}$

(a) we have that the advantage  $\text{Adv}_{\text{COM}, \mathcal{A}}^{\text{ExtGen}}(n)$  defined by

$$\left| \begin{array}{l} \Pr \left[ 1 \leftarrow \mathcal{A}(\text{CRS}_{\text{com}}) \right] \\ -\Pr \left[ 1 \leftarrow \mathcal{A}(\text{CRS}'_{\text{com}}) \right] \end{array} \middle| \begin{array}{l} \text{gp} \leftarrow \text{SetupGrp}(1^n), \\ \text{CRS}_{\text{com}} \leftarrow \text{Gen}(\text{gp}) \\ \text{gp} \leftarrow \text{SetupGrp}(1^n), \\ (\text{CRS}'_{\text{com}}, \text{td}_{\text{extcom}}) \leftarrow \text{ExtGen}(\text{gp}) \end{array} \right|$$

is negligible in  $n$ .

(b) we have that the advantage  $\text{Adv}_{\text{COM}, \mathcal{A}}^{\text{Ext}}(n)$  defined by

$$\Pr \left[ \text{Extract}(\text{CRS}'_{\text{com}}, \text{td}_{\text{extcom}}, c) \neq F_{\text{gp}}(m) \right] \middle| \begin{array}{l} \text{gp} \leftarrow \text{SetupGrp}(1^n), \\ (\text{CRS}'_{\text{com}}, \text{td}_{\text{extcom}}) \leftarrow \text{ExtGen}(\text{gp}), \\ c \leftarrow \mathcal{A}(\text{CRS}'_{\text{com}}), \\ \exists! m \in \mathcal{M}, r : c \leftarrow \text{Com}(\text{CRS}'_{\text{com}}, m; r) \end{array}$$

is zero.

Furthermore, assume that the message space of COM is an additive group. Then COM is called *additively homomorphic*, if there exist additional PPT algorithms  $c \leftarrow \text{CAdd}(\text{CRS}_{\text{com}}, c_1, c_2)$  and  $d \leftarrow \text{DAdd}(\text{CRS}_{\text{com}}, d_1, d_2)$  which on input of two commitments and corresponding decommitment values  $(c_1, d_1) \leftarrow \text{Com}(\text{CRS}_{\text{com}}, m_1)$  and  $(c_2, d_2) \leftarrow \text{Com}(\text{CRS}_{\text{com}}, m_2)$ , output a commitment  $c$  and decommitment  $d$ , respectively, such that  $\text{Open}(\text{CRS}_{\text{com}}, c, F_{\text{gp}}(m_1 + m_2), d) = 1$ .

Finally, we call COM *opening complete* if for all  $M \in \mathcal{M}'$  and arbitrary values  $c, d$  with  $\text{Open}(\text{CRS}_{\text{com}}, M, c, d) = 1$  holds that there exists  $m \in \mathcal{M}$  and randomness  $r$  such that  $(c, d) \leftarrow \text{Com}(\text{CRS}_{\text{com}}, m; r)$ .

*Our Instantiation.* We will make use of two commitment schemes that are both based on the SXDH assumption. We first use the shrinking  $\alpha$ -message-commitment scheme from Abe et al. [3]. This commitment scheme has message space  $\mathbb{Z}_p^\alpha$ , commitment space  $G_2$  and opening value space  $G_1$ . It is statistically hiding, additively homomorphic, equivocal, and  $F'_{\text{gp}}$ -Binding, for  $F'_{\text{gp}}(m_1, \dots, m_\alpha) := (g_1^{m_1}, \dots, g_1^{m_\alpha})$ . We use this commitment scheme as C1 with CRS  $\text{CRS}_{\text{com}}^1$  in the following ways in our system:

- In the Wallet Issuing task we use C1 for messages from  $\mathbb{Z}_p$  ( $\alpha := 1$ ),  $\mathbb{Z}_p^2$  ( $\alpha := 2$ ) and  $\mathbb{Z}_p^4$  ( $\alpha := 4$ ).
- In the Debt Accumulation task we use C1 for messages from  $\mathbb{Z}_p$  ( $\alpha := 1$ ) and  $\mathbb{Z}_p^4$  ( $\alpha := 4$ ).

We also use the (dual-mode) equivocal and extractable commitment scheme from Groth and Sahai [40]. This commitment scheme has message space  $G_1$ , commitment space  $G_1^2$  and opening value space  $\mathbb{Z}_p^2$ . It is equivocal,



extractable, hiding and  $F'_{\text{gp}}$ -Binding for  $F'_{\text{gp}}(m) := m$ . In our system, we use this commitment scheme as C2 with  $\text{CRS } \text{CRS}_{\text{com}}^2$  in the Wallet Issuing and Debt Accumulation tasks.

*C.2.4 Digital signatures.* A signature allows a signer to issue a signature  $\sigma$  on a message  $m$  using its secret signing key  $\text{sk}$  such that anybody can publicly verify that  $\sigma$  is a valid signature for  $m$  using the public verification key  $\text{pk}$  of the signer but nobody can feasibly forge a signature without knowing  $\text{sk}$ .

*Definition C.8.* A digital signature scheme  $S := (\text{SetupGrp}, \text{Gen}, \text{Sgn}, \text{Vfy})$  consists of four PPT algorithms:

- $\text{SetupGrp}$  takes as input a security parameter  $1^n$  and outputs public parameters  $\text{gp}$ . We assume that  $\text{gp}$  is given as implicit input to all algorithms.
- $\text{Gen}$  takes  $\text{gp}$  as input and outputs a key pair  $(\text{pk}, \text{sk})$ . The public key and  $\text{gp}$  define a message space  $\mathcal{M}$ .
- $\text{Sgn}$  takes as input the secret key  $\text{sk}$  and a message  $m \in \mathcal{M}$ , and outputs a signature  $\sigma$ .
- $\text{Vfy}$  takes as input the public key  $\text{pk}$ , a message  $m \in \mathcal{M}$ , and a purported signature  $\sigma$ , and outputs a bit.

We call  $S$  correct if for all  $n \in \mathbb{N}$ ,  $\text{gp} \leftarrow \text{SetupGrp}(1^n)$ ,  $m \in \mathcal{M}$ ,  $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(\text{gp})$ ,  $\sigma \leftarrow \text{Sgn}(\text{sk}, m)$  we have  $1 \leftarrow \text{Vfy}(\text{pk}, \sigma, m)$ .

We say that  $S$  is EUF-CMA secure if for all PPT adversaries  $\mathcal{A}$  it holds that the advantage  $\text{Adv}_{S, \mathcal{A}}^{\text{EUF-CMA}}(1^n)$  defined by

$$\Pr \left[ \text{Vfy}(\text{pk}, \sigma^*, m^*) = 1 \mid \begin{array}{l} \text{gp} \leftarrow \text{SetupGrp}(1^n) \\ (\text{pk}, \text{sk}) \leftarrow \text{Gen}(\text{gp}) \\ (m^*, \sigma^*) \leftarrow \mathcal{A}^{\text{Sgn}(\text{sk}, \cdot)}(1^n, \text{pk}) \\ m^* \notin \{m_1, \dots, m_q\} \end{array} \right]$$

is negligible in  $n$ , where  $\text{Sgn}(\text{sk}, \cdot)$  is an oracle that, on input  $m$ , returns  $\text{Sgn}(\text{sk}, m)$ , and  $\{m_1, \dots, m_q\}$  denotes the set of messages queried by  $\mathcal{A}$  to its oracle.

*Our Instantiation.* As we need to prove statements about signatures, the signature scheme has to be algebraic. For our construction, we use the structure-preserving signature scheme of Abe et al. [1], which is currently the most efficient structure-preserving signature scheme. Its EUF-CMA security proof is in the generic group model, a restriction we consider reasonable with respect to our goal of constructing a highly efficient P4TC scheme. An alternative secure in the plain model would be [53]. For the scheme in [1], one needs to fix two additional parameters  $\mu, \nu \in \mathbb{N}_0$  defining the actual message space  $G_1^\nu \times G_2^\mu$ . Then  $\text{sk} \in \mathbb{Z}_p^{\mu+\nu+2}$ ,  $\text{pk} \in G_1^{\mu+2} \times G_2^\nu$  and  $\sigma \in G_2^2 \times G_1$ .

We use the signature scheme  $S$  from Abe et al. [1] in the following ways in our system:

- In the Wallet Issuing and Debt Accumulation tasks we use  $S$  for messages from  $G_2 \times G_1$  ( $\nu = 1$  and  $\mu = 1$ ).
- In the Wallet Issuing task we use  $S$  for messages from  $G_1^{2\ell+2}$  ( $\nu = 2\ell + 2$  and  $\mu = 0$ ).
- In the RSU Certification and TSP Registration tasks we use  $S$  for messages from  $G_1^{3+y}$  ( $\nu = 3 + y$  and  $\mu = 0$ ).

*C.2.5 Asymmetric Encryption.* We use the standard definitions for asymmetric encryption schemes and corresponding security notions, except that we enhance them with a  $\text{SetupGrp}$  algorithm to fit our algebraic setting.

*Definition C.9 (Asymmetric Encryption).* An asymmetric encryption scheme  $E := (\text{SetupGrp}, \text{Gen}, \text{Enc}, \text{Dec})$  consists of four PPT algorithms:

- $\text{SetupGrp}$  takes as input a security parameter  $1^n$  and outputs public parameters  $\text{gp}$ . We assume that  $\text{gp}$  is given as implicit input to all algorithms.
- $\text{Gen}(\text{gp})$  outputs a pair  $(\text{pk}, \text{sk})$  of keys, where  $\text{pk}$  is the (public) encryption key and  $\text{sk}$  is the (secret) decryption key.
- $\text{Enc}(\text{pk}, m)$  takes a key  $\text{pk}$  and a plaintext message  $m \in \mathcal{M}$  and outputs a ciphertext  $c$ .

- $\text{Dec}(\text{sk}, c)$  takes a key  $\text{sk}$  and a ciphertext  $c$  and outputs a plaintext message  $m$  or  $\perp$ . We assume that  $\text{Dec}$  is deterministic.

Correctness is defined in the usual sense.

An asymmetric encryption scheme  $E$  is *IND-CCA2-secure* if for all PPT adversaries  $\mathcal{A}$  it holds that the advantage  $\text{Adv}_{E, \mathcal{A}}^{\text{IND-CCA-asym}}(1^n)$  defined by

$$\Pr \left[ b = b' \mid \begin{array}{l} \text{gp} \leftarrow \text{SetupGrp}(1^n) \\ (\text{pk}, \text{sk}) \leftarrow \text{Gen}(\text{gp}) \\ (\text{state}, m_0, m_1) \leftarrow \mathcal{A}^{\text{Dec}(\text{sk}, \cdot)}(1^n, \text{pk}) \\ b \xleftarrow{R} \{0, 1\} \\ c^* \leftarrow \text{Enc}(\text{pk}, m_b) \\ b' \leftarrow \mathcal{A}^{\text{Dec}'(\text{sk}, \cdot)}(\text{state}, c^*) \end{array} \right] - \frac{1}{2}$$

is negligible in  $n$ , where  $|m_0| = |m_1|$ ,  $\text{Dec}(\text{sk}, \cdot)$  is an oracle that gets a ciphertext  $c$  from the adversary and returns  $\text{Dec}(\text{sk}, c)$  and  $\text{Dec}'(\text{sk}, \cdot)$  is the same, except that it returns  $\perp$  on input  $c^*$ .

An asymmetric encryption scheme  $E$  is *NM-CCA2-secure* if for all PPT adversaries  $\mathcal{A}$  it holds that the advantage  $\text{Adv}_{E, \mathcal{A}}^{\text{NM-CCA}}(1^n)$  defined by

$$\left| \text{Succs}_{E, \mathcal{A}, \text{real}}^{\text{NM-CCA}}(1^n) - \text{Succs}_{E, \mathcal{A}, \text{random}}^{\text{NM-CCA}}(1^n) \right|$$

is negligible with

$$\text{Succs}_{E, \mathcal{A}, \text{real}}^{\text{NM-CCA}}(1^n) := \Pr \left[ \begin{array}{l} c \notin \mathcal{C} \wedge \\ \perp \notin \mathcal{M} \wedge \\ R(m, \mathbf{m}) = 1 \end{array} \mid \begin{array}{l} \text{gp} \leftarrow \text{SetupGrp}(1^n) \\ (\text{pk}, \text{sk}) \leftarrow \text{Gen}(\text{gp}) \\ (M, \text{state}) \leftarrow \mathcal{A}^{\text{Dec}(\text{sk}, \cdot)}(1^n, \text{pk}) \\ m \xleftarrow{R} M \\ c \leftarrow \text{Enc}(\text{pk}, m) \\ (R, \mathbf{c}) \leftarrow \mathcal{A}^{\text{Dec}'(\text{sk}, \cdot)}(1^n, \text{state}, c) \\ \mathbf{m} \leftarrow \text{Dec}(\text{sk}, c) \end{array} \right]$$

and

$$\text{Succs}_{E, \mathcal{A}, \text{random}}^{\text{NM-CCA}}(1^n) := \Pr \left[ \begin{array}{l} c \notin \mathcal{C} \wedge \\ \perp \notin \mathcal{M} \wedge \\ R(\tilde{m}, \mathbf{m}) = 1 \end{array} \mid \begin{array}{l} \text{gp} \leftarrow \text{SetupGrp}(1^n) \\ (\text{pk}, \text{sk}) \leftarrow \text{Gen}(\text{gp}) \\ (M, \text{state}) \leftarrow \mathcal{A}^{\text{Dec}(\text{sk}, \cdot)}(1^n, \text{pk}) \\ m, \tilde{m} \xleftarrow{R} M \\ c \leftarrow \text{Enc}(\text{pk}, m) \\ (R, \mathbf{c}) \leftarrow \mathcal{A}^{\text{Dec}'(\text{sk}, \cdot)}(1^n, \text{state}, c) \\ \mathbf{m} \leftarrow \text{Dec}(\text{sk}, c) \end{array} \right],$$

where  $M$  denotes a space of valid, equally long messages,  $R \subseteq M \times M^*$  denotes an relation,  $\text{Dec}(\text{sk}, \cdot)$  is an oracle that gets a ciphertext  $c$  from the adversary and returns  $\text{Dec}(\text{sk}, c)$  and  $\text{Dec}'(\text{sk}, \cdot)$  is the same, except that it returns  $\perp$  on input  $c$ .

An encryption is *IND-CCA2 secure* if and only if it is *NM-CCA2 secure* [13].

*Our Instantiation.* We will make use of two different *IND-CCA2-secure* encryption schemes:

- We implicitly use the encryption scheme by Cash et al. [26] to realize the secure channels underlying our model.
- We use a variant of Camenisch et al. [16] to instantiate the explicit encryption scheme  $E1$  for the deposit Wallet IDs.

The former scheme is based on the TWIN-DH assumption and is used to setup a session key for a symmetric encryption of all protocol messages (cp. [Appendix C.2.6](#)) in the usual way.

The latter scheme is an adapted variant of the structure-preserving, IND-CCA2 secure encryption scheme by Camenisch et al. [16]. Thus, some remarks are in order. The original scheme is formalized for a symmetric type-1 pairing, but we need a scheme that is secure in the asymmetric type-3 case. For the conversion we followed the generic transformation proposed by Abe et al. [2] with some additional, manual optimizations. The transformed scheme encrypts vectors of  $G_1$ -elements and is secure under the co-DLIN assumption (cp. [Definition C.4](#)) which holds in the generic group model. This follows automatically from [2] (or can also be easily seen by inspecting the original proof in [16]). We present the modified scheme in full detail.

*Definition C.10 (Type-3 variant of Camenisch et al. [16]).* Let  $\text{gp} := (G_1, G_2, G_T, e, \nu, g_1, g_2) \leftarrow \text{SetupGrp}(1^n)$  be as in [Definition C.1](#). Let  $\wp$  be the dimension of the message space  $G_1^\wp$ . The algorithms Gen, Enc and Dec are depicted in [Fig. 17](#).

We instantiate this scheme with  $\wp = \ell + 2$ .

**C.2.6 Symmetric Encryption.** We use standard definitions for symmetric encryption schemes and corresponding security notions.

*Definition C.11 (Symmetric Encryption).* A symmetric encryption scheme  $E := (\text{Gen}, \text{Enc}, \text{Dec})$  consists of three PPT algorithms:

- $\text{Gen}(1^n)$  outputs a (random) key  $k$ .
- $\text{Enc}(k, m)$  takes a key  $k$  and a plaintext message  $m \in \mathcal{M}$  and outputs a ciphertext  $c$ .
- $\text{Dec}(k, c)$  takes a key  $k$  and a ciphertext  $c$  and outputs a plaintext message  $m$  or  $\perp$ . We assume that Dec is deterministic.

As for asymmetric encryption, we require correctness in the usual sense.

We now define a multi-message version of IND-CCA2 security. It is a well-known fact that IND-CCA2 security in the multi-message setting is equivalent to standard IND-CCA2 security. (This can be shown via a standard hybrid argument.)

*Definition C.12 (IND-CCA2-Security for Symmetric Encryption).* A symmetric encryption scheme  $E$  is IND-CCA2-secure if for all PPT adversaries  $\mathcal{A}$  it holds that the advantage  $\text{Adv}_{E, \mathcal{A}}^{\text{IND-CCA-sym}}(1^n)$  defined by

$$\left| \Pr \left[ b = b' \mid \begin{array}{l} k \leftarrow \text{Gen}(1^n) \\ (state, j, \mathbf{m}_0, \mathbf{m}_1) \leftarrow \mathcal{A}^{\text{Enc}(k, \cdot), \text{Dec}(k, \cdot)}(1^n) \\ b \leftarrow \{0, 1\} \\ \mathbf{c}^* \leftarrow (\text{Enc}(k, m_{b,1}), \dots, \text{Enc}(k, m_{b,j})) \\ b' \leftarrow \mathcal{A}^{\text{Enc}(k, \cdot), \text{Dec}'(k, \cdot)}(state, \mathbf{c}^*) \end{array} \right] - \frac{1}{2} \right|$$

is negligible in  $n$ , where  $\mathbf{m}_0, \mathbf{m}_1$  are two vectors of  $j \in \mathbb{N}$  bitstrings each such that for all  $1 \leq i \leq j$ :  $|m_{0,i}| = |m_{1,i}|$ ,  $\text{Enc}(k, \cdot)$  and  $\text{Dec}(k, \cdot)$  denote oracles that return  $\text{Enc}(k, m)$  and  $\text{Dec}(k, c)$  for a  $m$  or  $c$  chosen by the adversary, and  $\text{Dec}'(k, \cdot)$  is the same as  $\text{Dec}(k, \cdot)$ , except that it returns  $\perp$  on input of any  $c_i^*$  that is contained in  $\mathbf{c}^*$ .

*Our Instantiation.* We use an IND-CCA2-secure symmetric encryption scheme in our protocol to encrypt the exchanged protocol messages. To this end, we combine an IND-CCA2-secure asymmetric encryption (see [Appendix C.2.5](#)) with an IND-CCA2-secure symmetric encryption in the usual KEM/DEM approach. The symmetric encryption can for example be instantiated with AES in CBC mode together with HMAC based on the SHA-256 hash function. The result will be IND-CCA2-secure if AES is a pseudo-random permutation and the SHA-256 compression function is a PRF when the data input is seen as the key [12].

**Gen(gp,  $\varphi$ )**

parse  $(G_1, G_2, G_T, e, p, g_1, g_2) := \text{gp}$

$\alpha_1, \dots, \alpha_\varphi, \beta_0, \dots, \beta_3, \gamma_1, \dots, \gamma_\varphi \xleftarrow{\mathbb{R}} \mathbb{Z}_p^3$

$\text{sk} := (\{\alpha_i\}_{i=1, \dots, \varphi}, \{\beta_i\}_{i=0, \dots, 3}, \{\gamma_i\}_{i=1, \dots, \varphi})$

$\xi_1, \dots, \xi_3 \xleftarrow{\mathbb{R}} \mathbb{Z}_p^*$

$\check{h}_1 := g_1^{\xi_1}, \quad \check{h}_2 := g_1^{\xi_2}, \quad \check{h}_3 := g_1^{\xi_3}$

$\hat{h}_1 := g_2^{\xi_1}, \quad \hat{h}_2 := g_2^{\xi_2}, \quad \hat{h}_3 := g_2^{\xi_3}$

$x_{i,1} := \check{h}_1^{\alpha_{i,1}} \check{h}_3^{\alpha_{i,3}}, \quad x_{i,2} := \check{h}_2^{\alpha_{i,2}} \check{h}_3^{\alpha_{i,3}}, \quad \text{for } i = 1, \dots, \varphi$

$y_{i,1} := \hat{h}_1^{\beta_{i,1}} \hat{h}_3^{\beta_{i,3}}, \quad y_{i,2} := \hat{h}_2^{\beta_{i,2}} \hat{h}_3^{\beta_{i,3}}, \quad \text{for } i = 0, \dots, 3$

$z_{i,1} := \hat{h}_1^{\gamma_{i,1}} \hat{h}_3^{\gamma_{i,3}}, \quad z_{i,2} := \hat{h}_2^{\gamma_{i,2}} \hat{h}_3^{\gamma_{i,3}}, \quad \text{for } i = 1, \dots, \varphi$

$\text{pk} := (\check{h}_1, \check{h}_2, \check{h}_3, \hat{h}_1, \hat{h}_2, \hat{h}_3,$   
 $\{x_{i,1}, x_{i,2}\}_{i=1, \dots, \varphi}, \{y_{i,1}, y_{i,2}\}_{i=0, \dots, 3}, \{z_{i,1}, z_{i,2}\}_{i=1, \dots, \varphi})$

**return** (pk, sk)

(a) Key generation algorithm Gen

**Enc(pk, m)**

parse  $(\check{h}_1, \check{h}_2, \check{h}_3, \hat{h}_1, \hat{h}_2, \hat{h}_3, \{x_{i,1}, x_{i,2}\}_{i=1, \dots, \varphi},$   
 $\{y_{i,1}, y_{i,2}\}_{i=0, \dots, 3}, \{z_{i,1}, z_{i,2}\}_{i=1, \dots, \varphi}) := \text{pk}$

$r, s \xleftarrow{\mathbb{R}} \mathbb{Z}_p$

$\check{u}_1 := \check{h}_1^r, \quad \check{u}_2 := \check{h}_2^s, \quad \check{u}_3 := \check{h}_3^{r+s}$

$\hat{u}_1 := \hat{h}_1^r, \quad \hat{u}_2 := \hat{h}_2^s, \quad \hat{u}_3 := \hat{h}_3^{r+s}$

$c_i = m_i x_{i,1}^r x_{i,2}^s$  for  $i = 1, \dots, \varphi$

$v = \prod_{i=0}^3 e(\check{u}_i, y_{i,1}^r y_{i,2}^s) \prod_{i=1}^{\varphi} e(c_i, z_{i,1}^r z_{i,2}^s)$  with  $\check{u}_0 := g_1$

$c := (u, c, v)$  with  $u := (\check{u}_1, \check{u}_2, \check{u}_3, \hat{u}_1, \hat{u}_2, \hat{u}_3)$  and  $c := (c_1, \dots, c_\varphi)$

**return** (c)

(b) Encryption algorithm Enc

**Dec(sk, c)**

parse  $(\{\alpha_i\}_{i=1, \dots, \varphi}, \{\beta_i\}_{i=0, \dots, 3}, \{\gamma_i\}_{i=1, \dots, \varphi}) := \text{sk}$

parse  $(u, c, v) := c, (\check{u}_1, \check{u}_2, \check{u}_3, \hat{u}_1, \hat{u}_2, \hat{u}_3) := u$  and  $(c_1, \dots, c_\varphi) := c$

$\check{u}_0 := g_1$

**if**  $v \neq \prod_{i=0}^3 e(\check{u}_i, \hat{u}_1^{\beta_{i,1}} \hat{u}_2^{\beta_{i,2}} \hat{u}_3^{\beta_{i,3}}) \prod_{i=1}^{\varphi} e(c_i, \hat{u}_1^{\gamma_{i,1}} \hat{u}_2^{\gamma_{i,2}} \hat{u}_3^{\gamma_{i,3}})$  **abort**

**if**  $e(\check{u}_i, g_2) \neq e(g_1, \hat{u}_i)$  for any  $i \in \{1, 2, 3\}$  **abort**

$m_i := c_i \check{u}_1^{-\alpha_{i,1}} \check{u}_2^{-\alpha_{i,2}} \check{u}_3^{-\alpha_{i,3}}$  for  $i \in \{1, \dots, \varphi\}$

$m := (m_1, \dots, m_\varphi)$

**return** (m)

(c) Decryption algorithm Dec

Fig. 17. The adapted CCA-secure encryption scheme by Camenisch et al. [16] with parameter  $\varphi$  and message space  $G_1^\varphi$ 

**C.2.7 Pseudo-Random Functions.** A pseudo-random function (PRF)  $F : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$  is a keyed function whose output cannot be distinguished from randomness, i.e., any PPT adversary given oracle access to either  $F(k, \cdot)$  or a randomly chosen function  $R : \mathcal{X} \rightarrow \mathcal{Y}$ , cannot distinguish between them with non-negligible probability. More precisely, a PRF—more precisely a family of PRF’s in the security parameter  $1^n$ —is defined as follows.

*Definition C.13.* A (group-based) *pseudo-random function* (PRF)  $\text{PRF} := (\text{SetupGrp}, \text{Gen}, \text{Eval})$  consists of three PPT algorithms:

- $\text{SetupGrp}$  takes as input a security parameter  $1^n$  and outputs public parameters  $\text{gp}$ . We assume that  $\text{gp}$  is given as implicit input to all algorithms. The input domain  $\mathcal{X}_{\text{gp}}$ , the key space  $\mathcal{K}_{\text{gp}}$ , and the codomain  $\mathcal{Y}_{\text{gp}}$  may all depend on  $\text{gp}$ .
- $\text{Gen}$  takes  $\text{gp}$  as input and outputs a key  $k \in \mathcal{K}_{\text{gp}}$ . (Typically, we have  $k \stackrel{\text{R}}{\leftarrow} \mathcal{K}_{\text{gp}}$ .)
- $\text{Eval}$  is a deterministic algorithm which takes as input a key  $k \in \mathcal{K}_{\text{gp}}$  and a value  $x \in \mathcal{X}_{\text{gp}}$ , and outputs some  $y \in \mathcal{Y}_{\text{gp}}$ . Usually, we simply write  $y = F(k, x)$  for short.

We say that PRF is secure if for all PPT adversaries  $\mathcal{A}$  it holds that the advantage  $\text{Adv}_{\mathcal{A}}^{\text{prf}}(1^n)$  defined by

$$\left| \Pr \left[ 1 \leftarrow \mathcal{A}^{F(k, \cdot)}(\text{gp}) \mid \begin{array}{l} \text{gp} \leftarrow \text{SetupGrp}(1^n), \\ k \leftarrow \text{Gen}(\text{gp}) \end{array} \right] - \Pr \left[ 1 \leftarrow \mathcal{A}^{R(\cdot)}(\text{gp}) \mid \begin{array}{l} \text{gp} \leftarrow \text{SetupGrp}(1^n), \\ R \stackrel{\text{R}}{\leftarrow} \{R : \mathcal{X}_{\text{gp}} \rightarrow \mathcal{Y}_{\text{gp}}\} \end{array} \right] \right|$$

is negligible in  $n$ .

*Our Instantiation.* As we want to efficiently prove statements about PRF outputs, we use an efficient algebraic construction, namely the Dodis-Yampolskiy PRF [36]. This function is defined by  $F(k, x) : \mathbb{Z}_p^2 \rightarrow G_1, (k, x) \mapsto g_1^{\frac{1}{x+k}}$ , where  $k \stackrel{\text{R}}{\leftarrow} \mathbb{Z}_p$  is the random PRF key. It is secure for inputs  $\{0, \dots, n_{\text{PRF}}\} \subset \mathbb{Z}_p$  under the  $n_{\text{PRF}}$ -DDHI assumption. This is a family of increasingly stronger assumptions which is assumed to hold for asymmetric bilinear groups.

### C.3 Types of Equations Supported by GS-NIZKs

Let  $\text{SetupGrp}$  be a bilinear group generator (cf. [Definition C.1](#)) for which the SXDH assumption holds and  $\text{gp} := (G_1, G_2, G_T, e, p, g_1, g_2, g_T) \leftarrow \text{SetupGrp}(1^n)$  denotes the output of  $\text{SetupGrp}$ . Furthermore, let  $X_1, \dots, X_{m_1} \in G_1, x_1, \dots, x_{m_2} \in \mathbb{Z}_p, Y_1, \dots, Y_{m_3} \in G_2,$  and  $y_1, \dots, y_{m_4} \in \mathbb{Z}_p$  denote variables in the following types of equations:

- *Pairing-Product Equation (PPE):*

$$\prod_{i=1}^{m_3} e(A_i, Y_i) \prod_{i=1}^{m_1} e(X_i, B_i) \prod_{i=1}^{m_1} \prod_{j=1}^{m_3} e(X_i, Y_i)^{y_{i,j}} = t_T$$

for constants  $A_i \in G_1, B_i \in G_2, t_T \in G_T, Y_{i,j} \in \mathbb{Z}_p$ .

- *Multi-Scalar Equation (MSE) over  $G_1$ :*

$$\prod_{i=1}^{m_4} A_i^{y_i} \prod_{i=1}^{m_1} X_i^{b_i} \prod_{i=1}^{m_1} \prod_{j=1}^{m_4} X_i^{y_i y_j} = t_1$$

for constants  $A_i, t_1 \in G_1, b_i, Y_{i,j} \in \mathbb{Z}_p$ .

- *Multi-Scalar Equation (MSE) over  $G_2$ :*

$$\prod_{i=1}^{m_2} B_i^{x_i} \prod_{i=1}^{m_3} Y_i^{a_i} \prod_{i=1}^{m_2} \prod_{j=1}^{m_3} Y_j^{y_i x_i} = t_2$$

for constants  $B_i, t_2 \in G_2, a_i, Y_{i,j} \in \mathbb{Z}_p$ .

- *Quadratic Equation (QE) over  $\mathbb{Z}_p$* :

$$\sum_{i=1}^{m_4} a_i y_i + \sum_{i=1}^{m_2} x_i b_i + \sum_{i=1}^{m_2} \sum_{j=1}^{m_4} \gamma_{i,j} x_i y_j = t$$

for constants  $a_i, b_i, \gamma_{i,j}, t \in \mathbb{Z}_p$ .

Let  $L_{\text{gp}}$  be a language containing statements described by the conjunction of  $n_1$  pairing-product equations over  $\text{gp}$ ,  $n_2$  multi-scalar equations over  $G_1$ ,  $n_3$  multi-scalar equations over  $G_2$ , and  $n_4$  quadratic equations over  $\mathbb{Z}_p$ , where  $n_i \in \mathbb{N}_0$  are constants, as well as by witnesses

$$w = (X_1, \dots, X_{m_1}, x_1, \dots, x_{m_2}, Y_1, \dots, Y_{m_3}, y_1, \dots, y_{m_4}),$$

where  $m_i \in \mathbb{N}_0$ . Then the Groth-Sahai proof system for  $L_{\text{gp}}$ , as introduced by [40], is perfectly correct, perfectly sound, and satisfies  $F_{\text{gp}}$ -extractability [37, 40] for

$$F_{\text{gp}} : G_1^{m_1} \times \mathbb{Z}_p^{m_2} \times G_2^{m_3} \times \mathbb{Z}_p^{m_4} \rightarrow G_1^{m_1} \times G_1^{m_2} \times G_2^{m_3} \times G_2^{m_4}$$

with

$$F_{\text{gp}}(w) := ((X_i)_{i \in [m_1]}, (g_1^{x_i})_{i \in [m_2]}, (Y_i)_{i \in [m_3]}, (g_2^{y_i})_{i \in [m_4]}).^{23}$$

It is also known to be composable zero-knowledge [37, 40] as long as for all PPEs in  $L_{\text{gp}}$  holds that either

- $t_T = 1$  or
- the right-hand side of the PPE can be written as  $\prod_{i=1}^k e(A_i, B_i)$  for constants  $A_i \in G_1, B_i \in G_2$ , such that for each  $i$   $\text{DLOG}(A_i)$  or  $\text{DLOG}(B_i)$  is known.

In the latter case, *hint* from Definition C.5 would contain these discrete logarithms which would simply be put (as additional elements) into the simulation trapdoor  $\text{td}_{\text{sPok}}$ . Also note that if these discrete logarithms are not known, there is a workaround which consists of adding new helper variables to  $L_{\text{gp}}$  [40].

## D FULL PROTOCOL DESCRIPTION

In this appendix we describe and define a real protocol  $\pi_{\text{P4TC}}$  that implements our toll collection system  $\mathcal{F}_{\text{P4TC}}$ . We say

*Definition D.1 (P4TC Scheme).* A protocol  $\pi$  is called a *privacy-preserving electronic toll collection* scheme, if it GUC-realizes  $\mathcal{F}_{\text{P4TC}}$ .

The proof that  $\pi_{\text{P4TC}}$  is a GUC-realization of  $\mathcal{F}_{\text{P4TC}}$  is postponed to Appendix F. The style of the presentation follows the same lines as the presentation of the ideal model  $\mathcal{F}_{\text{P4TC}}$  in Appendix B: Although  $\pi_{\text{P4TC}}$  is a single, monolithic protocol with different tasks, the individual tasks are presented as if they were individual protocols.

While in the ideal model all information is kept in a single, pervasive, trustworthy database, in the real model such a database does not exist. Instead, the state of the system is distributed across all parties. Each party locally stores a piece of information: The user owns a “user wallet”, which is updated during each transaction, the RSU collects “double-spending tags” as well as “proof of participation challenges”, which are periodically sent to the TSP, and the TSP creates and keeps “hidden user trapdoors” for each wallet issued. A precise definition what is stored by which party is depicted in Fig. 18. For typographic reasons we additionally split the presentation of most tasks into a *wrapper protocol* and a *core protocol*. Except for a few cases, there is a one-to-one correspondence between wrapper and core protocols. The wrapper protocols have the same input/output interfaces as their ideal counterparts and describe steps that are executed by each party locally before and after the respective core protocol. These steps include loading keys, parsing the previously stored state, persisting the new state after

<sup>23</sup> $F$  acts as identity function on group elements  $a \in G_1$  and  $b \in G_2$  but returns  $g_1^s \in G_1$  or  $g_2^s \in G_2$  for exponents  $s \in \mathbb{Z}_p$ .

**UC-Protocol**  $\pi_{P4TC}$ *I. Local State*

- (1) The TSP internally records:
  - It's public and private key  $(pk_{\mathcal{T}}, sk_{\mathcal{T}})$ .
  - A self-signed certificate  $cert_{\mathcal{T}}^{\mathcal{R}}$ .
  - A set  $HTD$  of hidden user trapdoors.
  - A (partial) mapping  $\{pk_{\mathcal{R}} \mapsto a_{\mathcal{R}}\}$  of RSU attributes.
- (2) Each RSU internally records:
  - It's public and private key  $(pk_{\mathcal{R}}, sk_{\mathcal{R}})$ .
  - A certificate  $cert_{\mathcal{R}}$  signed by the TSP.
  - Sets  $\Omega_{\mathcal{R}}^{dsp}$ ,  $\Omega_{\mathcal{R}}^{bl}$  and  $\Omega_{\mathcal{R}}^{pp}$  of transaction information for double-spending detection, blacklisting and prove participation respectively.
- (3) Each user internally records:
  - His public and private key  $(pk_{\mathcal{U}}, sk_{\mathcal{U}})$ .
  - A set  $\{\tau\}$  of all past tokens issued to him.
  - A set  $\Omega_{\mathcal{U}}^{pp}$  of transaction information for prove participation.

*II. Behavior*

- |                                      |                                      |  |
|--------------------------------------|--------------------------------------|--|
| • <i>DR Registration</i> (Fig. 20)   | • <i>Wallet Issuing</i> (Fig. 30)    | • <i>Prove Participation</i> (Fig. 36)       |
| • <i>TSP Registration</i> (Fig. 22)  | • <i>Debt Accumulation</i> (Fig. 32) | • <i>Double-Spending Detection</i> (Fig. 38) |
| • <i>RSU Registration</i> (Fig. 24)  | • <i>Debt Clearance</i> (Fig. 34)    | • <i>Guilt Verification</i> (Fig. 39)        |
| • <i>User Registration</i> (Fig. 26) |                                      | • <i>User Blacklisting</i> (Fig. 40)         |
| • <i>RSU Certification</i> (Fig. 28) |                                      |  |

Fig. 18. The UC-protocol  $\pi_{P4TC}$ 

the core protocol has returned, etc. The core protocols describe the actual interaction between parties and what messages are exchanged.

This dichotomy between wrapper and core protocols is lifted for four exceptions:

- (1) We give an algorithm for the setup of the system (cf. Fig. 19) which explains how the CRS is generated. Of course, there is no wrapper protocol because setup of the CRS is not even part of our protocol but part of the setup assumption and provided by  $\mathcal{F}_{CRS}$ .
- (2) We describe a “utility algorithm” *WalletVerification* (cf. Fig. 42). This algorithm has no purpose on its own, but simply collects some shared code of multiple tasks.
- (3)+(4) We only have “wrapper protocols” for the tasks *Double-Spending Detection* and *Guilt Verification* (cf. Figs. 38 and 39) because they are so simple that splitting it each into two yields no advantage.

**D.1 Secure Authenticated Channels**

In our system, all protocol messages are encrypted using CCA-secure encryption. For this purpose, a new session key chosen by the user is encrypted under the public key of an RSU/TSP for each interaction. Furthermore, we make use of fully authenticated channels. The only exception to this is the task of *Debt Accumulation* where only the participating RSU authenticates itself to the user who in turn remains anonymous. We omit these encryptions and authentications when describing the protocols.

$\text{Setup}(1^n, \mathcal{B})$ <hr style="border: 0.5px solid black;"/> $\text{gp} := (G_1, G_2, G_T, e, p, g_1, g_2) \leftarrow \text{SetupGrp}(1^n)$ $\text{CRS}_{\text{com}}^1 \leftarrow \text{C1.Gen}(\text{gp})$ $\text{CRS}_{\text{com}}^2 \leftarrow \text{C2.Gen}(\text{gp})$ $\text{CRS}_{\text{pok}} \leftarrow \text{SetupPoK}(\text{gp})$ $\text{CRS} := (\text{gp}, \mathcal{B}, \text{CRS}_{\text{com}}^1, \text{CRS}_{\text{com}}^2, \text{CRS}_{\text{pok}})$ $\text{return CRS}$
--

Fig. 19. System Setup Algorithm

## D.2 Wallets

A central component of our toll collection system is the wallet that is created during Wallet Issuing. It is of the form

$$\tau := (s, \phi, x^{\text{next}}, \lambda, \mathbf{a}_{\mathcal{U}}, c_{\mathcal{R}}, d_{\mathcal{R}}, \sigma_{\mathcal{R}}, \text{cert}_{\mathcal{R}}, c_{\mathcal{T}}, d_{\mathcal{T}}, \sigma_{\mathcal{T}}, b, u_1^{\text{next}}).$$

Some of the components are fixed after creation, some change after every transaction. The fixed components consist of the *wallet ID*  $\lambda$  (which is also used as the PRF seed), the *user attributes*  $\mathbf{a}_{\mathcal{U}}$ , the *TSP commitment*  $c_{\mathcal{T}}$  (a commitment on  $\lambda$  and the secret user key  $\text{sk}_{\mathcal{U}}$ ), its corresponding opening  $d_{\mathcal{T}}$  and a signature  $\sigma_{\mathcal{T}}$  on  $c_{\mathcal{T}}$  and  $\mathbf{a}_{\mathcal{U}}$  created by the TSP.

The alterable components consist of the *RSU commitment*  $c_{\mathcal{R}}$  (a commitment on  $\lambda$ ,  $b$ ,  $u_1^{\text{next}}$  and  $x^{\text{next}}$ ), its corresponding opening  $d_{\mathcal{R}}$ , a signature  $\sigma_{\mathcal{R}}$  on  $c_{\mathcal{R}}$  and  $s$  created by a RSU, the *balance*  $b$ , the *double-spending mask*  $u_1^{\text{next}}$  for the next transaction, the *PRF counter*  $x^{\text{next}}$  for the next interaction, a *RSU certificate*  $\text{cert}_{\mathcal{R}}$  and the *serial number*  $s$  and *fraud detection ID*  $\phi := \text{PRF}(\lambda, x^{\text{next}} - 1)$  for the current transaction. These components change after each interaction with an RSU via the Debt Accumulation task.

In the following, a protocol or algorithm for each task is presented. For a better overview, it is depicted in Fig. 1 which parties are involved in each task (except for some registration tasks). Also, the used variables are summarized in Tables 4 and 5.

## D.3 System Setup

To setup the system once (see Fig. 19), the public parameter CRS must be generated in a trustworthy way. The CRS consists of a description of the underlying algebraic framework, a splitting base  $\mathcal{B}$  and the individual CRSs for the used commitments and zero-knowledge proofs. We assume that the CRS is implicitly available to all protocols and algorithms. Either a number of mutually distrusting parties run a multi-party computation (using some other sort of setup assumption) to generate the CRS or a commonly trusted party is charged with this task. As a trusted-third party (the DR) explicitly participates in our system (for resolving disputes), this party could also run the system setup.

## D.4 Registration

The registration algorithms of  $\text{DR}$ ,  $\mathcal{T}$ ,  $\mathcal{R}$  and  $\mathcal{U}$  are all presented with a wrapper protocol  $\pi_{\text{P4TC}}$  (see Figs. 20, 22, 24 and 26) and a core protocol (see Figs. 21, 23, 25 and 27). The wrapper protocol interacts with other UC entities, pre-processes the inputs, post-processes the outputs and internally invokes the core protocols. In the following, only the mechanics of the core protocols are explicitly described.



Table 5. Notation that only occurs in the real protocol

Identifier	Domain	Description
$sk_{DR}$	$\mathbb{Z}_p^{2\ell+8}$	secret DR key
$sk_{\mathcal{T}}$	$\mathbb{Z}_p^{j+3} \times \mathbb{Z}_p^{y+6} \times \mathbb{Z}_p^4$	secret TSP key
$pk_{\mathcal{T}}^{\mathcal{T}}$	$G_1^{j+3}$	public TSP commitment signing key (part of $pk_{\mathcal{T}}$ )
$sk_{\mathcal{T}}^{\mathcal{T}}$	$\mathbb{Z}_p^{j+3}$	secret TSP commitment signing key (part of $sk_{\mathcal{T}}$ )
$pk_{\mathcal{T}}^{\text{cert}}$	$G_1^3 \times G_2^{y+3}$	public certification key (part of $pk_{\mathcal{T}}$ )
$sk_{\mathcal{T}}^{\text{cert}}$	$\mathbb{Z}_p^{y+6}$	secret certification key (part of $sk_{\mathcal{T}}$ )
$pk_{\mathcal{T}}^{\mathcal{R}}$	$G_1^3 \times G_2$	public TSP's RSU commitment signing key (part of $pk_{\mathcal{T}}$ )
$sk_{\mathcal{T}}^{\mathcal{R}}$	$\mathbb{Z}_p^4$	secret TSP's RSU commitment signing key (part of $sk_{\mathcal{T}}$ )
$sk_{\mathcal{R}}$	$\mathbb{Z}_p^4$	secret RSU key
$sk_{\mathcal{U}}$	$\mathbb{Z}_p$	secret user key
$c_{\mathcal{T}}$	$G_2$	TSP commitment
$d_{\mathcal{T}}$	$G_1$	decommitment of $c_{\mathcal{T}}$
$\sigma_{\mathcal{T}}$	$G_2^2 \times G_1$	signature on $c_{\mathcal{T}}$ and $a_{\mathcal{U}}$
$c_{\mathcal{R}}$	$G_2$	RSU commitment
$d_{\mathcal{R}}$	$G_1$	decommitment of $c_{\mathcal{R}}$
$\sigma_{\mathcal{R}}$	$G_2^2 \times G_1$	signature on $c_{\mathcal{R}}$ and $s$
$\text{cert}_{\mathcal{R}}$	$G_1^3 \times G_1^y \times (G_2^2 \times G_1)$	RSU certificate
$\text{cert}_{\mathcal{T}}^{\mathcal{R}}$	$G_1^3 \times G_1^y \times (G_2^2 \times G_1)$	TSP certificate
$\sigma_{\mathcal{R}}^{\text{cert}}$	$G_2^2 \times G_1$	signature on $pk_{\mathcal{R}}$ and $a_{\mathcal{R}}$
$\sigma_{\mathcal{T}}^{\text{cert}}$	$G_2^2 \times G_1$	signature on $pk_{\mathcal{T}}^{\mathcal{R}}$ and $a_{\mathcal{T}}$
$c_{\text{hid}}$	$G_2$	hidden ID
$d_{\text{hid}}$	$G_1$	opening of hidden ID
$c'_{\text{seed}}$	$G_2$	commitment on the user half of the wallet ID
$d'_{\text{seed}}$	$G_1$	opening of $c'_{\text{seed}}$
$c''_{\text{ser}}$	$G_1^2$	commitment on the RSU half of the serial number
$d''_{\text{ser}}$	$\mathbb{Z}_p^2$	opening of $c''_{\text{ser}}$
$\omega^{\text{dsp}}$	$G_1 \times \mathbb{Z}_p \times \mathbb{Z}_p$	transaction information for double-spending
$\omega^{\text{bl}}$	$G_1 \times \mathbb{Z}_p$	transaction information for blacklisting
$\omega_{\mathcal{U}}^{\text{PP}}$	$G_1 \times G_2 \times G_1$	user transaction information prove participation
$\omega_{\mathcal{R}}^{\text{PP}}$	$G_1 \times G_2$	RSU transaction information for prove participation
$u_1$	$\mathbb{Z}_p$	double-spending mask
$u_2$	$\mathbb{Z}_p$	double-spending randomness
$t$	$\mathbb{Z}_p$	double-spending tag
$htd$	$G_1 \times G_1 \times \mathbb{Z}_p \times ((G_1^3 \times G_2^3) \times G_1^{\ell+2} \times G_{\mathcal{T}})$	hidden user trapdoor
$HTD$	set of $(G_1 \times G_1 \times \mathbb{Z}_p \times ((G_1^3 \times G_2^3) \times G_1^{\ell+2} \times G_{\mathcal{T}}))$ elements	set of hidden user trapdoors
$n_{\text{PRF}}$	$\mathbb{N}$	maximum value of the PRF counter

**UC-Protocol  $\pi_{P4TC}$  (cont.) – Task DR Registration***DR input:* (register)

- (1) If a key pair  $(pk_{DR}, sk_{DR})$  has already been recorded, output  $\perp$  and abort.
- (2) Obtain CRS CRS from  $\mathcal{F}_{CRS}$ .
- (3) Run  $(pk_{DR}, sk_{DR}) \leftarrow \text{DRRegistration}(\text{CRS})$  (see Fig. 21).
- (4) Record  $(pk_{DR}, sk_{DR})$  internally and call  $\overline{\mathcal{G}}_{bb}$  with input (register,  $pk_{DR}$ ).

*DR output:*  $(pk_{DR})$ Fig. 20. The UC-protocol  $\pi_{P4TC}$  (cont. from Fig. 18)

DRRegistration(CRS)
parse $(gp, \mathcal{B}, \text{CRS}_{com}^1, \text{CRS}_{com}^2, \text{CRS}_{pok}) := \text{CRS}$ $(pk_{DR}, sk_{DR}) \leftarrow \text{E.Gen}(gp)$ return $(pk_{DR}, sk_{DR})$

Fig. 21. DR Registration Core Protocol

**UC-Protocol  $\pi_{P4TC}$  (cont.) – Task TSP Registration***TSP input:* (register,  $a_{\mathcal{T}}$ )

- (1) If a key pair  $(pk_{\mathcal{T}}, sk_{\mathcal{T}})$  has already been recorded, output  $\perp$  and abort.
- (2) Obtain CRS CRS from  $\mathcal{F}_{CRS}$ .
- (3) Run  $(pk_{\mathcal{T}}, sk_{\mathcal{T}}, \text{cert}_{\mathcal{T}}^{\mathcal{R}}) \leftarrow \text{TSPRegistration}(\text{CRS}, a_{\mathcal{T}})$  (see Fig. 23).
- (4) Record  $(pk_{\mathcal{T}}, sk_{\mathcal{T}})$  and  $(\text{cert}_{\mathcal{T}}^{\mathcal{R}})$  internally and call  $\overline{\mathcal{G}}_{bb}$  with input (register,  $pk_{\mathcal{T}}$ ).

*TSP output:*  $(pk_{\mathcal{T}})$ Fig. 22. The UC-protocol  $\pi_{P4TC}$  (cont. from Fig. 18)

TSPRegistration(CRS, $a_{\mathcal{T}}$ )
parse $(gp, \mathcal{B}, \text{CRS}_{com}^1, \text{CRS}_{com}^2, \text{CRS}_{pok}) := \text{CRS}$ $(pk_{\mathcal{T}}^{\mathcal{T}}, sk_{\mathcal{T}}^{\mathcal{T}}) \leftarrow \text{S.Gen}(gp)$ $(pk_{\mathcal{T}}^{\text{cert}}, sk_{\mathcal{T}}^{\text{cert}}) \leftarrow \text{S.Gen}(gp)$ $(pk_{\mathcal{T}}^{\mathcal{R}}, sk_{\mathcal{T}}^{\mathcal{R}}) \leftarrow \text{S.Gen}(gp)$ $(pk_{\mathcal{T}}, sk_{\mathcal{T}}) := ((pk_{\mathcal{T}}^{\mathcal{T}}, pk_{\mathcal{T}}^{\text{cert}}, pk_{\mathcal{T}}^{\mathcal{R}}), (sk_{\mathcal{T}}^{\mathcal{T}}, sk_{\mathcal{T}}^{\text{cert}}, sk_{\mathcal{T}}^{\mathcal{R}}))$ $\sigma_{\mathcal{T}}^{\text{cert}} \leftarrow \text{S.Sgn}(sk_{\mathcal{T}}^{\text{cert}}, (pk_{\mathcal{T}}^{\mathcal{R}}, a_{\mathcal{T}}))$ $\text{cert}_{\mathcal{T}}^{\mathcal{R}} := (pk_{\mathcal{T}}^{\mathcal{R}}, a_{\mathcal{T}}, \sigma_{\mathcal{T}}^{\text{cert}})$ return $(pk_{\mathcal{T}}, sk_{\mathcal{T}}, \text{cert}_{\mathcal{T}}^{\mathcal{R}})$

Fig. 23. TSP Registration Core Protocol

**UC-Protocol  $\pi_{P4TC}$  (cont.) – Task *RSU Registration****RSU input:* (register)

- (1) If a key pair  $(pk_{\mathcal{R}}, sk_{\mathcal{R}})$  has already been stored, output  $\perp$  and abort.
- (2) Obtain CRS CRS from  $\mathcal{F}_{CRS}$ .
- (3) Run  $(pk_{\mathcal{R}}, sk_{\mathcal{R}}) \leftarrow \text{RSURegistration}(\text{CRS})$  (see Fig. 25).
- (4) Store  $(pk_{\mathcal{R}}, sk_{\mathcal{R}})$  internally and call  $\overline{\mathcal{G}}_{bb}$  with input (register,  $pk_{\mathcal{R}}$ ).

*RSU output:*  $(pk_{\mathcal{R}})$ Fig. 24. The UC-protocol  $\pi_{P4TC}$  (cont. from Fig. 18)

RSURegistration(CRS)
$\text{parse } (gp, \mathcal{B}, \text{CRS}_{\text{com}}^1, \text{CRS}_{\text{com}}^2, \text{CRS}_{\text{pok}}) := \text{CRS}$ $(pk_{\mathcal{R}}, sk_{\mathcal{R}}) \leftarrow \text{S.Gen}(gp)$ $\text{return } (pk_{\mathcal{R}}, sk_{\mathcal{R}})$

Fig. 25. RSU Registration Core Protocol

**UC-Protocol  $\pi_{P4TC}$  (cont.) – Task *User Registration****User input:* (register)

- (1) If a key pair  $(pk_{\mathcal{U}}, sk_{\mathcal{U}})$  has already been stored, output  $\perp$  and abort.
- (2) Obtain CRS CRS from  $\mathcal{F}_{CRS}$ .
- (3) Run  $(pk_{\mathcal{U}}, sk_{\mathcal{U}}) \leftarrow \text{UserRegistration}(\text{CRS})$  (see Fig. 27).
- (4) Store  $(pk_{\mathcal{U}}, sk_{\mathcal{U}})$  internally and call  $\overline{\mathcal{G}}_{bb}$  with input (register,  $pk_{\mathcal{U}}$ ).

*User output:*  $(pk_{\mathcal{U}})$ Fig. 26. The UC-protocol  $\pi_{P4TC}$  (cont. from Fig. 18)

UserRegistration(CRS)
$\text{parse } (gp, \mathcal{B}, \text{CRS}_{\text{com}}^1, \text{CRS}_{\text{com}}^2, \text{CRS}_{\text{pok}}) := \text{CRS}$ $y \xleftarrow{R} \mathbb{Z}_p$ $(pk_{\mathcal{U}}, sk_{\mathcal{U}}) := (g_1^y, y)$ $\text{return } (pk_{\mathcal{U}}, sk_{\mathcal{U}})$

Fig. 27. User Registration Core Protocol

The DR computes a key pair  $(pk_{DR}, sk_{DR})$  which can be used to remove the unlinkability of user transactions in case of a dispute between the user and the TSP (see Figs. 20 and 21). The DR could be a (non-governmental) organization trusted by both, users to protect their privacy and the TSP to protect system security.

**UC-Protocol  $\pi_{P4TC}$  (cont.) – Task *RSU Certification***

*RSU input:* (certify)

*TSP input:* (certify,  $\mathbf{a}_R$ )

(1) At the RSU side:

- Load the internally recorded  $(pk_R, sk_R)$ .<sup>‡</sup>
- Receive  $pk_T$  from the bulletin-board  $\overline{\mathcal{G}}_{bb}$  for PID  $pid_T$ .<sup>‡</sup>

(2) At the TSP side:

- Load the internally recorded  $(pk_T, sk_T)$ .<sup>‡</sup>
- Receive  $pk_R$  from the bulletin-board  $\overline{\mathcal{G}}_{bb}$  for PID  $pid_R$ .<sup>‡</sup>
- Check that no mapping  $pk_R \mapsto \mathbf{a}'_R$  has been registered before, else output  $\perp$  and abort.

(3) Both sides: Run the code of RSUCertification between the RSU and the TSP (see Fig. 29)

$$((cert_R), (OK)) \leftarrow \text{RSUCertification} \langle \mathcal{R}(pk_T, pk_R), \mathcal{T}(pk_T, sk_T, pk_R, \mathbf{a}_R) \rangle.$$

(4) At the RSU side:

- Parse  $\mathbf{a}_R$  from  $cert_R$ .
- Record  $cert_R$  internally.

(5) At the TSP side:

- Record  $pk_R \mapsto \mathbf{a}_R$  internally.

*RSU output:* ( $\mathbf{a}_R$ )

*TSP output:* (OK)

<sup>‡</sup>If this does not exist, output  $\perp$  and abort.

Fig. 28. The UC-protocol  $\pi_{P4TC}$  (cont. from Fig. 18)

The TSP must also generate a key pair (see Figs. 22 and 23). Therefore, the TSP generates several signature key pairs  $(pk_T^T, sk_T^T)$ ,  $(pk_T^{\text{cert}}, sk_T^{\text{cert}})$ ,  $(pk_T^R, sk_T^R)$ , where  $sk_T^T$  is used in the Wallet Issuing task to sign the TSP commitment  $c_T$  and the user attributes  $\mathbf{a}_U$ ,  $sk_T^{\text{cert}}$  is used to sign RSU public keys in the RSU Certification task and  $sk_T^R$  is used in the Wallet Issuing task to sign the RSU commitment  $c_R$  and the serial number  $s$  in place of an RSU. The TSP also generates a certificate  $cert_T^R$  for its own key  $pk_T^R$ .

Each RSU must generate a key pair as well (see Figs. 24 and 25). For that purpose, each RSU generates a signature key pair  $(pk_R, sk_R)$  that is used in the Debt Accumulation task to sign the RSU commitment  $c_R$ .

Each user also has to generate a key pair (see Figs. 26 and 27). The public key  $pk_U$  will be used to identify the user in the system and is assumed to be bound to a physical ID such as a passport number, social security number, etc. Of course, for this purpose the public key needs to be unique. We assume that ensuring the uniqueness of user public keys as well as verifying and binding a physical ID to them is done “out-of-band” before participating in the Wallet Issuing task. A simple way to realize the latter could be to make use of external trusted certification authorities.

#### D.5 RSU Certification

The RSU Certification task is executed between  $\mathcal{R}$  and  $\mathcal{T}$  when a new RSU is deployed into the field. The task is presented in two parts: a wrapper protocol  $\pi_{P4TC}$  (see Fig. 28) and a core protocol RSUCertification (see Fig. 29). The wrapper protocol interacts with other UC entities, pre-processes the input and post-processes the output.

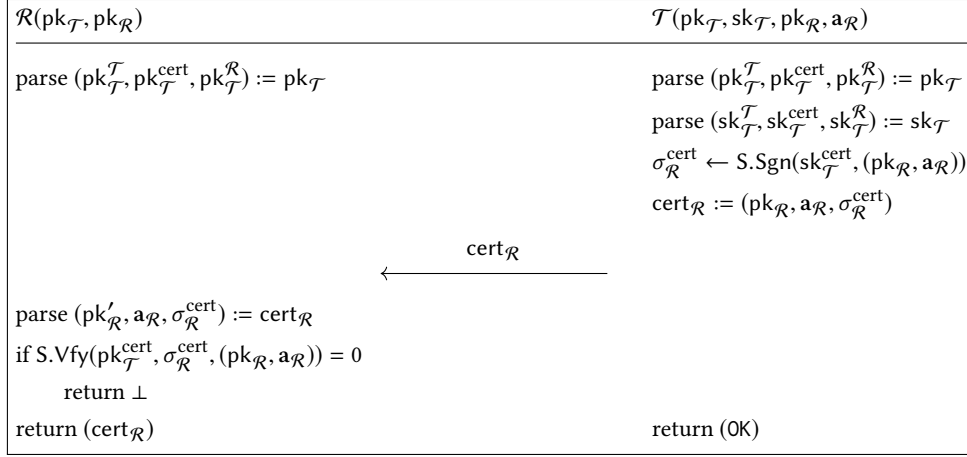


Fig. 29. RSU Certification Core Protocol

The wrapper protocol internally invokes the core protocol

$$((\text{cert}_{\mathcal{R}}), (\text{OK})) \leftarrow \text{RSUCertification} \langle \mathcal{R}(\text{pk}_{\mathcal{T}}, \text{pk}_{\mathcal{R}}), \mathcal{T}(\text{pk}_{\mathcal{T}}, \text{sk}_{\mathcal{T}}, \text{pk}_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}}) \rangle.$$

In the core protocol, the TSP certifies the validity of the RSU public key and stores the certificate on the RSU.

Note that the public key of an RSU  $\text{pk}_{\mathcal{R}}$  and the associated certificate  $\text{cert}_{\mathcal{R}}$  has to be refreshed from time to time. For the ease of presentation we assume that the same RSU (identified by its PID  $\text{pid}_{\mathcal{R}}$ ) can only be registered once. In other words, if the (physically identical) RSU is removed from the field, goes to maintenance and is re-deployed to the field, we consider this RSU a “new” RSU.

## D.6 Wallet Issuing

The Wallet Issuing task is executed between  $\mathcal{U}$  and  $\mathcal{T}$ . It is executed at the beginning of each billing period to generate a fresh wallet for the user. The task is presented in two parts: a wrapper protocol  $\pi_{\text{P4TC}}$  (see Fig. 30) and a core protocol `WalletIssuing` (see Fig. 31). The wrapper protocol interacts with other UC entities, pre-processes the input, post-processes the output and checks the validity of the created wallet by executing the `WalletVerification` algorithm (see Fig. 42) after the core protocol. The wrapper protocol internally invokes the core protocol

$$((\tau), (s, \text{htd})) \leftarrow \text{WalletIssuing} \langle \mathcal{U}(\text{pk}_{\text{DR}}, \text{pk}_{\mathcal{U}}, \text{sk}_{\mathcal{U}}), \mathcal{T}(\text{pk}_{\text{DR}}, \text{sk}_{\mathcal{T}}, \mathbf{a}_{\mathcal{U}}, \text{cert}_{\mathcal{T}}^{\mathcal{R}}, \text{bl}_{\mathcal{T}}) \rangle.$$

The joint input of the core protocol is the public key of the DR  $\text{pk}_{\text{DR}}$ . The user additionally obtains its public and secret key pair  $(\text{pk}_{\mathcal{U}}, \text{sk}_{\mathcal{U}})$ . The TSP also gets his own secret key  $\text{sk}_{\mathcal{T}}$ , the attribute vector  $\mathbf{a}_{\mathcal{U}}$  for the user, its own certificate  $\text{cert}_{\mathcal{T}}^{\mathcal{R}}$  and the TSP blacklist  $\text{bl}_{\mathcal{T}}$  as input.

The protocol fulfills four objectives:

- (1) Jointly computing a fresh and random wallet ID for the user that is only known to the user.
- (2) Storing this wallet ID in a secret form at the TSP such that it can only be recovered by the DR in the case that the user conducts a fraud.
- (3) Jointly computing a fresh and random serial number for this transaction.
- (4) Creating a new wallet for the user.

For the first objective, both parties randomly choose shares of the wallet ID  $\lambda'$  and  $\lambda''$ , respectively, that together form the wallet ID  $\lambda := \lambda' + \lambda''$ . To this end, the parties engage in the first two message of a Blum coin

**UC-Protocol  $\pi_{P4TC}$  (cont.) – Task *Wallet Issuing***

*User input:* (issue)

*TSP input:* (issue,  $\mathbf{a}_U, bl_T$ )

(1) At the user side:

- Load the internally recorded  $(pk_U, sk_U)$ .<sup>‡</sup>
- Receive  $pk_T$  from the bulletin-board  $\overline{\mathcal{G}}_{bb}$  for PID  $pid_T$ .<sup>‡</sup>
- Receive  $pk_{DR}$  from the bulletin-board  $\overline{\mathcal{G}}_{bb}$  for PID  $pid_{DR}$ .<sup>‡</sup>

(2) At the TSP side:

- Load the internally recorded  $(pk_T, sk_T)$ .<sup>‡</sup>
- Load the internally recorded  $cert_T^R$ .<sup>‡</sup>
- Receive  $pk_{DR}$  from the bulletin-board  $\overline{\mathcal{G}}_{bb}$  for PID  $pid_{DR}$ .<sup>‡</sup>

(3) Both sides: Run the code of *WalletIssuing* between the user and the TSP (see Fig. 31)

$$((\tau), (s, htd)) \leftarrow \text{WalletIssuing} \left( \mathcal{U}(pk_{DR}, pk_U, sk_U), \mathcal{T}(pk_{DR}, sk_T, \mathbf{a}_U, cert_T^R, bl_T) \right).$$

(4) At the user side:

- Run the code of *WalletVerification*( $pk_T, pk_U, \tau$ ) (see Fig. 42).
- If *WalletVerification* returns 0, output  $\perp$  and abort.
- Record  $\tau$  internally.
- Parse  $s$  and  $\mathbf{a}_U$  from  $\tau$ .

(5) At the TSP side:

- Insert  $htd$  into *HTD*.

*User output:* ( $s, \mathbf{a}_U$ )

*TSP output:* ( $s$ )

<sup>‡</sup>If this does not exist, output  $\perp$  and abort.

Fig. 30. The UC-protocol  $\pi_{P4TC}$  (cont. from Fig. 18)

toss and omit the last message. This way, the wallet ID  $\lambda$  is fixed and known by the user, but remains secret to the TSP.

The second objective requires some sort of key-escrow mechanism. Ideally, the user would simply encrypt the wallet ID  $\lambda$  under the public key  $pk_{DR}$  of the DR and prove to the TSP that the encrypted value is consistent to the committed shares in zero-knowledge. Unfortunately, we are unaware of a CCA-secure encryption scheme whose message space equals the key space of our PRF underlying the wallet (i.e.,  $\mathbb{Z}_p$ ) and that is compatible to the GS-NIZK proof system (i.e., is algebraic). Moreover, it is impossible to recover  $\lambda$  from  $g_1^\lambda$  due to the hardness of the DLOG problem in  $G_1$ . Therefore, the user splits its share  $\lambda'$  into  $\lambda'_0, \dots, \lambda'_{\ell-1} \in \{0, \dots, \mathcal{B} - 1\}$  s.t.  $\lambda' = \sum_{i=0}^{\ell-1} \lambda'_i \cdot \mathcal{B}^i$  for some base  $\mathcal{B}$ . The base  $\mathcal{B}$  is chosen in a way that it is feasible for the DR to recover  $\lambda'_i$  from  $g_1^{\lambda'_i}$  in a reasonable amount of time (e.g.,  $\mathcal{B} = 2^{32}$ ). Then the user encrypts a vector of all  $g_1^{\lambda'_i}$  together with the TSP's share  $g_1^{\lambda''}$  and its own public key  $pk_U$  under the public key  $pk_{DR}$  of the DR and sends the ciphertext  $e^*$  to the TSP. The TSP's share  $g_1^{\lambda''}$  and the public key  $pk_U$  are embedded into the ciphertext in order to bind it to the user and to rule out malleability attacks. The TSP creates the hidden user trapdoor as  $htd := (pk_U, s, \lambda'', e^*)$ . In the case that the user commits a fraud, the TSP sends the  $htd$  for each wallet of the fraudulent user to the DR and the DR recovers the wallet ID  $\lambda$  for each wallet. For more details see the User Blacklisting task in Fig. 41.

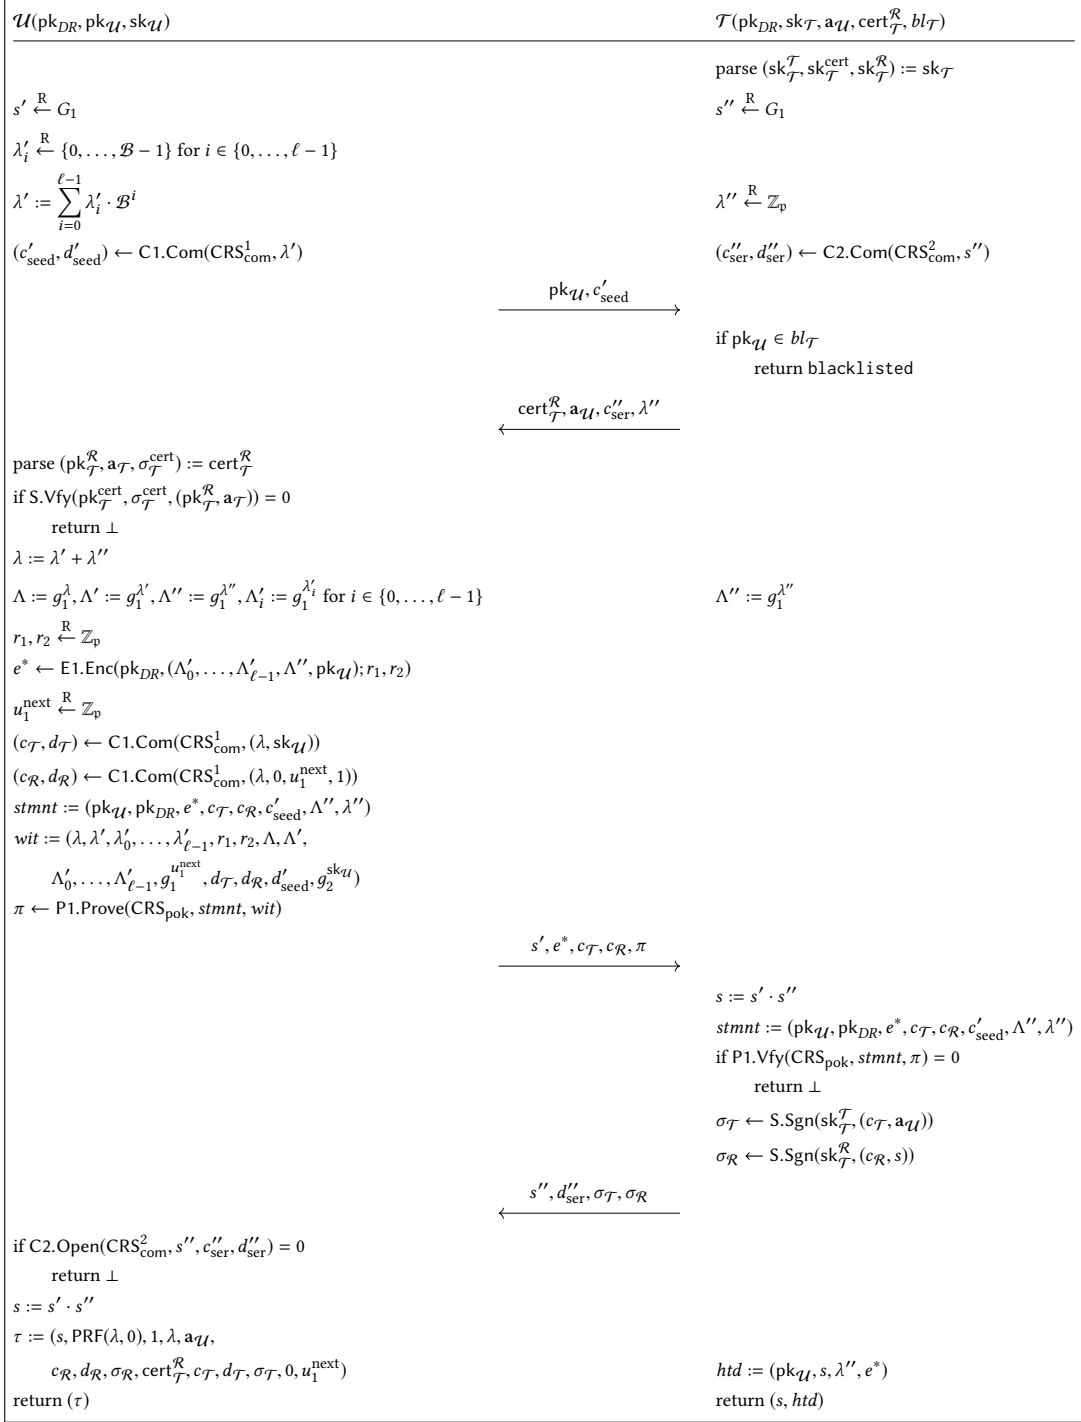


Fig. 31. Wallet Issuing Core Protocol

The goal of the third objective is to create a truly random serial number  $s \in G_1$  for this transaction. To ensure that the serial number is indeed random (and not maliciously chosen by either party), the user and the TSP engage in another and complete Blum coin toss.

For the last objective, the user generates the TSP and RSU commitments, i.e., the fixed and the updatable part of the wallet. He commits to the wallet ID  $\lambda$  and his secret user key  $sk_U$  for the TSP commitment  $c_{\mathcal{T}}$ . For the RSU commitment  $c_{\mathcal{R}}$ , he commits to the wallet ID  $\lambda$ , the balance  $b := 0$ , a fresh double-spending mask  $u_1^{\text{next}}$  and the PRF counter  $x^{\text{next}} := 1$ . He then computes a proof showing that these commitments are formed correctly. The proof also shows that the encryption  $e^*$  has been honestly created and that each  $\lambda'_i$  is smaller than  $\mathcal{B}$ . More precisely, P1 is used to compute a proof  $\pi$  for a statement *stmt* from the language  $L_{\text{gp}}^{(1)}$  defined by

$$L_{\text{gp}}^{(1)} := \left\{ \begin{array}{l} \left( \begin{array}{l} \text{pk}_U \\ \text{pk}_{DR} \\ e^* \\ c_{\mathcal{T}} \\ c_{\mathcal{R}} \\ c'_{\text{seed}} \\ \Lambda'' \\ \lambda'' \end{array} \right)^{\top} \quad \left. \begin{array}{l} \exists \lambda, \lambda', \lambda'_0, \dots, \lambda'_{\ell-1}, r_1, r_2 \in \mathbb{Z}_p; \\ \Lambda, \Lambda', \Lambda'_0, \dots, \Lambda'_{\ell-1}, U_1^{\text{next}}, d_{\mathcal{T}}, d_{\mathcal{R}}, d'_{\text{seed}} \in G_1; \\ \text{SK}_U \in G_2 : \\ \\ e(\text{pk}_U, g_2) = e(g_1, \text{SK}_U) \\ \text{C1.Open}(\text{CRS}_{\text{com}}^1, (\Lambda, \text{pk}_U), c_{\mathcal{T}}, d_{\mathcal{T}}) = 1 \\ \text{C1.Open}(\text{CRS}_{\text{com}}^1, (\Lambda, 1, U_1^{\text{next}}, g_1), c_{\mathcal{R}}, d_{\mathcal{R}}) = 1 \\ \text{C1.Open}(\text{CRS}_{\text{com}}^1, \Lambda', c'_{\text{seed}}, d'_{\text{seed}}) = 1 \\ e^* = \text{E1.Enc}(\text{pk}_{DR}, (\Lambda'_0, \dots, \Lambda'_{\ell-1}, \Lambda'', \text{pk}_U); r_1, r_2) \\ \lambda = \lambda' + \lambda'' \\ \Lambda = g_1^{\lambda}, \Lambda' = g_1^{\lambda'} \\ \lambda' = \sum_{i=0}^{\ell-1} \lambda'_i \cdot \mathcal{B}^i \\ \forall i \in \{0, \dots, \ell-1\} : \\ \quad \lambda'_i \in \{0, \dots, \mathcal{B}-1\} \\ \quad \Lambda'_i = g_1^{\lambda'_i} \end{array} \right\} \quad (6)$$

Note that the first equation in Eq. (6) actually proves the knowledge of  $g_2^{\text{sk}_U}$  (rather than  $\text{sk}_U$  itself).<sup>24</sup> However, computing  $g_2^{\text{sk}_U}$  without knowing  $\text{sk}_U$  (only given  $\text{pk}_U$ ) is assumed to be a hard problem (Co-CDH).

In temporal order, the protocol proceeds as follows: In the first message (from user to TSP) the user sends its own public key  $\text{pk}_U$  and starts the Blum coin toss for the wallet ID by sending  $c'_{\text{seed}}$ . The TSP checks if the user's public key is contained in the TSP blacklist  $bl_{\mathcal{T}}$  and potentially aborts. If not, the TSP replies with the second message of the Blum coin toss for the wallet ID by sending its own share  $\lambda''$  and starts the other Blum coin toss for the serial number by sending  $c'_{\text{ser}}$ . Moreover, the TSP sends its own certificate  $\text{cert}_{\mathcal{T}}^{\mathcal{R}}$  and the attributes  $\mathbf{a}_U$  the user is supposed to incorporate into its wallet. This completes the first Blum coin toss for the wallet ID. At this point the user knows all information to create the two commitments  $c_{\mathcal{T}}$  and  $c_{\mathcal{R}}$  of the wallet, the hidden user trapdoor  $e^*$  for the key-escrow mechanism and to create a proof  $\pi$  that everything is consistent. In the third message (again from user to TSP) the user sends all these elements ( $e^*$ ,  $c_{\mathcal{T}}$ ,  $c_{\mathcal{R}}$  and  $\pi$ ) to the TSP together with its share  $s'$  of the serial number as the second message of the second Blum coin toss. If the proof  $\pi$  verifies, the TSP creates two signatures:  $\sigma_{\mathcal{T}}$  on  $c_{\mathcal{T}}$  together with  $\mathbf{a}_U$  for the fixed part of the wallet, and  $\sigma_{\mathcal{R}}$  on  $c_{\mathcal{R}}$  together with  $s$  on the updatable part of the wallet. Please note that at this point the serial number  $s := s' \cdot s''$  is fixed and known to the TSP. In the fourth message (from TSP to user) the TSP sends the signatures  $\sigma_{\mathcal{T}}$  and  $\sigma_{\mathcal{R}}$  to the user and completes the Blum coin toss for the serial number by sending its share  $s''$  together with the opening information  $d''_{\text{ser}}$ . At this point the user assembles all part to obtain a fully functional wallet

$$\tau := (s, \phi := \text{PRF}(\lambda, 0), x^{\text{next}} := 1, \lambda := \lambda' + \lambda'', \mathbf{a}_U, c_{\mathcal{R}}, d_{\mathcal{R}}, \sigma_{\mathcal{R}}, \text{cert}_{\mathcal{T}}^{\mathcal{R}}, c_{\mathcal{T}}, d_{\mathcal{T}}, b := 0, u_1^{\text{next}}).$$

<sup>24</sup>Note that proving a statement  $\exists \text{sk}_U \in \mathbb{Z}_p : \text{pk}_U = g_1^{\text{sk}_U}$  instead would not help as we can only extract  $g_1^{\text{sk}_U}$  from the proof.



**UC-Protocol  $\pi_{P4TC}$  (cont.) – Task *Debt Accumulation****User input:*  $(\text{pay\_toll}, s^{\text{prev}})$ *RSU input:*  $(\text{pay\_toll}, bl_{\mathcal{R}})$ 

- (1) At the user side:
  - Load the internally recorded  $(pk_{\mathcal{U}}, sk_{\mathcal{U}})$ .<sup>‡</sup>
  - Receive  $pk_{\mathcal{T}}$  from the bulletin-board  $\overline{\mathcal{G}}_{\text{bb}}$  for PID  $pid_{\mathcal{T}}$ .<sup>‡</sup>
  - Load the internally recorded token  $\tau^{\text{prev}}$  for serial number  $s^{\text{prev}}$ .<sup>‡</sup>
- (2) At the RSU side:
  - Load the internally recorded  $(pk_{\mathcal{R}}, sk_{\mathcal{R}})$ .<sup>‡</sup>
  - Load the internally recorded  $\text{cert}_{\mathcal{R}}$ .<sup>‡</sup>
  - Receive  $pk_{\mathcal{T}}$  from the bulletin-board  $\overline{\mathcal{G}}_{\text{bb}}$  for PID  $pid_{\mathcal{T}}$ .<sup>‡</sup>
- (3) Both sides: Run the code of DebtAccumulation between the user and the RSU (see Fig. 33)

$$\left( \begin{array}{c} (\tau, \omega_{\mathcal{U}}^{\text{pp}}), \\ (\mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}}, \omega^{\text{dsp}}, \omega^{\text{bl}}, \omega_{\mathcal{R}}^{\text{pp}}) \end{array} \right) \leftarrow \text{DebtAccumulation} \left\langle \begin{array}{c} \mathcal{U}(pk_{\mathcal{T}}, pk_{\mathcal{U}}, sk_{\mathcal{U}}, \tau^{\text{prev}}), \\ \mathcal{R}^{\text{O}_{\text{pricing}}(\cdot, \cdot)}(pk_{\mathcal{T}}, \text{cert}_{\mathcal{R}}, sk_{\mathcal{R}}, bl_{\mathcal{R}}) \end{array} \right\rangle$$

and forward calls to the pricing oracle  $\text{O}_{\text{pricing}}$  of the form  $(\mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}})$  to the adversary and pass the result  $p$  back.

- (4) At the user side:
  - Run the code of WalletVerification( $pk_{\mathcal{T}}, pk_{\mathcal{U}}, \tau$ ) (see Fig. 42).
  - If WalletVerification returns 0, output  $\perp$  and abort.
  - Record  $\tau$  and  $\omega_{\mathcal{U}}^{\text{pp}}$  internally.
  - Parse  $s$ ,  $\text{cert}_{\mathcal{R}}$ ,  $p$  and  $b$  from  $\tau$ .
  - Parse  $\mathbf{a}_{\mathcal{R}}$  from  $\text{cert}_{\mathcal{R}}$ .
- (5) At the RSU side:
  - Record  $\omega^{\text{dsp}}, \omega^{\text{bl}}$  and  $\omega_{\mathcal{R}}^{\text{pp}}$  internally.
  - Parse  $s$  from  $\omega_{\mathcal{R}}^{\text{pp}}$ .
  - Parse  $\phi$  from  $\omega^{\text{bl}}$ .

*User output:*  $(s, \mathbf{a}_{\mathcal{R}}, p, b)$ *RSU output:*  $(s, \phi, \mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}})$ 

<sup>‡</sup>If this does not exist, output  $\perp$  and abort.

Fig. 32. The UC-protocol  $\pi_{P4TC}$  (cont. from Fig. 18)

At the end of the protocol, the user returns the wallet  $\tau$ . The TSP returns the serial number of this transaction and the hidden user trapdoor  $htd$ .

**D.7 Debt Accumulation**

When a driving car passes an RSU, the Debt Accumulation task is executed between  $\mathcal{U}$  and  $\mathcal{R}$ . The task is presented in two parts: a wrapper protocol  $\pi_{P4TC}$  (see Fig. 32) and a core protocol DebtAccumulation (see Fig. 33).

The wrapper protocol interacts with other UC entities, pre-processes the input, post-processes the output and lets the user execute the WalletVerification algorithm (see Fig. 42) after the core protocol has terminated. The

$\mathcal{U}(\text{pk}_{\mathcal{T}}, \text{pk}_{\mathcal{U}}, \text{sk}_{\mathcal{U}}, \tau^{\text{prev}})$	$\mathcal{R}^{\text{Opricing}(\cdot, \cdot, \cdot)}(\text{pk}_{\mathcal{T}}, \text{cert}_{\mathcal{R}}, \text{sk}_{\mathcal{R}}, \text{bl}_{\mathcal{R}})$
$\text{parse}(\text{pk}_{\mathcal{T}}^{\mathcal{T}}, \text{pk}_{\mathcal{T}}^{\text{cert}}, \text{pk}_{\mathcal{T}}^{\mathcal{R}}) := \text{pk}_{\mathcal{T}}$	$\text{parse}(\text{pk}_{\mathcal{T}}^{\mathcal{T}}, \text{pk}_{\mathcal{T}}^{\text{cert}}, \text{pk}_{\mathcal{T}}^{\mathcal{R}}) := \text{pk}_{\mathcal{T}}$
$\text{parse}(s^{\text{prev}}, \phi^{\text{prev}}, x, \lambda, \mathbf{a}_{\mathcal{U}}, c_{\mathcal{R}}^{\text{prev}}, d_{\mathcal{R}}^{\text{prev}}, \sigma_{\mathcal{R}}^{\text{prev}}, \text{cert}_{\mathcal{R}}^{\text{prev}},$ $c_{\mathcal{T}}, d_{\mathcal{T}}, \sigma_{\mathcal{T}}, b^{\text{prev}}, u_1) := \tau^{\text{prev}}$	
$\text{parse}(\text{pk}_{\mathcal{R}}^{\text{prev}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}}, \sigma_{\mathcal{R}}^{\text{certprev}}) := \text{cert}_{\mathcal{R}}^{\text{prev}}$	$\text{parse}(\text{pk}_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}}, \sigma_{\mathcal{R}}^{\text{cert}}) := \text{cert}_{\mathcal{R}}$
$\phi := \text{PRF}(\lambda, x)$	
$s' \xleftarrow{\mathcal{R}} G_1$	$s'' \xleftarrow{\mathcal{R}} G_1$
$u_1^{\text{next}} \xleftarrow{\mathcal{R}} \mathbb{Z}_p$	$u_2 \xleftarrow{\mathcal{R}} \mathbb{Z}_p$
$(c'_{\mathcal{R}}, d'_{\mathcal{R}}) \leftarrow \text{C1.Com}(\text{CRS}_{\text{com}}^1, (\lambda, b^{\text{prev}}, u_1^{\text{next}}, x))$	$(c''_{\text{ser}}, d''_{\text{ser}}) \leftarrow \text{C2.Com}(\text{CRS}_{\text{com}}^2, s'')$
	$\xleftarrow{u_2, c''_{\text{ser}}, \text{cert}_{\mathcal{R}}}$
$\text{parse}(\text{pk}_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}}, \sigma_{\mathcal{R}}^{\text{cert}}) := \text{cert}_{\mathcal{R}}$	
$\text{if } S.\text{Vfy}(\text{pk}_{\mathcal{T}}^{\text{cert}}, \sigma_{\mathcal{R}}^{\text{cert}}, (\text{pk}_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}})) = 0$	
$\quad \text{return } \perp$	
$t := \text{sk}_{\mathcal{U}} u_2 + u_1 \text{ mod } p$	
$(c_{\text{hid}}, d_{\text{hid}}) \leftarrow \text{C1.Com}(\text{CRS}_{\text{com}}^1, \text{sk}_{\mathcal{U}})$	
$\text{stmnt} := (\text{pk}_{\mathcal{T}}^{\mathcal{T}}, \text{pk}_{\mathcal{T}}^{\text{cert}}, \phi, \mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}}, c_{\text{hid}}, c'_{\mathcal{R}}, t, u_2)$	
$\text{wit} := (x, \lambda, \text{sk}_{\mathcal{U}}, u_1, s^{\text{prev}}, \phi^{\text{prev}}, g_1^x, g_1^\lambda, \text{pk}_{\mathcal{U}}, g_1^{b^{\text{prev}}}, g_1^{u_1}, g_1^{u_1^{\text{next}}},$ $d_{\text{hid}}, d_{\mathcal{R}}^{\text{prev}}, d'_{\mathcal{R}}, d_{\mathcal{T}}, \text{pk}_{\mathcal{R}}^{\text{prev}}, c_{\mathcal{R}}^{\text{prev}}, c_{\mathcal{T}}, \sigma_{\mathcal{R}}^{\text{prev}}, \sigma_{\mathcal{R}}^{\text{certprev}}, \sigma_{\mathcal{T}})$	
$\pi \leftarrow \text{P2.Prove}(\text{CRS}_{\text{pok}}, \text{stmnt}, \text{wit})$	
	$\xrightarrow{s', \pi, \phi, \mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}}, c_{\text{hid}}, c'_{\mathcal{R}}, t}$
	$\text{stmnt} := (\text{pk}_{\mathcal{T}}^{\mathcal{T}}, \text{pk}_{\mathcal{T}}^{\text{cert}}, \phi, \mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}}, c_{\text{hid}}, c'_{\mathcal{R}}, t, u_2)$
	$\text{if } \text{P2.Vfy}(\text{CRS}_{\text{pok}}, \text{stmnt}, \pi) = 0$
	$\quad \text{return } \perp$
	$\text{if } \phi \in \text{bl}_{\mathcal{R}}$
	$\quad \text{return blacklisted}$
	$\quad // \text{ obtains the price from the pricing oracle}$
	$\quad // \text{ based on } \mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}} \text{ and possibly}$
	$\quad // \text{ other public-verifiable, environmental information}$
	$p \leftarrow \text{Opricing}(\mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}})$
	$s := s' \cdot s''$
	$(c''_{\mathcal{R}}, d''_{\mathcal{R}}) \leftarrow \text{C1.Com}(\text{CRS}_{\text{com}}^1, (0, p, 0, 1))$
	$c_{\mathcal{R}} := c'_{\mathcal{R}} \cdot c''_{\mathcal{R}}$
	$\sigma_{\mathcal{R}} \leftarrow \text{S.Sgn}(\text{sk}_{\mathcal{R}}, (c_{\mathcal{R}}, s))$
	$\xleftarrow{s'', d''_{\text{ser}}, c_{\mathcal{R}}, d''_{\mathcal{R}}, \sigma_{\mathcal{R}}, p}$
$\text{if } \text{C2.Open}(\text{CRS}_{\text{com}}^2, s'', c''_{\text{ser}}, d''_{\text{ser}}) = 0$	
$\quad \text{return } \perp$	
$s := s' \cdot s''$	
$d_{\mathcal{R}} := d'_{\mathcal{R}} \cdot d''_{\mathcal{R}}$	
$b := b^{\text{prev}} + p$	
$x^{\text{next}} := x + 1$	
$\tau := (s, \phi, x^{\text{next}}, \lambda, \mathbf{a}_{\mathcal{U}}, c_{\mathcal{R}}, d_{\mathcal{R}}, \sigma_{\mathcal{R}}, \text{cert}_{\mathcal{R}}, c_{\mathcal{T}}, d_{\mathcal{T}}, \sigma_{\mathcal{T}}, b, u_1^{\text{next}})$	
$\omega_{\mathcal{U}}^{\text{pp}} := (s, c_{\text{hid}}, d_{\text{hid}})$	$\omega^{\text{dsp}} := (\phi, t, u_2)$
$\text{return } (\tau, \omega_{\mathcal{U}}^{\text{pp}})$	$\omega^{\text{bl}} := (\phi, p)$
	$\omega_{\mathcal{R}}^{\text{pp}} := (s, c_{\text{hid}})$
	$\text{return } (\mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}}, \omega^{\text{dsp}}, \omega^{\text{bl}}, \omega_{\mathcal{R}}^{\text{pp}})$

Fig. 33. Debt Accumulation Core Protocol

wrapper protocol internally invokes the core protocol

$$\left( \begin{array}{c} \left( \tau, \omega_{\mathcal{U}}^{\text{pp}} \right), \\ \left( \mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}}, \omega^{\text{dsp}}, \omega^{\text{bl}}, \omega_{\mathcal{R}}^{\text{pp}} \right) \end{array} \right) \leftarrow \text{DebtAccumulation} \left\langle \begin{array}{c} \mathcal{U}(\text{pk}_{\mathcal{T}}, \text{pk}_{\mathcal{U}}, \text{sk}_{\mathcal{U}}, \tau^{\text{prev}}), \\ \mathcal{R}^{\text{Opricing}(\cdot, \cdot, \cdot)}(\text{pk}_{\mathcal{T}}, \text{cert}_{\mathcal{R}}, \text{sk}_{\mathcal{R}}, \text{bl}_{\mathcal{R}}) \end{array} \right\rangle.$$

The user gets his public and private key, the public key of the TSP and his current wallet

$$\tau^{\text{prev}} := (s^{\text{prev}}, \phi^{\text{prev}}, x, \lambda, \mathbf{a}_{\mathcal{U}}, c_{\mathcal{R}}^{\text{prev}}, d_{\mathcal{R}}^{\text{prev}}, \sigma_{\mathcal{R}}^{\text{prev}}, \text{cert}_{\mathcal{R}}^{\text{prev}}, c_{\mathcal{T}}, d_{\mathcal{T}}, \sigma_{\mathcal{T}}, b^{\text{prev}}, u_1)$$

as input. The RSU gets its own secret key and certificate, the public key of the TSP and the RSU blacklist  $\text{bl}_{\mathcal{R}}$  as input. It has also access to a pricing oracle  $\text{O}_{\text{pricing}}(\cdot, \cdot, \cdot)$ , which helps it to determine the price the user has to pay.

Analogous to WalletIssuing, the RSU and the user utilize a Blum coin toss to jointly compute a fresh and random serial number  $s$  for this transaction. The detailed description of this coin toss is therefore omitted in the following protocol description.

The RSU starts the protocol by sending its certificate  $\text{cert}_{\mathcal{R}}$  and a fresh double-spending randomness  $u_2$  to the user. The user checks the validity of the certificate and uses the randomness to calculate the double-spending tag  $t := \text{sk}_{\mathcal{U}} \cdot u_2 + u_1 \bmod \mathfrak{p}$ . He then calculates the fraud detection ID for the current transaction as  $\phi := \text{PRF}(\lambda, x)$ . The user then proceeds by preparing the updated wallet. Therefore, he first chooses a fresh double-spending mask  $u_1^{\text{next}}$  and executes  $(c'_{\mathcal{R}}, d'_{\mathcal{R}}) \leftarrow \text{C1.Com}(\text{CRS}_{\text{com}}^1, (\lambda, b^{\text{prev}}, u_1^{\text{next}}, x))$  to commit to his wallet ID, the current balance, the fresh double-spending mask and the current counter. He also executes  $(c_{\text{hid}}, d_{\text{hid}}) \leftarrow \text{C1.Com}(\text{CRS}_{\text{com}}^1, \text{sk}_{\mathcal{U}})$  to create a fresh commitment on his secret user key. This commitment can be used at a later point in the Prove Participation task (cf. Figs. 36 and 37) to prove to the TSP that the user behaved honestly in this transaction.

The user continues by using P2 to compute a proof  $\pi$  for a statement  $\text{stmt}$  from the language  $L_{\text{gp}}^{(2)}$  defined by

$$L_{\text{gp}}^{(2)} := \left\{ \begin{array}{l} \left( \begin{array}{c} \text{pk}_{\mathcal{T}}^{\mathcal{T}} \\ \text{pk}_{\mathcal{T}}^{\text{cert}} \\ \phi \\ \mathbf{a}_{\mathcal{U}} \\ \mathbf{a}_{\mathcal{R}}^{\text{prev}} \\ c_{\text{hid}} \\ c'_{\mathcal{R}} \\ t \\ u_2 \end{array} \right)^{\top} \left\{ \begin{array}{l} \exists x, \lambda, \text{sk}_{\mathcal{U}}, u_1 \in \mathbb{Z}_{\mathfrak{p}}; s^{\text{prev}}, \phi^{\text{prev}}, X, \Lambda, \text{pk}_{\mathcal{U}}, \\ B^{\text{prev}}, U_1, U_1^{\text{next}}, d_{\text{hid}}, d_{\mathcal{R}}^{\text{prev}}, d'_{\mathcal{R}}, d_{\mathcal{T}} \in G_1; \\ \text{pk}_{\mathcal{R}}^{\text{prev}} \in G_1^3, c_{\mathcal{R}}^{\text{prev}}, c_{\mathcal{T}} \in G_2; \sigma_{\mathcal{R}}^{\text{prev}}, \sigma_{\mathcal{R}}^{\text{cert}^{\text{prev}}}, \\ \sigma_{\mathcal{T}} \in G_2^2 \times G_1; \\ \\ \text{C1.Open}(\text{CRS}_{\text{com}}^1, (\Lambda, \text{pk}_{\mathcal{U}}), c_{\mathcal{T}}, d_{\mathcal{T}}) = 1 \\ \text{C1.Open}(\text{CRS}_{\text{com}}^1, \text{pk}_{\mathcal{U}}, c_{\text{hid}}, d_{\text{hid}}) = 1 \\ \text{C1.Open}(\text{CRS}_{\text{com}}^1, (\Lambda, B^{\text{prev}}, U_1, X), c_{\mathcal{R}}^{\text{prev}}, d_{\mathcal{R}}^{\text{prev}}) = 1 \\ \text{C1.Open}(\text{CRS}_{\text{com}}^1, (\Lambda, B^{\text{prev}}, U_1^{\text{next}}, X), c'_{\mathcal{R}}, d'_{\mathcal{R}}) = 1 \\ \text{S.Vfy}(\text{pk}_{\mathcal{T}}^{\mathcal{T}}, \sigma_{\mathcal{T}}, (c_{\mathcal{T}}, \mathbf{a}_{\mathcal{U}})) = 1 \\ \text{S.Vfy}(\text{pk}_{\mathcal{R}}^{\text{prev}}, \sigma_{\mathcal{R}}^{\text{prev}}, (c_{\mathcal{R}}^{\text{prev}}, s^{\text{prev}})) = 1 \\ \text{S.Vfy}(\text{pk}_{\mathcal{T}}^{\text{cert}}, \sigma_{\mathcal{R}}^{\text{cert}^{\text{prev}}}, (\text{pk}_{\mathcal{R}}^{\text{prev}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}})) = 1 \\ \phi^{\text{prev}} = \text{PRF}(\lambda, x - 1), \phi = \text{PRF}(\lambda, x), \\ t = \text{sk}_{\mathcal{U}} u_2 + u_1 \\ \text{pk}_{\mathcal{U}} = g_1^{\text{sk}_{\mathcal{U}}}, U_1 = g_1^{u_1}, X = g_1^x, \Lambda = g_1^\lambda \end{array} \right. \end{array} \right. \quad (7)$$

This proof essentially shows that the wallet  $\tau$  is valid, i.e., that the commitments  $c_{\mathcal{T}}$  and  $c_{\mathcal{R}}^{\text{prev}}$  are valid, bound to the user and have valid signatures, that the certificate  $\text{cert}_{\mathcal{R}}^{\text{prev}}$  from the previous RSU is valid and that the fraud detection ID  $\phi^{\text{prev}}$  from the last transaction has been computed correctly. It also shows that  $c_{\text{hid}}$  is valid and contains the secret user key, that  $c_{\mathcal{R}}^{\text{prev}}$  and  $c'_{\mathcal{R}}$  contain the same values (except for the double-spending mask) and that the fraud detection ID  $\phi$  and the double-spending tag  $t$  are computed correctly.

The user then sends  $(\pi, \phi, \mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}}, c_{\text{hid}}, c'_{\mathcal{R}}, t)$  to the RSU, who first checks whether the proof  $\pi$  verifies and the fraud detection ID  $\phi$  is on the RSU blacklist  $\text{bl}_{\mathcal{R}}$  or not. If one of the checks fails, the RSU aborts the communication with the user and takes certain measures. These measures should include instructing the connected camera to take a picture of the cheating vehicle.

If the tests have been passed, the RSU calculates with the help of the pricing oracle the price  $p$  the user has to pay, depending on factors like the user's attributes, the attributes of the current and previous RSU and auxiliary information (e.g., the time of the day, the current traffic volume). Then the RSU does its part to update the user's wallet. It blindly adds the price  $p$  to the wallet balance  $b$  and increases the PRF counter  $x$  by 1 by calculating  $(c''_{\mathcal{R}}, d''_{\mathcal{R}}) := \text{C1.Com}(\text{CRS}_{\text{com}}^1, (0, p, 0, 1))$ . Then it computes the new RSU commitment  $c_{\mathcal{R}}$  by adding  $c'_{\mathcal{R}}$  and  $c''_{\mathcal{R}}$  (remember that C1 is homomorphic) and also signs it along with the serial number  $s$ . The RSU then sends  $(c_{\mathcal{R}}, d''_{\mathcal{R}}, \sigma_{\mathcal{R}}, p)$  to the user. It also stores several transaction information: The blacklisting transaction information  $\omega^{\text{bl}} := (\phi, p)$  can be used at a later point in the User Blacklisting task (cf. Fig. 40) to calculate the total fee of a fraudulent user. The double-spending transaction information  $\omega^{\text{dsp}} := (\phi, t, u_2)$  enables the TSP to identify the user if he uses the old state of the wallet (with unchanged balance) in another transaction (cf. Fig. 38). The RSU prove participation transaction information  $\omega_{\mathcal{R}}^{\text{pp}} := (s, c_{\text{hid}})$  can be used later in the Prove Participation task (cf. Fig. 36).

The user can then calculate the remaining values needed to update the wallet, e.g., increasing the counter and the balance. Then he can construct the updated wallet

$$\tau := (s, \phi, x^{\text{next}}, \lambda, \mathbf{a}_{\mathcal{U}}, c_{\mathcal{R}}, d_{\mathcal{R}}, \sigma_{\mathcal{R}}, \text{cert}_{\mathcal{R}}, c_{\mathcal{T}}, d_{\mathcal{T}}, \sigma_{\mathcal{T}}, b, u_1^{\text{next}}).$$

The user's output is the updated wallet along with the user prove participation transaction information  $\omega_{\mathcal{U}}^{\text{pp}} := (s, c_{\text{hid}}, d_{\text{hid}})$ . The RSU's output are the user's attributes  $\mathbf{a}_{\mathcal{U}}$ , the attributes  $\mathbf{a}_{\mathcal{R}}^{\text{prev}}$  of the RSU from the previous transaction of the user and the three transaction information  $\omega^{\text{bl}}$ ,  $\omega^{\text{dsp}}$  and  $\omega_{\mathcal{R}}^{\text{pp}}$ .

## D.8 Debt Clearance

After the end of a billing period, the Debt Clearance task is executed between  $\mathcal{U}$  and  $\mathcal{T}$ . The task is presented in two parts: a wrapper protocol  $\pi_{\text{P4TC}}$  (see Fig. 34) and a core protocol DebtClearance (see Fig. 35). The wrapper protocol interacts with other UC entities, pre-processes the input and post-processes the output. The wrapper protocol internally invokes the core protocol

$$\left( (b^{\text{bill}}), (\text{pk}_{\mathcal{U}}, \omega^{\text{bl}}, \omega^{\text{dsp}}) \right) \leftarrow \text{DebtClearance} \langle \mathcal{U} (\text{pk}_{\mathcal{T}}, \text{pk}_{\mathcal{U}}, \text{sk}_{\mathcal{U}}, \tau^{\text{prev}}), \mathcal{T} (\text{pk}_{\mathcal{T}}) \rangle.$$

The user gets the public key  $\text{pk}_{\mathcal{T}}$  of the TSP, his own public and private key  $(\text{pk}_{\mathcal{U}}, \text{sk}_{\mathcal{U}})$  and his current wallet

$$\tau^{\text{prev}} := (s^{\text{prev}}, \phi^{\text{prev}}, x, \lambda, \mathbf{a}_{\mathcal{U}}, c_{\mathcal{R}}^{\text{prev}}, d_{\mathcal{R}}^{\text{prev}}, \sigma_{\mathcal{R}}^{\text{prev}}, \text{cert}_{\mathcal{R}}^{\text{prev}}, c_{\mathcal{T}}, d_{\mathcal{T}}, \sigma_{\mathcal{T}}, b^{\text{prev}}, u_1)$$

as input. The TSP gets its own public key  $\text{pk}_{\mathcal{T}}$  as input.

The protocol is similar to the DebtAccumulation protocol, with the difference that the user here is not anonymous, the wallet balance is no secret and the TSP does not give the user an updated wallet. Like in DebtAccumulation, the TSP first sends a fresh double-spending randomness  $u_2$  to the user and the user calculates the double-spending tag  $t := \text{sk}_{\mathcal{U}}u_2 + u_1 \bmod p$  and the fraud detection ID  $\phi$  for this transaction. He then continues by preparing a proof of knowledge. More precisely, P3 is used to compute a proof  $\pi$  for a statement

**UC-Protocol  $\pi_{P4TC}$  (cont.) – Task *Debt Clearance***

*User input:* (clear\_debt,  $s^{\text{prev}}$ )

*TSP input:* (clear\_debt)

- (1) At the user side:
  - Load the internally recorded  $(pk_u, sk_u)$ .<sup>‡</sup>
  - Receive  $pk_{\mathcal{T}}$  from the bulletin-board  $\overline{\mathcal{G}}_{\text{bb}}$  for PID  $pid_{\mathcal{T}}$ .<sup>‡</sup>
  - Load the internally recorded token  $\tau^{\text{prev}}$  for serial number  $s^{\text{prev}}$ .<sup>‡</sup>
- (2) At the TSP side:
  - Load the internally recorded  $(pk_{\mathcal{T}}, sk_{\mathcal{T}})$ .<sup>‡</sup>
- (3) Both sides: Run the code of DebtClearance between the user and the TSP (see Fig. 35)

$$\left( (b^{\text{bill}}), (pk_u, \omega^{\text{bl}}, \omega^{\text{dsp}}) \right) \leftarrow \text{DebtClearance} \left( \mathcal{U}(pk_{\mathcal{T}}, pk_u, sk_u, \tau^{\text{prev}}), \mathcal{T}(pk_{\mathcal{T}}) \right).$$

- (4) At the TSP side:
  - Record  $\omega^{\text{bl}}$  and  $\omega^{\text{dsp}}$  internally.
  - Parse  $\phi$  and  $-b^{\text{bill}}$  from  $\omega^{\text{bl}}$ .

*User output:* ( $b^{\text{bill}}$ )

*TSP output:*  $(pk_u, \phi, b^{\text{bill}})$

<sup>‡</sup>If this does not exist, output  $\perp$  and abort.

Fig. 34. The UC-protocol  $\pi_{P4TC}$  (cont. from Fig. 18)

*stmt* from the language  $L_{\text{gp}}^{(3)}$  defined by

$$L_{\text{gp}}^{(3)} := \left( \begin{array}{l} \left( \begin{array}{l} pk_u \\ pk_{\mathcal{T}} \\ pk_{\mathcal{T}}^{\text{cert}} \\ \phi \\ a_u \\ a_{\mathcal{R}}^{\text{prev}} \\ B^{\text{prev}} \\ t \\ u_2 \end{array} \right)^{\top} \\ \left( \begin{array}{l} \exists \lambda, x, u_1, sk_u \in \mathbb{Z}_p; \phi^{\text{prev}}, s^{\text{prev}}, X, \Lambda, U_1, \\ d_{\mathcal{R}}^{\text{prev}}, d_{\mathcal{T}} \in G_1; pk_{\mathcal{R}}^{\text{prev}} \in G_1^3, c_{\mathcal{R}}^{\text{prev}}, c_{\mathcal{T}} \in G_2; \\ \sigma_{\mathcal{R}}^{\text{prev}}, \sigma_{\mathcal{R}}^{\text{certprev}}, \sigma_{\mathcal{T}} \in G_2^2 \times G_1 : \\ \\ \text{C1.Open}(\text{CRS}_{\text{com}}^1, (\Lambda, pk_u), c_{\mathcal{T}}, d_{\mathcal{T}}) = 1 \\ \text{C1.Open}(\text{CRS}_{\text{com}}^1, (\Lambda, B^{\text{prev}}, U_1, X), c_{\mathcal{R}}^{\text{prev}}, d_{\mathcal{R}}^{\text{prev}}) = 1 \\ \text{S.Vfy}(pk_{\mathcal{T}}^{\mathcal{T}}, \sigma_{\mathcal{T}}, (c_{\mathcal{T}}, a_u)) = 1 \\ \text{S.Vfy}(pk_{\mathcal{R}}^{\text{prev}}, \sigma_{\mathcal{R}}^{\text{prev}}, (c_{\mathcal{R}}^{\text{prev}}, s^{\text{prev}})) = 1 \\ \text{S.Vfy}(pk_{\mathcal{T}}^{\text{cert}}, \sigma_{\mathcal{R}}^{\text{certprev}}, (pk_{\mathcal{R}}^{\text{prev}}, a_{\mathcal{R}}^{\text{prev}})) = 1 \\ \phi^{\text{prev}} = \text{PRF}(\lambda, x - 1), \phi = \text{PRF}(\lambda, x), \\ t = sk_u u_2 + u_1 \\ pk_u = g_1^{sk_u}, U_1 = g_1^{u_1}, X = g_1^x, \Lambda = g_1^\lambda \end{array} \right) \end{array} \right) \quad (8)$$

The proof is a simplified version of the one in the DebtAccumulation protocol. The balance and the public user key are now in the statement and not in the witness and one does not need to prove anything about  $c'_{\mathcal{R}}$  and  $c_{\text{hid}}$ .

The user then sends  $(pk_u, \pi, \phi, a_u, a_{\mathcal{R}}^{\text{prev}}, b^{\text{prev}}, t)$  to the TSP. Unlike in DebtAccumulation, the balance and the user's public key are transmitted. The TSP then checks the validity of the proof and signals the user that the proof successfully verified.

At the end of the protocol, the TSP outputs the user's public key  $pk_u$ , the blacklisting transaction information  $\omega^{\text{bl}} := (\phi, -b^{\text{prev}})$  and the double-spending transaction information  $\omega^{\text{dsp}} := (\phi, t, u_2)$ . The user just outputs his final debt  $b^{\text{prev}}$  for this billing period.

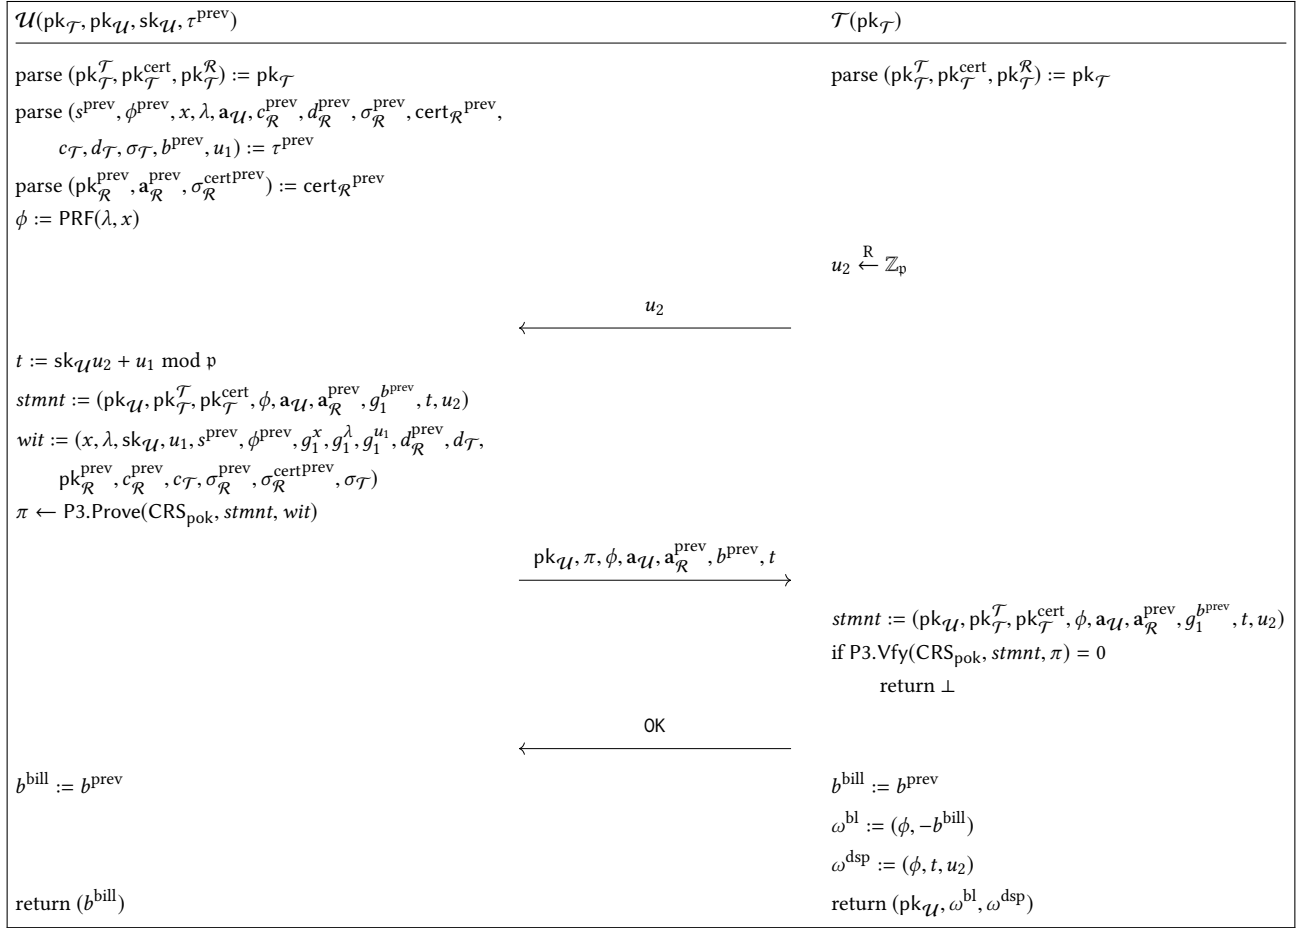


Fig. 35. Debt Clearance Core Protocol

Note that the wallet itself is discarded. It is expected that the user and the TSP execute the Wallet Issuing task next, to give the user a fresh wallet. After the protocol has ended, the TSP can issue an invoice to the user for the current billing period. The specifics of the actual payment process are out of scope.

#### D.9 Prove Participation

The Prove Participation task is used by a user to prove to the SA that he behaved honestly at a specific Debt Accumulation transaction. In the case that more than one vehicle is captured on a photograph taken by a RSU camera after a fraud occurred, this protocol can be used to identify the fraudulent driver.<sup>25</sup> The SA recovers the identities of the users captured on the photograph and executes the Prove Participation task which each of them. The user that is not able to prove that he honestly participated in a corresponding RSU transaction is found guilty.

<sup>25</sup>Of course, the task can also be used in the case that only one vehicle was captured on the photograph to eliminate the possibility that the RSU falsely instructed the camera to take a photograph.

**UC-Protocol  $\pi_{P4TC}$  (cont.) – Task *Prove Participation****User input:* (prove\_participation)*SA input:* (prove\_participation,  $pk_{\mathcal{U}}, S_{\mathcal{R}}^{PP}$ )

- (1) At the SA side:
  - Load the internally recorded set  $\Omega_{\mathcal{R}}^{PP}$  of all prove participation transaction information with serial numbers in  $S_{\mathcal{R}}^{PP}$ .
- (2) At the user side:
  - Load the internally recorded set  $\Omega_{\mathcal{U}}^{PP}$  of all prove participation transaction information.
- (3) Both sides: Run the code of ProveParticipation between the User and the SA (see Fig. 37)

$$((out_{\mathcal{U}}), (out_{SA})) \leftarrow \text{ProveParticipation} \langle \mathcal{U}(\Omega_{\mathcal{U}}^{PP}), SA(pk_{\mathcal{U}}, S_{\mathcal{R}}^{PP}, \Omega_{\mathcal{R}}^{PP}) \rangle.$$


*User output:* ( $out_{\mathcal{U}}$ )*SA output:* ( $out_{SA}$ )<sup>†</sup>If this does not exist, output  $\perp$  and abort.Fig. 36. The UC-protocol  $\pi_{P4TC}$  (cont. from Fig. 18)

Fig. 37. Prove Participation Core Protocol

The task is presented in two parts: a wrapper protocol  $\pi_{P4TC}$  (see Fig. 36) and a core protocol ProveParticipation (see Fig. 37). In the wrapper protocol, the SA interacts with other UC entities and processes the set  $S_{\mathcal{R}}^{PP}$  of all serial numbers that were recorded by the RSU that took the photo at roughly the time the photo was taken. In particular, the SA loads the internally recorded set  $\Omega_{\mathcal{R}}^{PP}$  of all RSU prove participation transaction information that contain serial numbers that are also in  $S_{\mathcal{R}}^{PP}$ . Afterwards, the wrapper protocol invokes the core protocol

$$((out_{\mathcal{U}}), (out_{SA})) \leftarrow \text{ProveParticipation} \langle \mathcal{U}(\Omega_{\mathcal{U}}^{PP}), SA(pk_{\mathcal{U}}, S_{\mathcal{R}}^{PP}, \Omega_{\mathcal{R}}^{PP}) \rangle.$$

**UC-Protocol  $\pi_{P4TC}$  (cont.) – Task *Double-Spending Detection***

*TSP input:* (scan\_for\_fraud,  $\phi$ )

- (1) Load the internally recorded set  $\Omega^{\text{dsp}}$  of all double-spending transaction information.
- (2) Pick transaction information  $\omega^{\text{dsp}} = (\phi, t, u_2)$  and  $\omega^{\text{dsp}'} = (\phi', t', u_2')$  from the database  $\Omega^{\text{dsp}}$ , such that  $\phi = \phi'$  and  $u_2 \neq u_2'$ .<sup>‡</sup>
- (3)  $\text{sk}_{\mathcal{U}} := (t - t') \cdot (u_2 - u_2')^{-1} \bmod p$ .
- (4)  $\text{pk}_{\mathcal{U}} := g_1^{\text{sk}_{\mathcal{U}}}$ .
- (5)  $\pi := \text{sk}_{\mathcal{U}}$ .

*TSP output:* ( $\text{pk}_{\mathcal{U}}, \pi$ )

<sup>‡</sup>If this does not exist, output  $\perp$  and abort.

Fig. 38. The UC-protocol  $\pi_{P4TC}$  (cont. from Fig. 18)

The SA gets the user's public key  $\text{pk}_{\mathcal{U}}$ , the set  $S_{\mathcal{R}}^{\text{PP}}$  of serial numbers that were observed roughly at the time the photograph was taken and the set of corresponding RSU prove participation transaction information  $\Omega_{\mathcal{R}}^{\text{PP}}$  as input. The user gets his internally recorded set  $\Omega_{\mathcal{U}}^{\text{PP}}$  of all prove participation transaction information as input.

The protocol itself is simple. The SA first sends  $S_{\mathcal{R}}^{\text{PP}}$  to the user who searches  $\Omega_{\mathcal{U}}^{\text{PP}}$  for a user prove participation transaction information  $\omega_{\mathcal{U}}^{\text{PP}}$  for which the serial number  $s$  in  $\omega_{\mathcal{U}}^{\text{PP}}$  is also in  $S_{\mathcal{R}}^{\text{PP}}$ . If one is found, the user's output bit  $\text{out}_{\mathcal{U}}$  is set to OK, if none is found,  $\text{out}_{\mathcal{U}}$  is set to NOK ("not ok"). The user then sends  $\omega_{\mathcal{U}}^{\text{PP}} := (s, c_{\text{hid}}, d_{\text{hid}})$  to the SA and the SA checks if  $(s, c_{\text{hid}}) \in \Omega_{\mathcal{R}}^{\text{PP}}$  and if  $d_{\text{hid}}$  is the opening of  $c_{\text{hid}}$  under  $\text{pk}_{\mathcal{U}}$ . If both checks succeed, the SA's output bit  $\text{out}_{\text{SA}}$  is set to OK, if at least one check fails,  $\text{out}_{\text{SA}}$  is set to NOK. At the end of the protocol both parties output their bits  $\text{out}_{\mathcal{U}}$  and  $\text{out}_{\text{SA}}$ , respectively.<sup>26</sup>

If the SA's output equals NOK, the user is found guilty and appropriate measures are taken (e.g., the user gets blacklisted).

#### D.10 Double-Spending Detection

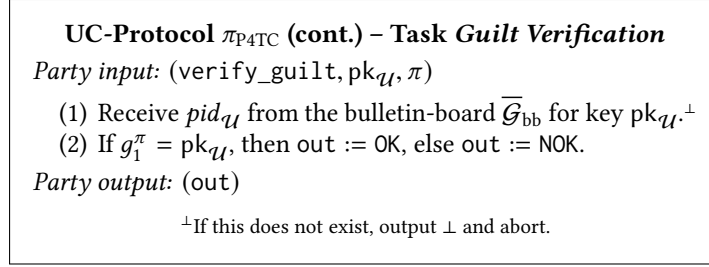
The double-spending transaction information  $\omega^{\text{dsp}}$  collected by the RSUs are periodically transmitted to the TSP's database, which is regularly checked for two double-spending transaction information associated with the same serial number. If the database contains two such records, the Double-Spending Detection task (see Fig. 38) can be used by the TSP to extract the public key of the user these double-spending transaction information belong to as well as a proof (such as his secret user key) that the user is guilty.

In particular, the task gets a serial number  $s$  as input and searches the internal database for two double-spending transaction information  $\omega^{\text{dsp}} = (\phi, t, u_2)$  and  $\omega^{\text{dsp}'} = (\phi', t', u_2')$  that contain the same serial number  $s = s'$  but not the same double-spending randomness  $u_2 \neq u_2'$ . Then the fraudulent user's secret key can be recovered as  $\text{sk}_{\mathcal{U}} := (t - t') \cdot (u_2 - u_2')^{-1} \bmod p$ . His public key is then  $\text{pk}_{\mathcal{U}} := g_1^{\text{sk}_{\mathcal{U}}}$ . The secret key  $\text{sk}_{\mathcal{U}}$  can be used as a proof of guilt in the Guilt Verification task (cf. Fig. 39).

Every user that is convicted of double-spending is added to the blacklist via the User Blacklisting task (cf. Figs. 40 and 41) and additional measures are taken (these are out of scope).

<sup>26</sup>Note that after a successful (both parties output OK) execution of the Prove Participation protocol a single transaction can be linked to the user. But as long as the user does not get guiltlessly photographed at every RSU he passes, tracking is not possible.



Fig. 39. The UC-protocol  $\pi_{P4TC}$  (cont. from Fig. 18)

### D.11 Guilt Verification

Whether a user is indeed guilty of double-spending can be verified using the Guilt Verification task depicted in Fig. 39. This algorithm may be run by anyone, in particular by justice. Essentially, the algorithm checks if a given public user key  $pk_u$  and a proof of guilt  $\pi$  match. This is easily accomplished because they match if and only if  $g_1^\pi = pk_u$  holds. This equation holds if and only if  $\pi$  equals the user's secret key  $sk_u$  that was recovered using the Double-Spending Detection task (cf. Fig. 38).

### D.12 User Blacklisting

The User Blacklisting task executed between the DR and TSP is used to put a user on the blacklist. There are several reasons why a user is entered on the blacklist:

- (1) A user did not submit his balance at the end of the billing period.
- (2) A user did not physically pay his debt after submitting his balance.
- (3) A user has been convicted of double-spending.
- (4) The wallet (or the vehicle including the wallet) of a user has been stolen and the user wants to prevent the thief from paying tolls from his account.

Blacklisted users are unable to get issued a new wallet and they also get photographed at every RSU they pass. They may also be punished by other means (which are out of scope).

The User Blacklisting task is presented in two parts: a wrapper protocol  $\pi_{P4TC}$  (see Fig. 40) and a core protocol UserBlacklisting (see Fig. 41). The wrapper protocol interacts with other UC entities, pre-processes the input and afterwards internally invokes the core protocol

$$((OK), (\Phi_u)) \leftarrow \text{UserBlacklisting} \langle DR(pk_{DR}, sk_{DR}, pk_u^{DR}), \mathcal{T}(HTD_u) \rangle.$$

The DR gets as input its public and private key  $(pk_{DR}, sk_{DR})$  together with the public key  $pk_u^{DR}$  of the user to be blacklisted. The TSP gets a set  $HTD_u$  containing all his hidden user trapdoors from the current billing period as input.<sup>27</sup> We assume that the DR and TSP agreed upon which user is going to be blacklisted before the protocol out-of-band.

At the beginning of the protocol the TSP sends its input  $HTD$  to the DR. The DR then recovers the corresponding wallet ID  $\lambda$  for every  $htd := (pk_u^T, s, \lambda'', e^*) \in HTD$ . To this end, the DR decrypts  $e^*$  to get  $(\Lambda'_0, \dots, \Lambda'_{\ell-1}, \Lambda'', pk_u^T)$ . Firstly, the DR checks if the decrypted public key  $pk_u^T$  equals the expected public  $pk_u^{DR}$  of the correct user in

<sup>27</sup>In the case that a user owns more than one vehicle he can have more than one wallet and hence more than one hidden user trapdoor is stored at the TSP for this user.

**UC-Protocol  $\pi_{P4TC}$  (cont.) – Task *User Blacklisting***

*DR input:* (blacklist\_user,  $\text{pk}_{\mathcal{U}}^{DR}$ )

*TSP input:* (blacklist\_user,  $\text{pk}_{\mathcal{U}}^{\mathcal{T}}$ )

(1) At the DR side:

- Load the internally recorded  $(\text{pk}_{DR}, \text{sk}_{DR})^\perp$ .
- Receive  $\text{pk}_{\mathcal{U}}^{DR}$  from the bulletin-board  $\overline{\mathcal{G}}_{\text{bb}}$  for PID  $\text{pid}_{\mathcal{U}}^{DR}$ .<sup>‡</sup>

(2) At the TSP side:

- Load internally recorded set  $HTD$  of all hidden user trapdoors and set  $HTD_{\mathcal{U}} := \{htd \mid (\text{pk}_{\mathcal{U}}^{\mathcal{T}}, \cdot, \cdot, \cdot) \in HTD\}$ .

(3) Both sides: Run the code of UserBlacklisting between the DR and the TSP (see Fig. 41)

$$((\text{OK}), (\Phi_{\mathcal{U}})) \leftarrow \text{UserBlacklisting} \langle DR(\text{pk}_{DR}, \text{sk}_{DR}, \text{pk}_{\mathcal{U}}^{DR}), \mathcal{T}(HTD_{\mathcal{U}}) \rangle.$$

(4) At the TSP side:

- Load the internally recorded set  $\Omega^{\text{bl}}$  of all blacklisting transaction information.
- Let  $\Omega_{\mathcal{U}}^{\text{bl}}$  be the subset of transaction entries  $\omega^{\text{bl}} = (\phi, p)$  with fraud detection IDs  $\phi \in \Phi_{\mathcal{U}}$ .
- $b^{\text{bill}} := \sum_{\omega^{\text{bl}} \in \Omega_{\mathcal{U}}^{\text{bl}}} p$ .

*DR output:* (OK)

*TSP output:* ( $b^{\text{bill}}, \Phi_{\mathcal{U}}$ )

<sup>‡</sup>If this does not exist, output  $\perp$  and abort.

Fig. 40. The UC-protocol  $\pi_{P4TC}$  (cont. from Fig. 18)

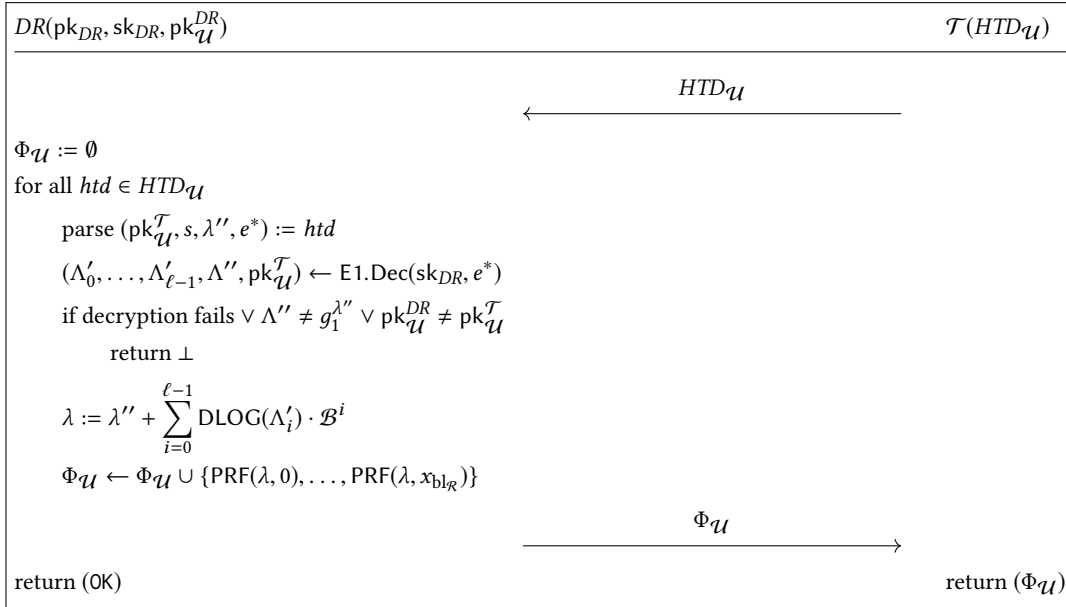


Fig. 41. User Blacklisting Core Protocol

question.<sup>28</sup> This way, the DR cannot be tricked into recovering the wallet ID of another (possibly innocent) user. Since each  $\lambda'_i$  is small ( $\lambda'_i < \mathcal{B}$ ), the DR can compute the discrete logarithms of  $(\Lambda'_0, \dots, \Lambda'_{\ell-1})$  in a reasonable amount of time to recover  $(\lambda'_0, \dots, \lambda'_{\ell-1})$ . This algorithm is also not time-critical and is expected to be executed only a few times per billing period. Therefore, the amount of required computation should be acceptable. Secondly, the DR checks if the claimed TSP's share  $\lambda''$  of the wallet ID is consistent to the decrypted  $\Lambda'' \stackrel{?}{=} g_1^{\lambda''}$ . (Remember that  $\lambda''$  is directly stored in *htd*.) The DR calculates the wallet ID as  $\lambda = \lambda'' + \sum_{i=0}^{\ell-1} \lambda'_i \cdot \mathcal{B}^i$ . Finally, the DR sends the union of all sets  $\{\text{PRF}(\lambda, 0), \dots, \text{PRF}(\lambda, x_{\text{bl}_R})\}$  of fraud detection IDs for every *htd*  $\in$  *HTD* to the TSP. The blacklist parameter  $x_{\text{bl}_R}$  is chosen in such a way that a user is expected to perform at most  $x_{\text{bl}_R}$  executions of the Debt Accumulation task in a single billing period. The DR's output of the core protocol is simply OK, while the TSP's output is a set of fraud detection IDs.

After the core protocol has terminated, the TSP calculates the total toll amount the user owes for the billing period in question in the wrapper protocol. To this end, the TSP calculates the subset  $\Omega_{\mathcal{U}}^{\text{bl}} \subseteq \Omega^{\text{bl}}$  of all blacklisting transaction information that correspond to the unveiled fraud detection IDs in  $\Phi_{\mathcal{U}}$ . By summing up all the prices in  $\Omega_{\mathcal{U}}^{\text{bl}}$ , the TSP can calculate the total fee  $b^{\text{bill}}$  the user owes.

After the wrapper protocol has terminated, the fraud detection IDs  $\Phi_{\mathcal{U}}$  are added to the RSU blacklist  $\text{bl}_R$  and the user's public key  $\text{pk}_{\mathcal{U}}$  is added on the TSP blacklist  $\text{bl}_T$ .

In the Wallet Issuing task the TSP uses the blacklist  $\text{bl}_T$  to prevent a user that did not pay its invoice from receiving a fresh wallet. In the Debt Accumulation task a RSU checks if the fraud detection ID that the current user presents is on the RSU blacklist  $\text{bl}_R$ .

### D.13 Wallet Verification

The WalletVerification algorithm is depicted in Fig. 42. A user can verify with this algorithm that the wallet he stores at the end of a transaction is valid. In particular, the algorithm verifies that the commitments  $c_T$  and  $c_R$  are valid and contain the values they are supposed to contain, that  $\sigma_T$  is a valid signature under  $\text{sk}_T^T$  of  $c_T$  and  $\mathbf{a}_{\mathcal{U}}$ , that  $\sigma_R$  is a valid signature under  $\text{sk}_R$  of  $c_R$  and  $s$ , that the certificate  $\text{cert}_R$  containing  $\text{pk}_R$  is valid and that the fraud detection id  $\phi$  was calculated using the correct values.

Of course, this algorithm can also be run by a third party to verify the validity of a wallet (since no secret keys are needed to run this algorithm).

## E ADVERSARIAL MODEL

For our security analysis to hold we consider a restricted class of adversarial environments  $\mathcal{Z}$  and will argue why these restrictions are reasonable.

### E.1 Restricted Corruption

Firstly, we only consider security under static corruption. This is a technical necessity to enable the use of PRFs to generate fraud detection IDs. With adaptive corruption the simulator would be required to come up with a consistent PRF that could explain the up to the point of corruption uniformly and randomly drawn fraud detection IDs. We deem static corruption to provide a sufficient level of security as a statically corrupted party may always decide to interact honestly first and then deviate from the protocol later. Adaptive corruption only comes into play with deniability which is not part of our desired properties.

<sup>28</sup>N.b.: The public key  $\text{pk}_{\mathcal{U}}^T$  is stored redundantly: in clear as part of  $\text{htd} := (\text{pk}_{\mathcal{U}}^T, s, \lambda'', e^*)$  and encrypted as part of the cipher text  $e^* = \text{Enc}(\Lambda'_0, \dots, \Lambda'_{\ell-1}, \Lambda'', \text{pk}_{\mathcal{U}}^T)$ . The DR only considers the latter version as this is protected by the non-malleability of the CCA encryption. The outer public key is only required as management information and utilized by the TSP which cannot look inside  $e^*$ . The DR ignores the outer  $\text{pk}_{\mathcal{U}}^T$ .

<pre> WalletVerification(pk<sub>T</sub>, pk<sub>U</sub>, τ) ----- parse (pk<sub>T</sub><sup>T</sup>, pk<sub>T</sub><sup>cert</sup>, pk<sub>T</sub><sup>R</sup>) := pk<sub>T</sub> parse (s, φ, x<sup>next</sup>, λ, a<sub>U</sub>, c<sub>R</sub>, d<sub>R</sub>, σ<sub>R</sub>, cert<sub>R</sub>, c<sub>T</sub>, d<sub>T</sub>, σ<sub>T</sub>, b, u<sub>1</sub><sup>next</sup>) := τ parse (pk<sub>R</sub>, a<sub>R</sub>, σ<sub>R</sub><sup>cert</sup>) := cert<sub>R</sub> if   C1.Open(CRS<sub>com</sub><sup>1</sup>, (g<sub>1</sub><sup>λ</sup>, pk<sub>U</sub>), c<sub>T</sub>, d<sub>T</sub>) = 0 ∨   S.Vfy(pk<sub>T</sub><sup>T</sup>, σ<sub>T</sub>, (c<sub>T</sub>, a<sub>U</sub>)) = 0 ∨   C1.Open(CRS<sub>com</sub><sup>1</sup>, (g<sub>1</sub><sup>λ</sup>, g<sub>1</sub><sup>b</sup>, g<sub>1</sub><sup>u<sub>1</sub><sup>next</sup></sup>, g<sub>1</sub><sup>x<sup>next</sup></sup>), c<sub>R</sub>, d<sub>R</sub>) = 0 ∨   S.Vfy(pk<sub>R</sub>, σ<sub>R</sub>, (c<sub>R</sub>, s)) = 0 ∨   S.Vfy(pk<sub>T</sub><sup>cert</sup>, σ<sub>R</sub><sup>cert</sup>, (pk<sub>R</sub>, a<sub>R</sub>)) = 0 ∨   PRF(λ, x<sup>next</sup> - 1) ≠ φ then return 0 else return 1 </pre>
--

Fig. 42. Algorithm for Wallet Verification

Secondly, we only consider adversaries  $\mathcal{Z}$  that corrupt one of the following sets<sup>29</sup>:

- (1) A subset of users.
- (2) All users and a subset of RSUs, TSP and SA.
- (3) A subset of RSUs, TSP and SA.
- (4) All of RSUs, TSP and SA as well as a subset of users.

We subsume the cases (1) and (2) under the term *Operator Security* and the cases (3) and (4) under the term *User Security*. For both Operator Security and User Security the two subordinate cases are collectively treated by the same proof. It is best to picture the cases inversely: To prove Operator Security we consider a scenario in which at least some parties at the operator's side remain honest; to prove User Security we consider a scenario in which at least some users remain honest. Please note that both scenarios also commonly cover the case in which all parties are corrupted, however, this extreme case is tedious as it is trivially simulatable.

One might believe that the combination of all cases above should already be sufficient to guarantee privacy, security and correctness under arbitrary corruption. For example, case (4) guarantees that privacy and correctness of accounting are still provided for honest users, even if all of the operator's side and some fellow users are corrupted. This ought to be the worst case from a honest user's perspective. Further note that the proof of indistinguishability quantifies over all environments  $\mathcal{Z}$ . This includes environments that—still in case (4)—first corrupt all the operator's side but then let some (formally corrupted) parties follow the protocol honestly.

However, consider a scenario in which a party acts as a Man-in-the-Middle (MitM) playing the roles of a user and an RSU at the same time while interacting with an honest user in the left interaction and an honest RSU in the right interaction. The MitM simply relays messages back and forth unaltered. If the MitM approaches the RSU and the RSU requests the MitM to participate in Debt Accumulation, the MitM relays all messages of the RSU to the honest user (possibly driving the same road behind the MitM). The honest user replies and the MitM forwards the messages to the honest RSU. The MitM passes by the RSU unnoticed and untroubled, while the honest user pays for the MitM.

<sup>29</sup>Note that “subset” also includes the empty or full set.

This scenario is not captured by any of the above cases and is the missing gap towards arbitrary corruption. As the MitM is corrupted and plays both roles of a user and RSU, this falls into case (2) or (4). But either all users are corrupted in case (2), which contradicts the existence of an honest user in the left interaction, or all of RSUs, TSP and SA are corrupted in case (4), which does not allow for an honest RSU in the right interaction.

This attack is known as relay attack. Please note that the MitM does not need to break any cryptographic assumption for this kind of attack as it just poses as a prolonged communication channel. There are some possible counter measures that can be applied in the real world. For example, using distance-bounding the honest user could refuse to participate in the protocol, if the RSU is known to be too far away. However, these are physical counter measures and thus are not captured by the UC notion nor any other cryptographic notion. Actually, it is a strength of the UC model that this gap is made explicit. For example, a set of game-based security notions that cover a list of individual properties would most likely not unveil this issue.

## E.2 Channel Model

Most of the time we assume channels to be secure and authenticated. The only exception is Debt Accumulation, which uses a secure but only half-authenticated channel. Half-authenticated channel means only the RSU authenticates itself and the user does not. These channels exempt us from the burden of defining a simulator for the case where only honest parties interact with each other and rules out some trivial replay attacks. Of course, the authentication of the channels must be tied to the parties' credentials used in the toll collection system. In other words, the same key registration service that registers the public keys for the toll collection system must also be used to register the public keys to authenticate the communication.

## E.3 Handling of Aborts

Lastly, we assume our functionality also uses the implicit writing conventions for ideal functionalities [20]. In particular, our simulator can delay outputs and abort at any point. Beyond that, the simulator has the power to override the output to honest parties with an abort reason (e.g., "blacklisting") if it decides to abort.

Another important aspect with respect to aborts is that privacy may be partially lost for an honest user if a task aborts prematurely. In this case the user's identity can be unveiled if he chooses to take part in another transaction with the same wallet. The reason for this is the double-spending detection mechanism. If the (honest) user has not correctly updated his previous state, the user must start the next interaction from the same state as before. Thus an (honest) user can be tricked into committing a double-spending without any fault of his own. We explicitly model this artifact into the simulator.

Again, as in [Appendix E.1](#) this kind of "attack" does not require to break any cryptographic assumptions. Hence, the UC notion as well as any other cryptographic security notion does not take aborts into account—the more so as aborts can occur at any time due to technical problems. To mitigate the effect of aborts we propose that each user has one or more backup wallets and switches over to another wallet if the principal wallet becomes unusable and the user does not accept to become linkable for a single transaction. Of course, at the end of a billing period a user must always clear all of his wallets. If aborts occur more frequently than one would reasonably expect due to technical problems and the RSU/TSP is suspected to purposely abort in order to lift privacy, the user (or some NGO) is expected to file a claim. However, these are non-technical countermeasures.

## F SECURITY PROOF

In this appendix we show that  $\pi_{\text{P4TC}}$  UC-realizes  $\mathcal{F}_{\text{P4TC}}$  in the  $(\mathcal{F}_{\text{CRS}}, \overline{\mathcal{G}}_{\text{bb}})$ -hybrid model for static corruption. More precisely, we show the following theorem:

**THEOREM F.1 (SECURITY STATEMENT).** *Assume that the SXDH-problem is hard for  $\text{gp} := (G_1, G_2, G_T, e, p, g_1, g_2)$ , the Co-CDH problem is hard for  $(G_1, G_2)$ , the  $n_{\text{PRF}}$ -DDHI problem is hard for  $G_1$ , and the DLOG-problem is hard*

for  $G_1$  and our building blocks (NIZK, commitment schemes, signature scheme, encryption schemes and PRF) are instantiated as described in [Appendix C.2](#). Then

$$\pi_{\text{P4TC}}^{\mathcal{F}_{\text{CRS}}, \bar{\mathcal{G}}_{\text{bb}}} \geq_{\text{UC}} \mathcal{F}^{\bar{\mathcal{G}}_{\text{bb}}},$$

if either only a subset of the users or only a subset of the RSUs, the TSP and the SA is statically corrupted.

Please note, that the hardness of the Co-CDH problem and DLOG-problem is already implied by the SXDH-assumption. For a discussion of the reasons and why this limited corruption model is not a severe restriction from a practical vantage point see [Appendix E.1](#).

We prove [Theorem F.1](#) in three steps:

- In [Appendix F.1](#) we first show in [Theorem F.21](#) that  $\pi_{\text{P4TC}}^{\mathcal{F}_{\text{CRS}}, \bar{\mathcal{G}}_{\text{bb}}}$  correctly implements  $\mathcal{F}^{\bar{\mathcal{G}}_{\text{bb}}}$ , if all parties follow the protocol honestly.
- In the next step, [Theorem F.22](#) proves  $\pi_{\text{P4TC}}^{\mathcal{F}_{\text{CRS}}, \bar{\mathcal{G}}_{\text{bb}}} \geq_{\text{UC}} \mathcal{F}^{\bar{\mathcal{G}}_{\text{bb}}}$ , if only a subset of the users is statically corrupted. We call this case ‘‘System Security’’. [Theorem F.22](#) is presented in [Appendix F.2](#).
- In the last step, [Theorem F.36](#) proves  $\pi_{\text{P4TC}}^{\mathcal{F}_{\text{CRS}}, \bar{\mathcal{G}}_{\text{bb}}} \geq_{\text{UC}} \mathcal{F}^{\bar{\mathcal{G}}_{\text{bb}}}$ , if only a subset of the RSUs is statically corrupted. We call this case ‘‘User Security and Privacy’’ and present it in [Appendix F.3](#).

PROOF OF [THEOREM F.1](#). The theorem is immediately implied by [Theorems F.21, F.22](#) and [F.36](#).  $\square$

Before giving the proof in full detail and complexity in [Appendices F.1](#) to [F.3](#), [Section 5](#) explains our course of action and outlines the proof on a high level.

## F.1 Proof of Correctness

Many papers that show some protocol to be UC-secure consider rather simple cases (e.g., a commitment, an oblivious transfer, a coin toss) and correctness is mostly obvious. In contrast, our protocol is a complex system with polynomially many parties that can reactively interact forever, i.e., the protocol itself has no inherent exit point except that at some point the polynomially bounded runtime of the environment is exhausted. In order to prove the security of our system later (see [Appendices F.2](#) and [F.3](#)) we make use of a theoretical construct we call the *Augmented Transaction Graph*. This graph helps us to link the interactions of the parties across the various tasks our protocol provides. However, the Augmented Transaction Graph is also useful to prove correctness and hence we use this proof to introduce the concept.


In this section no security reduction occurs, because we assume all parties to be honest and solely show that the subroutine input/output of the main parties is indistinguishable between the ideal and the real model.

We start with a series of simple lemmas about the ideal functionality  $\mathcal{F}_{\text{P4TC}}$  which also help to develop a good conception about how the individual tasks/transactions are connected.

Internally,  $\mathcal{F}_{\text{P4TC}}$  stores a pervasive database *TRDB* whose entries *trdb* are of the form

$$\text{trdb} = (s^{\text{prev}}, s, \phi, x, \lambda, \text{pid}_{\mathcal{U}}, \text{pid}_{\mathcal{R}}, p, b).$$

This set can best be visualized as a directed graph in which each node represents the state of a user *after* the respective transaction, i.e., at the end of an execution of Wallet Issuing, Debt Accumulation or Debt Clearance, and the edges correspond to the transition from the previous to the next state. Each *trdb* entry represents a node together with an edge pointing to its predecessor node. The node is labeled with  $(s, \phi, x, \lambda, \text{pid}_{\mathcal{U}}, b)$  and identified by  $s$ . The edge to the predecessor is identified by  $(s^{\text{prev}}, s)$  and labeled with  $(\text{pid}_{\mathcal{R}}, p)$ . See [Fig. 43](#) for a depiction. Transaction entries or nodes that are inserted by Wallet Issuing do not have a predecessor, therefore  $s^{\text{prev}} = \perp$  and also  $p = 0$  holds. All other tasks besides Wallet Issuing, Debt Accumulation and Debt Clearance do not alter the graph but only query it. We show that the graph is a directed forest, i.e., a set of directed trees. Wallet Issuing

Fig. 43. An entry  $trdb \in TRDB$  visualized as an element of a directed graph

creates a new tree by inserting a new root node. Debt Accumulation and Debt Clearance extend a tree. Debt Clearance results in a leaf node from where no further extension is possible. As long as no double spending occurs, each tree is a path graph.

*Definition F.2 (Ideal Transaction Graph (informal)).* The transaction database  $TRDB = \{trdb_i\}$  with  $trdb = (s^{prev}, s, \phi, x, \lambda, pid_U, pid_R, p, b)$  is a directed, labeled graph as defined above. This graph is called the *Ideal Transaction Graph*.

LEMMA F.3. *The Ideal Transaction Graph  $TRDB$  is a forest.*

PROOF.  $TRDB$  is a forest, if and only if it is cycle-free and every node has in-degree at most one. A new node is only inserted in the scope of Wallet Issuing, Debt Accumulation or Debt Clearance. Proof by Induction: The statement is correct for the empty  $TRDB$ . If Wallet Issuing is invoked, a new node with no predecessor is inserted. Moreover, the serial number  $s$  of the new node is randomly chosen from the set of unused serial numbers, i.e., it is unique and no existing node can point to the new node as its predecessor. If Debt Accumulation or Debt Clearance is invoked, a new node is inserted that points to an existing node. Again, the serial number  $s$  of the new node is randomly chosen from the set of unused serial numbers, i.e., it is unique and no existing node can point to the new node as its predecessor. Hence, no cycle can be closed. Since the only incoming edge of a node is defined by the stated predecessor  $s^{prev}$  (which may also be  $\perp$ ), each vertex has in-degree at most one.  $\square$

LEMMA F.4. *The wallet ID  $\lambda$  maps one-to-one and onto a connected component (i.e., tree) of the Ideal Transaction Graph.*

PROOF. “ $\Leftarrow$ ”: Let  $trdb_i$  be an arbitrary node in  $TRDB$  and  $\lambda$  be its wallet ID. Furthermore let  $trdb_i^*$  be the root of the tree containing  $trdb_i$ . Then on the (unique) path from  $trdb_i^*$  to  $trdb_i$ , every node apart from  $trdb_i^*$  was inserted by means of either Debt Accumulation or Debt Clearance, both of which ensure the inserted node has the same  $\lambda$  as its predecessor. By induction over the length of the path,  $trdb_i$  has the same wallet ID as  $trdb_i^*$  and hence the wallet ID is a locally constant function on  $TRDB$ .

“ $\Rightarrow$ ”: For contradiction assume there are two nodes  $trdb_i$  and  $trdb_j$  with equal wallet IDs  $\lambda_i = \lambda_j$  in two different connected components. Pick the root nodes  $trdb_i^*$  and  $trdb_j^*$  of their respective trees. By “ $\Leftarrow$ ” it we get  $\lambda_i^* = \lambda_i = \lambda_j = \lambda_j^*$ , i.e., the root nodes have equal wallet IDs, too. Both root nodes are inserted in the scope of Wallet Issuing and the wallet ID is randomly drawn from the set of *unused* wallet IDs, i.e., they can not both have the same wallet ID. Contradiction!  $\square$

LEMMA F.5. *Within a tree of the Ideal Transaction Graph the PID  $pid_U$  of the corresponding user is constant.*

PROOF. Same proof as “ $\Leftarrow$ ” in the proof of Lemma F.4.  $\square$

In other words, Lemma F.5 states that a wallet (a tree in  $TRDB$ ) is always owned by a distinct user. But a user can own multiple wallets.

LEMMA F.6. *Within a tree of  $TRDB$ , every node  $trdb = (s^{prev}, s, \phi, x, \lambda, pid_U, pid_R, p, b)$  has depth  $x$  and all nodes of the same depth in the same tree have the same fraud detection ID  $\phi$ . Conversely, nodes with the same fraud detection ID are in the same tree and have the same depth within this tree.*

PROOF. Proof by Induction. The statement is true for the empty  $TRDB$ . In the scope of Wallet Issuing a new root node is inserted, Wallet Issuing sets  $x := 0$  and an unused  $\phi$  is chosen. In the scope of Debt Accumulation or Debt Clearance,  $x$  is calculated as  $x := x^{\text{prev}} + 1$ , where by induction  $x^{\text{prev}}$  is the depth of its predecessor. With respect to  $\phi$  we note that when inserted, every node gets as fraud detection ID the value stored in  $f_\phi(\lambda, x)$  which only depends on the node's wallet ID and depth. When this value is set (in either Wallet Issuing, Debt Accumulation, Debt Clearance or User Blacklisting) it is chosen from the set of *unused* fraud detection IDs and therefore unique for given  $\lambda$  and  $x$ .  $\square$

LEMMA F.7. Let  $trdb = (s^{\text{prev}}, s, \phi, x, \lambda, pid_{\mathcal{U}}, pid_{\mathcal{R}}, p, b)$  be an arbitrary but fixed node. If  $trdb$  is not a root let  $trdb^{\text{prev}} = (s^{\text{prev, prev}}, s^{\text{prev}}, \phi^{\text{prev}}, x^{\text{prev}}, \lambda, pid_{\mathcal{U}}, pid_{\mathcal{R}}^{\text{prev}}, p^{\text{prev}}, b^{\text{prev}})$  be its predecessor. Then  $b = b^{\text{prev}} + p$  holds for non-root nodes and  $p = \perp, b = 0$  for root nodes.

PROOF. Same induction argument as in proof of Lemma F.6.  $\square$

Before we start to show that  $\pi_{\text{P4TC}}$  correctly implements  $\mathcal{F}_{\text{P4TC}}$ , we make note of two additional simple statements about the ideal functionality itself.

- LEMMA F.8 (PROTECTION AGAINST FALSE ACCUSATION). (1) The task Double-Spending Detection returns a proof  $\pi \neq \perp$  if and only if the user has committed double-spending.  
 (2) The task Verify Guilt returns OK if and only if its input  $(pid_{\mathcal{U}}, \pi)$  has been output at a previous invocation of Double-Spending Detection.

PROOF. The first part obviously follows by the definition of the task Double-Spending Detection (cp. Fig. 14). Note for the second part that users are assumed to be honest. If  $f_\Pi(pid_{\mathcal{U}}, \pi)$  is undefined, then out is set to NOK and the result is recorded (cp. Steps 3 and 4 in Fig. 15). Guilt Verification only returns OK, if  $f_\Pi(pid_{\mathcal{U}}, \pi) = \text{OK}$  has already been defined (cp. Step 2). As (in case of honest users) Guilt Verification extends  $f_\Pi$  by nothing but invalid proofs, Double-Spending Detection exclusively sets  $f_\Pi(pid_{\mathcal{U}}, \pi) = \text{OK}$  (cp. Step 2, in Fig. 14).  $\square$

LEMMA F.9 (CORRECTNESS OF BLACKLISTING). Let  $\mathcal{U}$  be an arbitrary but fixed user with  $pid_{\mathcal{U}}$ . Under the assumption that  $\mathcal{U}$  participates in less than  $x_{bl_{\mathcal{R}}}$  transactions, i.e., in less than  $x_{bl_{\mathcal{R}}}$  invocations of Wallet Issuing, Debt Accumulation and Debt Clearance, the following two statements hold:

- (1) The set  $\Phi_{\mathcal{U}}$  returned to  $\mathcal{T}$  by User Blacklisting contains all fraud detection IDs that have ever been used by  $\mathcal{U}$ .  
 (2) Any invocation of Debt Accumulation for  $\mathcal{U}$  with input  $bl_{\mathcal{R}} = \Phi_{\mathcal{U}}$  aborts with message blacklisted.

PROOF. (1) Let  $TRDB_{\mathcal{U}} \subseteq TRDB$  be the subset of all transaction entries  $trdb = (\cdot, \cdot, \cdot, \cdot, \cdot, pid_{\mathcal{U}}, \cdot, p, \cdot)$  corresponding to  $pid_{\mathcal{U}}$  and let  $\mathcal{L}_{\mathcal{U}}$  denote the set of wallet IDs occurring in  $TRDB_{\mathcal{U}}$ . For  $\lambda \in \mathcal{L}_{\mathcal{U}}$  the depth of the tree associated to the wallet id  $\lambda$  is given by  $x_\lambda := \max\{x \mid f_\phi(\lambda, x) \neq \perp\}$ . If  $\mathcal{U}$  with  $pid_{\mathcal{U}}$  participated in less than  $x_{bl_{\mathcal{R}}}$  transaction, then the maximum depth  $x_{\max} = \max_{\lambda \in \mathcal{L}_{\mathcal{U}}} x_\lambda$  is smaller than  $x_{bl_{\mathcal{R}}}$ . The set of used fraud detection IDs is given by  $\{f_\phi(\lambda, x) \mid \lambda \in \mathcal{L}_{\mathcal{U}}, 0 \leq x \leq x_\lambda\}$  which is a subset of  $\Phi_{\mathcal{U}} := \{f_\phi(\lambda, x) \mid \lambda \in \mathcal{L}_{\mathcal{U}}, 0 \leq x \leq x_{bl_{\mathcal{R}}}\}$ .  
 (2) Let  $s^{\text{prev}}$  denote the serial number for which Debt Accumulation is invoked and let  $trdb^{\text{prev}} = (\cdot, s^{\text{prev}}, \phi^{\text{prev}}, x^{\text{prev}}, \lambda, \dots)$  be the corresponding transaction entry. By assumption  $x^{\text{prev}} < x_{bl_{\mathcal{R}}}$  holds. As User Blacklisting has previously been called,  $\phi = f_\phi(\lambda, x^{\text{prev}} + 1)$  is already fixed. Moreover,  $\phi \in \Phi_{\mathcal{U}} = bl_{\mathcal{R}}$  holds and thus Debt Accumulation aborts.  $\square$

After these preliminaries that only dealt with the ideal functionality on its own, we are now ready to argue why  $\pi_{\text{P4TC}}$  implements  $\mathcal{F}_{\text{P4TC}}$  correctly. Please note that to this extent all parties are assumed to be honest. The simulator  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{correct}}$  is depicted in Fig. 44. As before in Appendix E we assume encrypted and one-sidedly



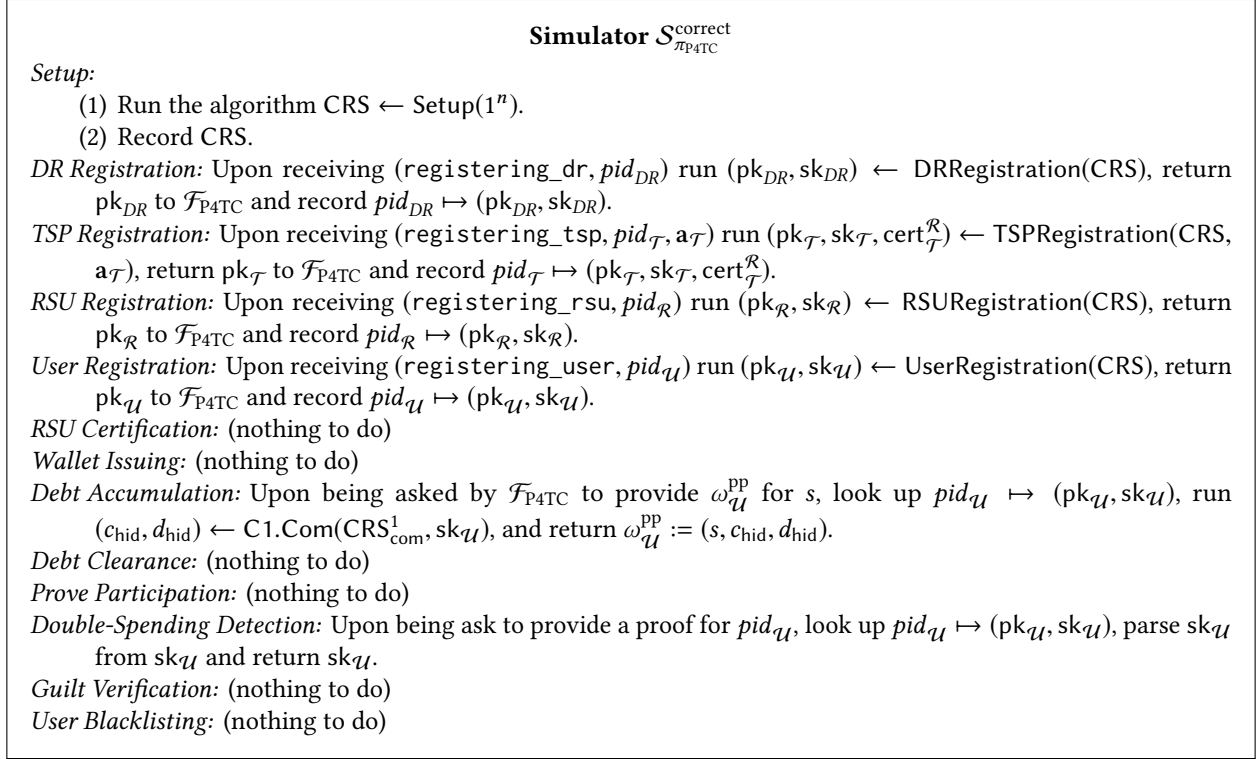


Fig. 44. The simulator for System Correctness

authenticated channels. Hence, the simulator is not required to simulate any messages, but only needs to provide the ideal function with implementation-specific values if being asked to do so. This also implies that the only option for the environment  $\mathcal{Z}$  to distinguish between the real and the ideal model is by means of different outputs from the different tasks. Luckily, either these tasks run independent from each other (e.g., registration) or their output is information-theoretically independent from any shared state (e.g., various OK-messages). Hence, only a handful of options remain and those can be handled individually.

LEMMA F.10 (REGISTRATION CORRECTNESS). *Let  $\mathcal{Z}^{\text{correct}}$  be an environment that does not corrupt any parties. Then the subroutine input/output of the main parties of  $\pi_{P4TC}$  and  $\mathcal{F}_{P4TC}$  resp. in the scope of the tasks DR/TSP/RSU/User Registration is perfectly indistinguishable between the experiments*

$$\begin{aligned} & \text{EXEC}_{\pi_{P4TC}^{\text{CRS}}, \bar{\mathcal{G}}_{\text{bb}}, \mathcal{S}_{\pi_{P4TC}}^{\text{correct}}, \mathcal{Z}^{\text{correct}}}(1^n) \quad \text{and} \\ & \text{EXEC}_{\mathcal{F}_{P4TC}, \bar{\mathcal{G}}_{\text{bb}}, \mathcal{S}_{\pi_{P4TC}}^{\text{correct}}, \mathcal{Z}^{\text{correct}}}(1^n). \end{aligned}$$

PROOF. Obvious.  $\mathcal{S}_{\pi_{P4TC}}^{\text{correct}}$  runs the real algorithms DRRegistration, TSPRegistration, RSURegistration, and UserRegistration in the ideal experiment.  $\square$

LEMMA F.11 (RSU CERTIFICATION CORRECTNESS). *Let  $\mathcal{Z}^{\text{correct}}$  be an environment that does not corrupt any parties. Under the assumption that  $S$  is correct, the subroutine input/output of the main parties of  $\pi_{P4TC}$  and  $\mathcal{F}_{P4TC}$*

resp. in the scope of the task *RSU Certification* is perfectly indistinguishable between the experiments

$$\begin{aligned} & \text{EXEC}_{\pi_{P4TC}^{\mathcal{F}_{CRS}}, \overline{\mathcal{G}}_{bb}, \mathcal{S}_{\pi_{P4TC}}^{\text{correct}}, \mathcal{Z}^{\text{correct}}}(1^n) \quad \text{and} \\ & \text{EXEC}_{\mathcal{F}_{P4TC}, \overline{\mathcal{G}}_{bb}, \mathcal{S}_{\pi_{P4TC}}^{\text{correct}}, \mathcal{Z}^{\text{correct}}}(1^n). \end{aligned}$$

PROOF. Obvious. In the ideal model the TSP inputs  $\mathbf{a}_R$  and the RSU outputs  $\mathbf{a}_R$  again. In the real model, the TSP sends  $\mathbf{a}_R$  to the RSU together with a signature. As  $\mathcal{S}$  is correct and all parties honest, the RSU does not abort.  $\square$

LEMMA F.12 (WALLET ISSUING CORRECTNESS). *Let  $\mathcal{Z}^{\text{correct}}$  be an environment that does not corrupt any parties and does not conduct blacklisting.<sup>30</sup> Under the assumption that all building blocks are correct and PRF is a pseudo-random function, the subroutine input/output of the main parties of  $\pi_{P4TC}$  and  $\mathcal{F}_{P4TC}$  resp. in the scope of the task *Wallet Issuing* is perfectly indistinguishable between the experiments*

$$\begin{aligned} & \text{EXEC}_{\pi_{P4TC}^{\mathcal{F}_{CRS}}, \overline{\mathcal{G}}_{bb}, \mathcal{S}_{\pi_{P4TC}}^{\text{correct}}, \mathcal{Z}^{\text{correct}}}(1^n) \quad \text{and} \\ & \text{EXEC}_{\mathcal{F}_{P4TC}, \overline{\mathcal{G}}_{bb}, \mathcal{S}_{\pi_{P4TC}}^{\text{correct}}, \mathcal{Z}^{\text{correct}}}(1^n). \end{aligned}$$

PROOF. In the ideal model the TSP inputs  $\mathbf{a}_U$  which is output by the user. In the real model the TSP sends  $\mathbf{a}_U$  together with a signature. As  $\mathcal{S}$  is correct and all parties honest, the user does not abort. In the ideal model the user outputs  $s$  which is uniformly drawn. In the real model  $s$  is drawn by means of a Blum coin toss.  $\square$

Until now, all previous lemmas have been rather trivial which is likely one of the reasons why correctness is usually not considered separately in UC proofs. Generally, it is quite obvious that a real protocol does what it is expected to do. The following lemma is slightly more involved and the main reason why we decided to show correctness first.

We start with more notation and augment the entries  $\overline{trdb}$  of the Ideal Transaction Graph  $TRDB$  with additional information. An *Augmented Transaction Entry*  $\overline{trdb}$  has the form


$$\begin{aligned} \overline{trdb} = & (s^{\text{prev}}, s, \phi, x, \lambda, pid_U, pid_R, p, b, \\ & c_{\mathcal{T}}^{\text{in}}, d_{\mathcal{T}}^{\text{in}}, M_{\mathcal{T}}^{\text{in}}, c_{\mathcal{R}}^{\text{in}}, d_{\mathcal{R}}^{\text{in}}, M_{\mathcal{R}}^{\text{in}}, \\ & c_{\mathcal{T}}^{\text{out}}, d_{\mathcal{T}}^{\text{out}}, M_{\mathcal{T}}^{\text{out}}, c_{\mathcal{R}}^{\text{out}}, d_{\mathcal{R}}^{\text{out}}, M_{\mathcal{R}}^{\text{out}}) \end{aligned} \quad (9)$$

with  $c$ ,  $d$  and  $M$  with equal suffixes denoting a commitment, its decommitment information and the opening in the implicit message space. At the beginning of a transaction in the scope of Debt Accumulation or Debt Clearance the user loads his token  $\tau^{\text{prev}}$  which contains two commitments  $c_{\mathcal{T}}$  and  $c_{\mathcal{R}}^{\text{prev}}$ , randomizes the commitments and at the end the user possesses two updated commitments  $c_{\mathcal{T}}$ ,  $c_{\mathcal{R}}$  which are stored in  $\tau$  again. We call the initial commitments the *in*-commitments of the transaction and the resulting commitments the *out*-commitments.

*Definition F.13 (Augmented Transaction Graph (informal)).* The set  $\overline{TRDB} = \{\overline{trdb}_i\}$  with  $\overline{trdb}_i$  defined as in Eq. (9) is called the *Augmented Transaction Graph*. It inherits the graph structure of the Ideal Transaction Graph and augments each edge by additional labels, called the *in*-commitments and *out*-commitments.

Two remarks are in order: Firstly, none of the (commitment, decommitment, message)-triples is neither completely received not sent by the RSU or TSP, respectively. The RSU receives a randomized version of the in-commitment and no decommitment at all. In the reverse direction, the RSU sends the out-commitment and a share of the decommitment. The complete triples only exist inside the user's token. Secondly, it is tempting but misleading to assume that  $c_{\mathcal{R}}^{\text{in}} = c_{\mathcal{R}}^{\text{prev}}$  (or similar equations) hold. Note that we do not make any of these assumptions for the definition. Hence we decided on a new notion and coined the term in-/out-commitments

<sup>30</sup>The case that *Wallet Issuing* aborts due to blacklisting is postponed to Lemma F.20.

Fig. 45. An entry  $\overline{trdb} \in \overline{TRDB}$  visualized as an element of a directed graph

instead of re-using the term “previous commitment”. Actually, these kind of equalities is what we have to show. While this is rather easy to show in the scope of the proof of correctness if all parties are honest, these equalities are not obviously true as soon as corrupted parties are considered.

The augmented transaction information is depicted in Fig. 45.

The next proposition will be used several times below and maps an execution of the real protocol one-to-one and onto an execution of the ideal functionality. In order to be formally correct we do not only claim that an Augmented Transaction Graph exists mathematically, but also show how it can be obtained computationally. To this end we describe a UC-like experiment  $\text{EXEC}_{\mathcal{F}_{\text{P4TC}}, \overline{\mathcal{G}}_{\text{bb}}, \mathcal{S}^*, \mathcal{Z}^*}(1^n)$ . We consider a special environment  $\mathcal{Z}^*$  that corrupts all parties but plays them honestly, i.e.,  $\mathcal{Z}^*$  is an honest adversary. This is a technical requirement such that the simulator  $\mathcal{S}^*$  is able to “see” the messages.  $\mathcal{S}^*$  simulates a CRS such that the ZK-proofs become extractable and otherwise behaves like the dummy adversary. Please note that  $\mathcal{S}^*$  does not need to do anything with respect to  $\mathcal{Z}^*$  except message delivery, as  $\mathcal{Z}^*$  plays with itself. In addition to the dummy adversary  $\mathcal{S}^*$  performs some bookkeeping (not required for simulation). It keeps track of all interactions and maintains a shadow copy of the Transaction Graph perfectly in sync with the Ideal Transaction Graph that  $\mathcal{F}_{\text{P4TC}}$  maintains internally. Moreover,  $\mathcal{S}^*$  augments its own copy of the Transaction Graph with the commitments it sees.  $\mathcal{S}^*$  can obtain the in-commitments via extraction of the ZK-proofs, and the out-commitments are sent by the TSP/RSU in the clear.

PROPOSITION F.14. Consider the experiment

$$\text{EXEC}_{\mathcal{F}_{\text{P4TC}}, \overline{\mathcal{G}}_{\text{bb}}, \mathcal{S}^*, \mathcal{Z}^*}(1^n)$$

as defined earlier, and the Augmented Transaction Graph  $\overline{TRDB}$  corresponding to this experiment.

- (1) Let  $\overline{trdb}^{\text{prev}} = (s^{\text{prev, prev}}, s^{\text{prev}}, \dots)$  and  $\overline{trdb} = (s^{\text{prev}}, s, \dots)$  be two transaction entries that are connected by an edge. Then  $c_{\mathcal{T}}^{\text{out}^{\text{prev}}} = c_{\mathcal{T}}^{\text{in}}$  and  $c_{\mathcal{R}}^{\text{out}^{\text{prev}}} = c_{\mathcal{R}}^{\text{in}}$  hold.
- (2) Let  $\overline{trdb}^* = (s^{\text{prev, *}}, s^*, \dots)$  and  $\overline{trdb} = (s^{\text{prev}}, s, \dots)$  be two transaction entries with  $c_{\mathcal{T}}^{\text{out}^*} = c_{\mathcal{T}}^{\text{in}}$  and  $c_{\mathcal{R}}^{\text{out}^*} = c_{\mathcal{R}}^{\text{in}}$ . Then  $s^* = s^{\text{prev}}$  holds with overwhelming probability.

In other words, the in-/out-commitments also induce a graph structure on the Augmented Transaction Graph and for an honest execution of the protocol this structure coincides with the graph structure inherited from the Ideal Transaction Graph by means of serial numbers  $s$ .

- PROOF. (1) Follows from the fact that the users are honest. At the end of some transaction with serial number  $s^{\text{prev}}$  the user obtains the out-commitments  $c_{\mathcal{T}}^{\text{out}^{\text{prev}}}$ ,  $c_{\mathcal{R}}^{\text{out}^{\text{prev}}}$  and records them locally. If the user invokes a new transaction with serial  $s$  and  $s^{\text{prev}}$  as its predecessor, he loads the recorded commitments  $c_{\mathcal{T}}^{\text{out}^{\text{prev}}}$ ,  $c_{\mathcal{R}}^{\text{out}^{\text{prev}}}$  and uses them as in-commitments  $c_{\mathcal{T}}^{\text{in}}$ ,  $c_{\mathcal{R}}^{\text{in}}$ .
- (2) Follows from the fact that the out-commitments are honestly generated and therefore unique with overwhelming probability. Assume the contrary, i.e.,  $c_{\mathcal{T}}^{\text{out}^*} = c_{\mathcal{T}}^{\text{in}}$  and  $c_{\mathcal{R}}^{\text{out}^*} = c_{\mathcal{R}}^{\text{in}}$  hold but  $s^* \neq s^{\text{prev}}$ . This can only happen due to one of two reasons:
- (a) At the beginning of a transaction the environment uses an in-commitment that is not the out-commitment of the preceding transaction. This contradicts the assumption that the user behaves honestly.

- (b) At the end of some transaction an out-commitment is accidentally generated that equals an in-commitment elsewhere. Again, as the out-commitments are honestly generated, this happens only with negligible probability.  $\square$

With [Proposition F.14](#) we can show the next lemma on our way to proving correctness.

LEMMA F.15 (DEBT ACCUMULATION CORRECTNESS). *Let  $\mathcal{Z}^{correct}$  be an environment that does not corrupt any parties and does not conduct blacklisting.<sup>31</sup> Under the assumption that all building blocks are correct, C1 is homomorphic, and PRF is a pseudo-random function the subroutine input/output of the main parties of  $\pi_{P4TC}$  and  $\mathcal{F}_{P4TC}$  resp. in the scope of the task Debt Accumulation is computationally indistinguishable between the experiments*

$$\begin{aligned} & \text{EXEC}_{\pi_{P4TC}, \overline{\mathcal{G}}_{bb}, \mathcal{S}_{\pi_{P4TC}}^{correct}, \mathcal{Z}^{correct}}(1^n) \quad \text{and} \\ & \text{EXEC}_{\mathcal{F}_{P4TC}, \overline{\mathcal{G}}_{bb}, \mathcal{S}_{\pi_{P4TC}}^{correct}, \mathcal{Z}^{correct}}(1^n). \end{aligned}$$

PROOF. With respect to the output of  $s$ ,  $\mathbf{a}_{\mathcal{R}}$  and  $\mathbf{a}_{\mathcal{U}}$  we refer to the proofs of [Lemmas F.11](#) and [F.12](#). Here, the same arguments hold. It remains to show that the output with respect to  $b$ ,  $p$  is indistinguishable. Consider the path in the Augmented Transaction Graph  $\overline{TRDB}$  from the root to this invocation of the Debt Accumulation task. Let  $\overline{trdb}^{prev}$  and  $trdb$  be two arbitrary but fixed transaction entries with  $\overline{trdb}^{prev}$  being the predecessor of  $trdb$ . Due to [Proposition F.14](#)  $c_{\mathcal{T}}^{out^{prev}} = c_{\mathcal{T}}^{in}$  and  $c_{\mathcal{R}}^{out^{prev}} = c_{\mathcal{R}}^{in}$  holds. Moreover,  $c_{\mathcal{R}}^{in}$  is updated to an out-commitment  $c_{\mathcal{R}}^{out}$  using the homomorphism of the commitment scheme C1 that is consistent with the modification of the semantic information in the ideal model, i.e.,  $b = b^{prev} + p$ .  $\square$

LEMMA F.16 (DEBT CLEARANCE CORRECTNESS). *Let  $\mathcal{Z}^{correct}$  be an environment that does not corrupt any parties. Under the assumption that all building blocks are correct, C1 is homomorphic, and PRF is a pseudo-random function the subroutine input/output of the main parties of  $\pi_{P4TC}$  and  $\mathcal{F}_{P4TC}$  resp. in the scope of the task Debt Clearance is computationally indistinguishable between the experiments*

$$\begin{aligned} & \text{EXEC}_{\pi_{P4TC}, \overline{\mathcal{G}}_{bb}, \mathcal{S}_{\pi_{P4TC}}^{correct}, \mathcal{Z}^{correct}}(1^n) \quad \text{and} \\ & \text{EXEC}_{\mathcal{F}_{P4TC}, \overline{\mathcal{G}}_{bb}, \mathcal{S}_{\pi_{P4TC}}^{correct}, \mathcal{Z}^{correct}}(1^n). \end{aligned}$$

PROOF. Uses the same arguments as the proof of [Lemma F.15](#).  $\square$

The remaining [Lemmas F.17](#) to [F.20](#) deal with correctness of the auxiliary tasks Prove Participation, Double-Spending Detection, Guilt Verification, and User Blacklisting. The proofs are actually neither complicated nor give any technical insight, because correctness simply follows from comparing the code of the real implementation with the pseudo-code of the ideal functionality under the assumption that the building blocks are correct. We merely include those proofs for the sake of completeness.

LEMMA F.17 (PROVE PARTICIPATION CORRECTNESS). *Let  $\mathcal{Z}^{correct}$  be an environment that does not corrupt any parties. Under the assumption that all building blocks are correct, the subroutine input/output of the main parties of  $\pi_{P4TC}$  and  $\mathcal{F}_{P4TC}$  resp. in the scope of the task Prove Participation is indistinguishable between the experiments*

$$\begin{aligned} & \text{EXEC}_{\pi_{P4TC}, \overline{\mathcal{G}}_{bb}, \mathcal{S}_{\pi_{P4TC}}^{correct}, \mathcal{Z}^{correct}}(1^n) \quad \text{and} \\ & \text{EXEC}_{\mathcal{F}_{P4TC}, \overline{\mathcal{G}}_{bb}, \mathcal{S}_{\pi_{P4TC}}^{correct}, \mathcal{Z}^{correct}}(1^n). \end{aligned}$$

<sup>31</sup>The case that Debt Accumulation aborts due to blacklisting is postponed to [Lemma F.20](#).

PROOF. The statement obviously follows by comparison of what the ideal functionality and the real implementation do. If Prove Participation returns OK in the ideal experiment, then there exists  $s$  s.t.  $(\cdot, s, \cdot, x, \lambda, pid_{\mathcal{U}}, \cdot, \cdot, \cdot) \in TRDB$ . This means the respective user has successfully participated in the  $x$ th transaction for a wallet with ID  $\lambda$  in the scope of Debt Accumulation. In the real experiment the user sends an honestly generated commitment  $c_{hid}$  which is recorded for  $s$  by both parties. During Prove Participation this commitment is re-sent and successfully opened. The reverse direction is likewise obvious.  $\square$

LEMMA F.18 (DOUBLE-SPENDING DETECTION CORRECTNESS). *Let  $\mathcal{Z}^{correct}$  be an environment that does not corrupt any parties. Under the assumption that all building blocks are correct, the subroutine input/output of the main parties of  $\pi_{P4TC}$  and  $\mathcal{F}_{P4TC}$  resp. in the scope of the task Double-Spending Detection is indistinguishable between the experiments*

$$\begin{aligned} & \text{EXEC}_{\pi_{P4TC}, \overline{\mathcal{G}}_{bb}, \mathcal{S}_{\pi_{P4TC}}^{correct}, \mathcal{Z}^{correct}}(1^n) \quad \text{and} \\ & \text{EXEC}_{\mathcal{F}_{P4TC}, \overline{\mathcal{G}}_{bb}, \mathcal{S}_{\pi_{P4TC}}^{correct}, \mathcal{Z}^{correct}}(1^n). \end{aligned}$$

PROOF. In the ideal model Double-Spending Detection returns the correct  $sk_{\mathcal{U}}$  for  $pid_{\mathcal{U}}$  in case of a fraud by definition of  $\mathcal{S}_{\pi_{P4TC}}^{correct}$ . In the real model Double-Spending Detection returns  $sk_{\mathcal{U}}$  with overwhelming probability. As the honest user is committed to the same  $u_1$ , the RSU picks two different values  $u_2, u'_2$  which result into two different points  $(u_2, t), (u'_2, t')$  with overwhelming probability. This means the linear equation system has a unique solution which equals  $sk_{\mathcal{U}}$ .  $\square$

LEMMA F.19 (GUILT VERIFICATION CORRECTNESS). *Let  $\mathcal{Z}^{correct}$  be an environment that does not corrupt any parties. Under the assumption that all building blocks are correct, and DLOG is a hard problem for  $G_1$ , the subroutine input/output of the main parties of  $\pi_{P4TC}$  and  $\mathcal{F}_{P4TC}$  resp. in the scope of the task Guilt Verification is indistinguishable between the experiments*

$$\begin{aligned} & \text{EXEC}_{\pi_{P4TC}, \overline{\mathcal{G}}_{bb}, \mathcal{S}_{\pi_{P4TC}}^{correct}, \mathcal{Z}^{correct}}(1^n) \quad \text{and} \\ & \text{EXEC}_{\mathcal{F}_{P4TC}, \overline{\mathcal{G}}_{bb}, \mathcal{S}_{\pi_{P4TC}}^{correct}, \mathcal{Z}^{correct}}(1^n). \end{aligned}$$

PROOF. In the ideal model Guilt Verification returns OK if and only if  $\pi$  has been an output of Double-Spending Detection (by Lemma F.8) and thus  $g_1^\pi = pk_{\mathcal{U}}$  holds (by definition of  $\mathcal{S}_{\pi_{P4TC}}^{correct}$ , cp. Fig. 44). In the real model Guilt Verification returns OK if and only if  $g_1^\pi = pk_{\mathcal{U}}$ . Hence, the only case in which the ideal and the real model differ is if  $g_1^\pi = pk_{\mathcal{U}}$  holds but  $\pi$  is not an output of Double-Spending Detection. This event is negligible, as DLOG is assumed to be a hard problem for  $G_1$ .  $\square$

LEMMA F.20 (USER BLACKLISTING CORRECTNESS). *Let  $\mathcal{Z}^{correct}$  be an environment that does not corrupt any parties. Assume that all building blocks are correct and PRF is a pseudo-random function.*

- (1) *If no user uses his wallet more than  $x_{bl_R}$  times, the subroutine input/output of the main parties of  $\pi_{P4TC}$  and  $\mathcal{F}_{P4TC}$  resp. in the scope of the task User Blacklisting is indistinguishable between the experiments*

$$\begin{aligned} & \text{EXEC}_{\pi_{P4TC}, \overline{\mathcal{G}}_{bb}, \mathcal{S}_{\pi_{P4TC}}^{correct}, \mathcal{Z}^{correct}}(1^n) \quad \text{and} \\ & \text{EXEC}_{\mathcal{F}_{P4TC}, \overline{\mathcal{G}}_{bb}, \mathcal{S}_{\pi_{P4TC}}^{correct}, \mathcal{Z}^{correct}}(1^n). \end{aligned}$$

- (2) *The abort behavior of  $\pi_{P4TC}$  and  $\mathcal{F}_{P4TC}$  resp. in the scope of the task Debt Accumulation is computationally indistinguishable between the experiments*

$$\begin{aligned} & \text{EXEC}_{\pi_{P4TC}, \overline{\mathcal{G}}_{bb}, \mathcal{S}_{\pi_{P4TC}}^{correct}, \mathcal{Z}^{correct}}(1^n) \quad \text{and} \\ & \text{EXEC}_{\mathcal{F}_{P4TC}, \overline{\mathcal{G}}_{bb}, \mathcal{S}_{\pi_{P4TC}}^{correct}, \mathcal{Z}^{correct}}(1^n). \end{aligned}$$

PROOF. (1) In the ideal model, the output of User Blacklisting depends on the transaction graph  $TRDB$ . For this reason, we proceed by a proof of induction for a fixed but arbitrary PID  $pid_{\mathcal{U}}$ . There are only three tasks, namely Wallet Issuing, Debt Accumulation and Debt Clearance, that modify  $TRDB$ . Assume that  $\mathcal{Z}^{\text{correct}}$  has not triggered any of these tasks yet for  $pid_{\mathcal{U}}$ . Then, in the ideal model  $TRDB_{\mathcal{U}}$  is empty and User Blacklisting returns  $b^{\text{bill}} = 0$  and  $\Phi_{\mathcal{U}} = \emptyset$ . Likewise in the real model the set of hidden user trapdoors  $HTD$  is empty and thus  $b^{\text{bill}} = 0$  and  $\Phi_{\mathcal{U}} = \emptyset$  is returned.

Assume that correctness holds for  $n$  transactions before User Blacklisting is called and we consider an execution with  $n + 1$  previous transactions for the fixed  $pid_{\mathcal{U}}$ . We look at the three cases of Wallet Issuing, Debt Accumulation and Debt Clearance separately.

- (a) *The  $(n + 1)^{\text{th}}$  transaction is Wallet Issuing:* In the ideal model a new entry  $trdb^*$  with a unique  $\lambda^*$  and  $p^* = 0$  is inserted into  $TRDB$ , i.e.,  $b^{\text{bill}}$  does not change. In the real model  $b^{\text{bill}}$  does not change either, because Wallet Issuing does not insert a  $\omega^{\text{bl}}$  into  $\Omega^{\text{bl}}$ . In the ideal model the partial mapping  $f_{\phi}$  is extended onto the domain  $\{(\lambda^*, 0), \dots, (\lambda^*, x_{\text{bl}_{\mathcal{R}}})\}$  and mapped to uniformly drawn fraud detection IDs. The output  $\Phi_{\mathcal{U}}$  increases by  $x_{\text{bl}_{\mathcal{R}}} + 1$  elements as the probability to pick the same element twice is negligible. In the real model Wallet Issuing inserts a new element  $htd$  into  $HTD$  which is sent to  $DR$  and—by assumption that every party is honest—verifies correctly and is extracted to a unique  $\lambda^*$ .<sup>32</sup> Consequently,  $DR$  inserts  $x_{\text{bl}_{\mathcal{R}}} + 1$  elements  $\{\text{PRF}(\lambda^*, 0), \dots, \text{PRF}(\lambda^*, x_{\text{bl}_{\mathcal{R}}})\}$  into  $\Phi_{\mathcal{U}}$ . Under the assumption that PRF is a pseudo-random function, these elements are indistinguishable from those drawn in the ideal experiment.
- (b) *The  $(n + 1)^{\text{th}}$  transaction is Debt Accumulation:* In the ideal model a new entry  $trdb^*$  with fraud detection ID  $\phi^*$  and price  $p^*$  is inserted into  $TRDB$ , i.e.,  $b^{\text{bill}}$  increases by  $p^*$ . In the real model  $\mathcal{R}$  outputs  $\omega^{\text{bl}*} = (\phi^*, p^*)$  with  $\phi^* = \text{PRF}(\lambda^*, x^*)$  at the end of Debt Accumulation. By assumption that the wallet has been generated honestly, a hidden trapdoor  $htd^*$  that can be extracted to  $\lambda^*$  is contained in  $HTD$ . By assumption that the user does not use his wallet more than  $x_{\text{bl}_{\mathcal{R}}}$  times,  $x^* \leq x_{\text{bl}_{\mathcal{R}}}$  holds. Hence,  $\phi^* \in \Phi_{\mathcal{U}}$  and  $\omega^{\text{bl}*} \in \Omega_{\mathcal{U}}^{\text{bl}}$  follows whereby  $b^{\text{bill}}$  increases by  $p^*$ .
- (c) *The  $(n + 1)^{\text{th}}$  transaction is Debt Clearance:* Same argument as for Debt Accumulation.
- (2) Let  $s^{\text{prev}}$  denote the serial number for which Debt Accumulation is invoked and let

$$\overline{trdb}^{\text{prev}} = (\cdot, s^{\text{prev}}, \phi^{\text{prev}}, x^{\text{prev}}, \lambda, \dots)$$

be the corresponding transaction entry. We distinguish two cases:

- (a) *In the ideal model  $f_{\phi}(\lambda, x^{\text{prev}} + 1)$  has already been defined:* This can happen due to two cases; either there is a  $\overline{trdb}^*$  with  $\lambda^* = \lambda$  and  $x^* = x^{\text{prev}} + 1$  or User Blacklisting has been invoked and  $x^{\text{prev}} < x_{\text{bl}_{\mathcal{R}}}$  holds. In either case the fraud detection ID  $\phi = f_{\phi}(\lambda, x^{\text{prev}} + 1)$  associated with the considered transaction has already been fixed and given as part of some output to  $\mathcal{Z}^{\text{correct}}$ . In the real model  $\phi = \text{PRF}(\lambda, x^{\text{prev}} + 1)$  has likewise been output to  $\mathcal{Z}^{\text{correct}}$ . Debt Accumulation aborts with `blacklisted`, if and only if  $\phi$ <sup>33</sup> is contained in  $bl_{\mathcal{R}}$ .
- (b) *In the ideal model  $f_{\phi}(\lambda, x^{\text{prev}} + 1)$  has not yet been defined:* In the ideal model a fresh  $\phi$  is uniformly drawn, the probability that  $\phi \in bl_{\mathcal{R}}$  holds is negligible and thus Debt Accumulation negligibly aborts with `blacklisted`. In the real model,  $\phi$  is set as  $\text{PRF}(\lambda, x + 1)$  and  $\phi$  has never been output before. If  $\mathcal{Z}^{\text{correct}}$  inputs a  $bl_{\mathcal{R}}$  such that `blacklisted` occurs with non-negligible probability this immediately yields an efficient adversary  $\mathcal{B}$  against the unpredictability of PRF. Assume such a  $\mathcal{Z}^{\text{correct}}$  exists, then  $\mathcal{B}$  can be constructed as follows:  $\mathcal{B}$  internally executes the real experiment, i.e., it runs  $\mathcal{Z}^{\text{correct}}$  and plays all

<sup>32</sup>N.b.: By abuse of notation we use the same variable here as in the ideal experiment, but of course they are not necessarily equal—only computationally indistinguishable over the randomness of the experiment.

<sup>33</sup>N.b.: By abuse of notation we use the same variable for the real and the ideal experiment, but of course they are not necessarily equal—only computationally indistinguishable over the randomness of the experiment

parties and the adversary for  $\mathcal{Z}^{\text{correct}}$  honestly in its head. Externally  $\mathcal{B}$  plays the unpredictability game.  $\mathcal{B}$  needs to guess for which wallet  $\mathcal{Z}^{\text{correct}}$  eventually causes an abort. In the scope of Wallet Issuing  $\mathcal{B}$  picks an unused seed  $\tilde{\lambda}$  but queries its external PRF-oracle for  $x = 0$ . N.b.,  $\mathcal{B}$  never uses  $\tilde{\lambda}$ , it is only a placeholder for the unknown challenge  $\lambda^C$  used by the external oracle. Whenever  $\mathcal{B}$  needs to evaluate  $\text{PRF}(\tilde{\lambda}, x)$  it queries its external oracle instead. Note, that  $\mathcal{Z}^{\text{correct}}$  cannot distinguish, if it is executed within a reduction, because  $\tilde{\lambda}$  and  $\lambda^C$  are both randomly chosen.  $\mathcal{B}$  guesses the interaction in which  $\mathcal{Z}^{\text{correct}}$  causes the abort, takes the blacklist  $bl_R$  from  $\mathcal{Z}^{\text{correct}}$ , guesses what element of the blacklist is the next image of PRF and outputs this externally.  $\square$

**THEOREM F.21 (SYSTEM CORRECTNESS).** *Let  $\mathcal{Z}^{\text{correct}}$  be an environment that does not corrupt any parties. Under the assumption that all building blocks are correct, C1 is homomorphic, PRF is a pseudo-random function, and DLOG is a hard problem for  $G_1$*

$$\pi_{\text{P4TC}}^{\mathcal{F}_{\text{CRS}}, \overline{\mathcal{G}}_{\text{bb}}} \geq_{\text{UC}} \mathcal{F}^{\overline{\mathcal{G}}_{\text{bb}}}$$

holds. In other words,  $\pi_{\text{P4TC}}^{\mathcal{F}_{\text{CRS}}, \overline{\mathcal{G}}_{\text{bb}}}$  is a correct implementation of  $\mathcal{F}^{\overline{\mathcal{G}}_{\text{bb}}}$ .

**PROOF.** A direct consequence of **Lemmas F.10 to F.12** and **F.15 to F.20**.  $\square$

## F.2 Proof of System Security

In this section we show the following theorem.

**THEOREM F.22 (SYSTEM SECURITY).** *Under the assumptions of [Theorem F.1](#) and static corruption of a subset of the users*

$$\pi_{\text{P4TC}}^{\mathcal{F}_{\text{CRS}}, \overline{\mathcal{G}}_{\text{bb}}} \geq_{\text{UC}} \mathcal{F}^{\overline{\mathcal{G}}_{\text{bb}}}$$

holds.

The definition of the UC-simulator  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$  for [Theorem F.22](#) can be found in [Figs. 46 to 49](#). Please note that while the real protocol  $\pi_{\text{P4TC}}$  lives in the  $(\mathcal{F}_{\text{CRS}}, \overline{\mathcal{G}}_{\text{bb}})$ -model the ideal functionality  $\mathcal{F}_{\text{P4TC}}$  has no CRS. Hence, the CRS (but not the bulletin board) is likewise simulated by  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$ , giving it a lever to extract the ZK proofs P1, P2, and P3 and to equivocate the commitment C2. Please note that the simulator records more information during its execution—namely the Augmented Transaction Graph (cp. [Definition F.13](#)) and the set of simulated proof-of-participation information  $\overline{\Omega}_{\mathcal{U}}^{\text{pp}}$ —than is technically required by the simulator for an indistinguishable simulation from the viewpoint of the environment. This additional information only simplifies the presentation of the reduction argument within the security proof.

By skipping ahead we shortly explain how this additional tracking of information supports the proof by giving a concrete example for  $\overline{\Omega}_{\mathcal{U}}^{\text{pp}}$  and the binding property of the commitment scheme. Each of the upcoming security proofs roughly follows the same lines of argument. If an environment  $\mathcal{Z}$  can efficiently distinguish between the real and the ideal experiment, then we can construct an efficient adversary  $\mathcal{B}$  against one of the underlying cryptographic building blocks, i.e., the binding property of the commitment scheme for the example at hand. To this end,  $\mathcal{B}$  plays the adversary against the binding property in the outer game and internally executes the UC-experiment in its head while mimicking the role of the simulator. It is important to note that although  $\mathcal{B}$  emulates the environment and the ideal functionality internally, it only has *black-box access* to them. In other words, although everything happens inside “the head of  $\mathcal{B}$ ” it cannot somehow magically extract  $\mathcal{Z}$ ’s attack strategy. Moreover, in order for  $\mathcal{B}$  to win the external game against the binding property the mere existence of a break does not suffice, but  $\mathcal{B}$  actually has to present two different openings for the same commitment. In order to do so,  $\mathcal{B}$  observes all of  $\mathcal{Z}$ ’s communication and systematically records parts of the observed transcript such

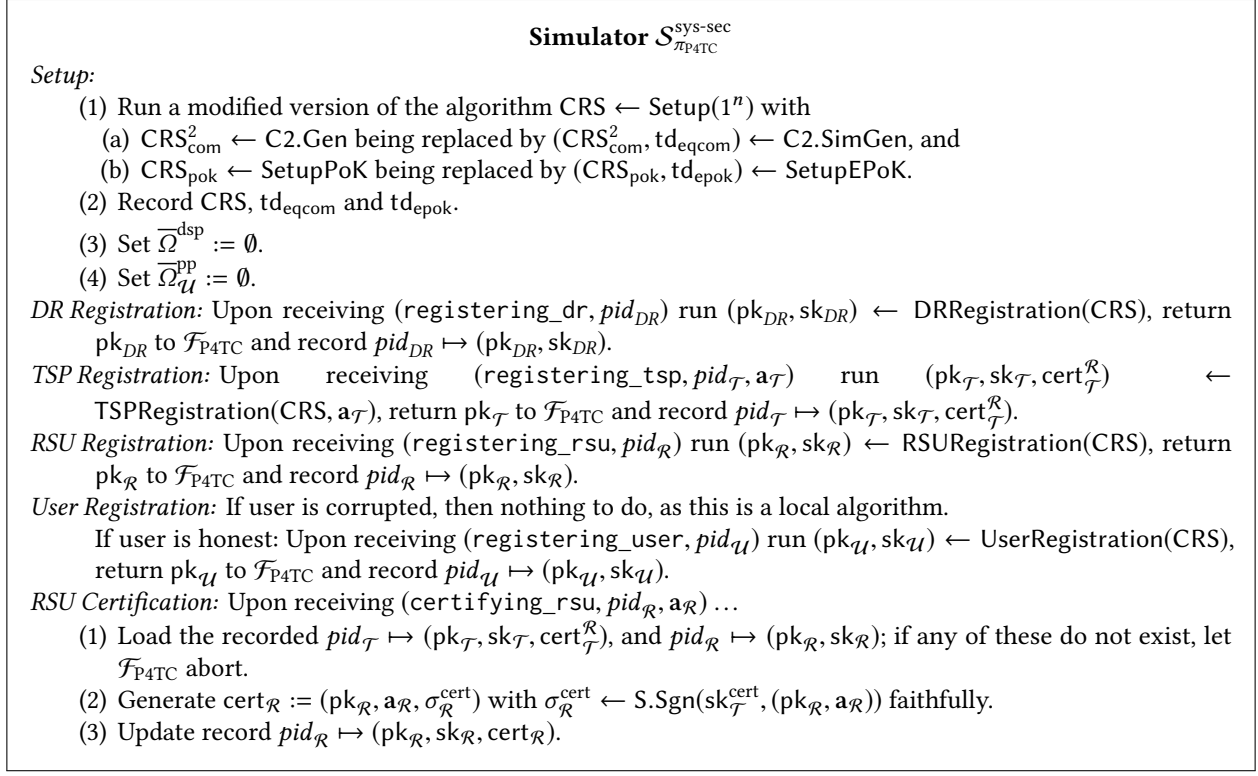


Fig. 46. The simulator for System Security

that it can efficiently look up a commitment and its (former) opening if the same commitment with a different opening occurs a second time. This is the only reason why the simulator additionally records the observed opening messages in  $\overline{\mathcal{Q}}_{\mathcal{U}}^{\text{pp}}$  but never uses them for its own simulation. These opening are only required in the outer game during the reduction argument.

As already stated in [Appendix E](#) we assume half-authenticated secure channels. In this model the adversary only learns the length of the messages being sent between honest parties. Secure channels can be realized by any IND-CCA secure encryption scheme (cp. [21]). Secure channels release us from the burden to spell out a simulator for messages between honest parties.

First, note the following simple observation.

**LEMMA F.23.** *The messages simulated by  $\mathcal{S}_{\pi_{P4TC}}^{\text{sys-sec}}$  in the name of honest parties and sent to  $\mathcal{Z}^{\text{sys-sec}}$  are perfect, i.e., statistically identical to real messages.*

**PROOF.** Follows from the definition of  $\mathcal{S}_{\pi_{P4TC}}^{\text{sys-sec}}$  (cp. [Figs. 46 to 49](#)). □

The messages generated by  $\mathcal{S}_{\pi_{P4TC}}^{\text{sys-sec}}$  do not depend on any unknown secrets. Moreover, another simple observation is

**LEMMA F.24.** *Whenever any  $\mathcal{Z}^{\text{sys-sec}}$  triggers an abort of one of the tasks in the real experiment  $\text{EXEC}_{\pi_{P4TC}, \overline{\mathcal{G}}_{bb}, \mathcal{S}_{\pi_{P4TC}}^{\text{sys-sec}}, \mathcal{Z}^{\text{sys-sec}}}(1^n)$ , the simulator  $\mathcal{S}_{\pi_{P4TC}}^{\text{sys-sec}}$  lets  $\mathcal{F}_{P4TC}$  also abort in the ideal experiment  $\text{EXEC}_{\mathcal{F}_{P4TC}, \overline{\mathcal{G}}_{bb}, \mathcal{S}_{\pi_{P4TC}}^{\text{sys-sec}}, \mathcal{Z}^{\text{sys-sec}}}(1^n)$ .*



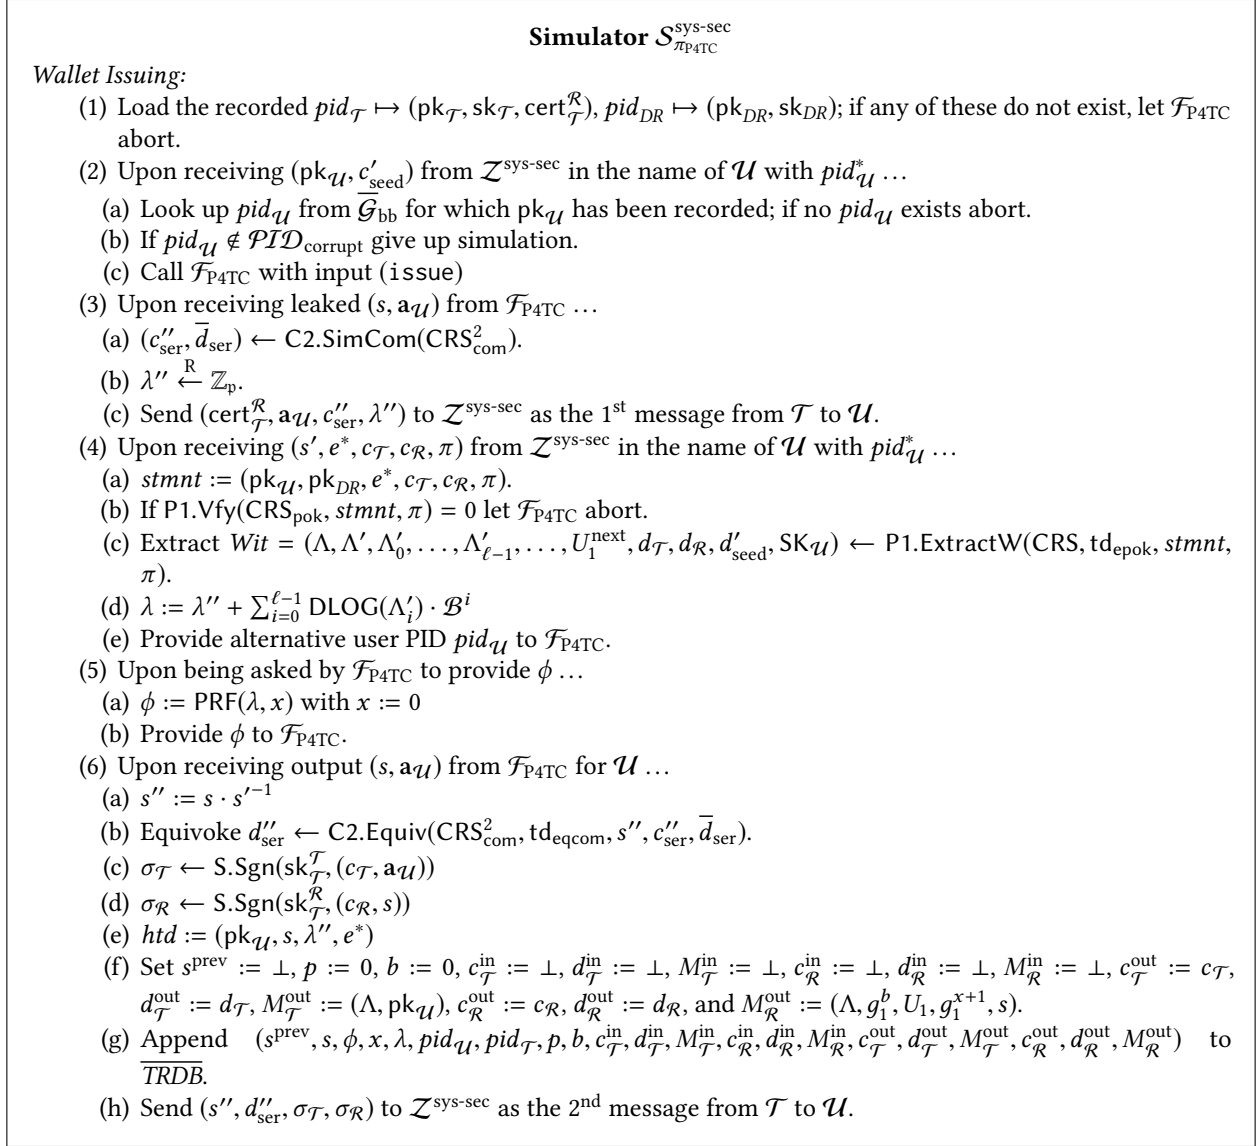


Fig. 47. The simulator for System Security (cont. from Fig. 46)

PROOF. Follows from the definition of  $\mathcal{S}_{\pi_{P4TC}}^{\text{sys-sec}}$  (cp. Figs. 46 to 49). □

An abort in the real model only happens if one of the messages is invalid, e.g., a signature does not verify or a ZK proof is not valid. If  $\mathcal{Z}^{\text{sys-sec}}$  playing the corrupted parties sends a message in the ideal experiment, the simulator  $\mathcal{S}_{\pi_{P4TC}}^{\text{sys-sec}}$  playing the honest receiver performs the same checks and lets  $\mathcal{F}_{P4TC}$  abort, if necessary.

Due to Lemmas F.23 and F.24 there is only a limited number of ways how  $\mathcal{Z}^{\text{sys-sec}}$  can distinguish between the real and the ideal experiment. All options fall into one out of two categories:

**Simulator  $\mathcal{S}_{\pi_{P4TC}}^{\text{sys-sec}}$  (cont.)**

*Debt Accumulation:*

- (1) Load the recorded  $pid_{\mathcal{T}} \mapsto (pk_{\mathcal{T}}, sk_{\mathcal{T}}, cert_{\mathcal{T}}^{\mathcal{R}})$ , and  $pid_{\mathcal{R}} \mapsto (pk_{\mathcal{R}}, sk_{\mathcal{R}}, cert_{\mathcal{R}})$ ; if any of these do not exist, let  $\mathcal{F}_{P4TC}$  abort.
- (2) Pick  $u_2 \xleftarrow{R} \mathbb{Z}_p$ .
- (3)  $(c''_{\text{ser}}, \bar{d}_{\text{ser}}) \leftarrow \text{C2.SimCom}(\text{CRS}_{\text{com}}^2)$ .
- (4) Output  $(u_2, c''_{\text{ser}}, cert_{\mathcal{R}})$  to  $\mathcal{Z}$  as the first message from  $\mathcal{R}$  to  $\mathcal{U}$ .
- (5) Upon receiving  $(s', \pi, \phi, a_{\mathcal{U}}, a_{\mathcal{R}}^{\text{prev}}, c_{\text{hid}}, c'_{\mathcal{R}}, t)$  from  $\mathcal{Z}$  in the name of  $\mathcal{U}$  ...
  - (a)  $stmtnt := (pk_{\mathcal{T}}^{\mathcal{T}}, pk_{\mathcal{T}}^{\text{cert}}, \phi, a_{\mathcal{U}}, a_{\mathcal{R}}^{\text{prev}}, c_{\text{hid}}, c'_{\mathcal{R}}, t, u_2)$ .
  - (b) If  $\text{P2.Vfy}(\text{CRS}_{\text{pok}}, stmtnt, \pi) = 0$  let  $\mathcal{F}_{P4TC}$  abort.
  - (c) Extract  $Wit = (X, \Lambda, pk_{\mathcal{U}}, U_1, s^{\text{prev}}, \phi^{\text{prev}}, X, \Lambda, pk_{\mathcal{U}}, B^{\text{prev}}, U_1, U_1^{\text{next}}, d_{\text{hid}}, d_{\mathcal{R}}^{\text{prev}}, d'_{\mathcal{R}}, d_{\mathcal{T}}, pk_{\mathcal{R}}^{\text{prev}}, c_{\mathcal{R}}^{\text{prev}}, c_{\mathcal{T}}, \sigma_{\mathcal{R}}^{\text{prev}}, \sigma_{\mathcal{R}}^{\text{cert}^{\text{prev}}}, \sigma_{\mathcal{T}}) \leftarrow \text{P2.ExtractW}(\text{CRS}, td_{\text{epok}}, stmtnt, \pi)$ .
  - (d) Look up  $\overline{trdb}^* := (s^{\text{prev},*}, s^*, \phi^*, x^*, \lambda^*, pid_{\mathcal{U}}^*, pid_{\mathcal{T}}^*, p^*, b^*, c_{\mathcal{T}}^{\text{in}*}, d_{\mathcal{T}}^{\text{in}*}, M_{\mathcal{T}}^{\text{in}*}, c_{\mathcal{R}}^{\text{in}*}, d_{\mathcal{R}}^{\text{in}*}, M_{\mathcal{R}}^{\text{in}*}, c_{\mathcal{T}}^{\text{out}*}, d_{\mathcal{T}}^{\text{out}*}, M_{\mathcal{T}}^{\text{out}*}, c_{\mathcal{R}}^{\text{out}*}, d_{\mathcal{R}}^{\text{out}*}, M_{\mathcal{R}}^{\text{out}*})$  with  $c_{\mathcal{R}}^{\text{out}*} = c_{\mathcal{R}}^{\text{prev}}$  and  $c_{\mathcal{T}}^{\text{out}*} = c_{\mathcal{T}}$  being used as keys; if no unique entry exists, give up simulation (event F1).
  - (e) Give up simulation if any of these conditions meet:  $\Lambda \neq g_1^{\lambda^*}$  (event F2),  $pk_{\mathcal{U}} \neq pk_{\mathcal{U}}^*$  for  $M_{\mathcal{T}}^{\text{out}*} = (\Lambda^*, pk_{\mathcal{U}}^*)$  (event F3),  $s^{\text{prev}} \neq s^*$  (event F4),  $B^{\text{prev}} \neq g_1^{b^*}$  (event D1), or  $X \neq g_1^{x^*+1}$  (event D2).
  - (f) Retrieve  $pid_{\mathcal{U}}$  from  $\overline{\mathcal{G}}_{\text{bb}}$  for  $pk_{\mathcal{U}}$  and assert  $pid_{\mathcal{U}}^* = pid_{\mathcal{U}}$  else abort (event F5).
  - (g) Call  $\mathcal{F}_{P4TC}$  with input  $(\text{pay\_toll}, s^{\text{prev}})$ ; upon being asked by  $\mathcal{F}_{P4TC}$  to provide an alternative PID, return  $pid_{\mathcal{U}}$  to  $\mathcal{F}_{P4TC}$ .
  - (h) If  $\mathcal{F}_{P4TC}$  aborts for some other reason than blacklisted, give up simulation (event F6).
    - (i) If being asked by  $\mathcal{F}_{P4TC}$  to provide  $\phi$ , return  $\phi := \text{PRF}(\lambda, x)$  with  $x := x^* + 1$  to  $\mathcal{F}_{P4TC}$ .<sup>a</sup>
    - (j) Upon being ask to provide a price for  $(a_{\mathcal{U}}, a_{\mathcal{R}}, a_{\mathcal{R}}^{\text{prev}})$  return a price  $p$  as the dummy adversary would do.
  - (k) Obtain the user's output  $(s, a_{\mathcal{R}}, p, b)$  from  $\mathcal{F}_{P4TC}$ .
    - (l) Set  $\overline{\omega}_{\mathcal{U}}^{\text{pp}} := (s, c_{\text{hid}}, d_{\text{hid}}, pk_{\mathcal{U}})$  and append  $\overline{\omega}_{\mathcal{U}}^{\text{pp}}$  to  $\overline{\Omega}_{\mathcal{U}}^{\text{pp}}$ .
  - (m) Run  $(c''_{\mathcal{R}}, d''_{\mathcal{R}}) \leftarrow \text{C1.Com}(\text{CRS}_{\text{com}}^1, (0, p, 0, 1))$ ,  $c_{\mathcal{R}} := c'_{\mathcal{R}} \cdot c''_{\mathcal{R}}$ , and  $\sigma_{\mathcal{R}} \leftarrow \text{S.Sgn}(sk_{\mathcal{R}}, (c_{\mathcal{R}}, s))$  honestly as the real protocol would do.
  - (n) Set  $s'' := s \cdot s'^{-1}$  and equivocate  $d''_{\text{ser}} \leftarrow \text{C2.Equiv}(\text{CRS}_{\text{com}}^2, s'', c''_{\text{ser}}, \bar{d}_{\text{ser}})$ .
  - (o) Set  $c_{\mathcal{T}}^{\text{in}} := c_{\mathcal{T}}, d_{\mathcal{T}}^{\text{in}} := d_{\mathcal{T}}, M_{\mathcal{T}}^{\text{in}} := (\Lambda, pk_{\mathcal{U}}), c_{\mathcal{R}}^{\text{in}} := c_{\mathcal{R}}^{\text{prev}}, d_{\mathcal{R}}^{\text{in}} := d_{\mathcal{R}}^{\text{prev}}, M_{\mathcal{R}}^{\text{in}} := (\Lambda, B^{\text{prev}}, U_1, X, s^{\text{prev}}), c_{\mathcal{T}}^{\text{out}} := c_{\mathcal{T}}, d_{\mathcal{T}}^{\text{out}} := d'_{\mathcal{T}} \cdot d''_{\mathcal{T}}, M_{\mathcal{T}}^{\text{out}} := (\Lambda, pk_{\mathcal{U}}), c_{\mathcal{R}}^{\text{out}} := c_{\mathcal{R}}, d_{\mathcal{R}}^{\text{out}} := d'_{\mathcal{R}} \cdot d''_{\mathcal{R}}$ , and  $M_{\mathcal{R}}^{\text{out}} := (\Lambda, g_1^b, U_1, g_1^{x^*+1}, s)$ .
  - (p) Append  $(s^{\text{prev}}, s, \phi, x, \lambda, pid_{\mathcal{U}}, pid_{\mathcal{R}}, p, b, c_{\mathcal{T}}^{\text{in}}, d_{\mathcal{T}}^{\text{in}}, M_{\mathcal{T}}^{\text{in}}, c_{\mathcal{R}}^{\text{in}}, d_{\mathcal{R}}^{\text{in}}, M_{\mathcal{R}}^{\text{in}}, c_{\mathcal{T}}^{\text{out}}, d_{\mathcal{T}}^{\text{out}}, M_{\mathcal{T}}^{\text{out}}, c_{\mathcal{R}}^{\text{out}}, d_{\mathcal{R}}^{\text{out}}, M_{\mathcal{R}}^{\text{out}})$  to  $\overline{TRDB}$ .
  - (q) Append  $\overline{\omega}^{\text{dsp}} = (\phi, t, u_2)$  to  $\overline{\Omega}^{\text{dsp}}$ .
  - (r) Check if  $\overline{\omega}^{\text{dsp}^{\ddagger}} = (\phi^{\ddagger}, t^{\ddagger}, u_2^{\ddagger}) \in \overline{\Omega}^{\text{dsp}}$  exists with  $\phi = \phi^{\ddagger}$  and  $u_2 \neq u_2^{\ddagger}$ ; in this case
    - (i)  $sk_{\mathcal{U}} := (t - t^{\ddagger}) \cdot (u_2 - u_2^{\ddagger})^{-1} \bmod p$
    - (ii) Record  $pid_{\mathcal{U}} \mapsto (pk_{\mathcal{U}}, sk_{\mathcal{U}})$  internally
  - (s) Output  $(s'', d''_{\text{ser}}, c_{\mathcal{R}}, d''_{\mathcal{R}}, \sigma_{\mathcal{R}}, p)$  to  $\mathcal{Z}$  as the second message from  $\mathcal{R}$  to  $\mathcal{U}$ .

<sup>a</sup>N.b.:  $\mathcal{F}_{P4TC}$  does not always ask for the next serial number. If the corrupted user re-uses an old token, then  $\mathcal{F}_{P4TC}$  internally picks the next serial number which has already been determined in some earlier interaction. Hence, the simulator only needs to provide the next serial number, if the chain of transactions is extended.

Fig. 48. The simulator for System Security (cont. from Fig. 46)

**Simulator  $\mathcal{S}_{\pi_{P4TC}}^{\text{sys-sec}}$  (cont.)**

*Debt Clearance:*

- (1) Load the recorded  $pid_{\mathcal{T}} \mapsto (pk_{\mathcal{T}}, sk_{\mathcal{T}}, cert_{\mathcal{T}}^{\mathcal{R}})$  if this does not exist, let  $\mathcal{F}_{P4TC}$  abort.
- (2) Pick  $u_2 \xleftarrow{\mathcal{R}} \mathbb{Z}_p$ .
- (3) Output  $u_2$  to  $\mathcal{Z}$  as the first message from  $\mathcal{T}$  to  $\mathcal{U}$ .
- (4) Upon receiving  $(pk_{\mathcal{U}}, \pi, \phi, a_{\mathcal{U}}, a_{\mathcal{R}}^{\text{prev}}, b^{\text{prev}}, t)$  from  $\mathcal{Z}$  in the name of  $\mathcal{U}$  ...
  - (a)  $stmt := (pk_{\mathcal{U}}, pk_{\mathcal{T}}, pk_{\mathcal{T}}^{\text{cert}}, \phi, a_{\mathcal{U}}, a_{\mathcal{R}}^{\text{prev}}, g_1^{\text{prev}}, t, u_2)$ .
  - (b) If  $P3.Vfy(\text{CRS}_{\text{pok}}, stmt, \pi) = 0$  let  $\mathcal{F}_{P4TC}$  abort.
  - (c) Extract  $Wit = (X, \Lambda, pk_{\mathcal{U}}, U_1, s^{\text{prev}}, \phi^{\text{prev}}, X, \Lambda, U_1, d_{\mathcal{R}}^{\text{prev}}, d_{\mathcal{T}}, pk_{\mathcal{R}}^{\text{prev}}, c_{\mathcal{R}}^{\text{prev}}, c_{\mathcal{T}}, \sigma_{\mathcal{R}}^{\text{prev}}, \sigma_{\mathcal{R}}^{\text{certprev}}, \sigma_{\mathcal{T}}) \leftarrow P3.ExtractW(\text{CRS}, td_{\text{epok}}, stmt, \pi)$ .
  - (d) Lookup  $\overline{trdb}^* := (s^{\text{prev},*}, s^*, \phi^*, x^*, \lambda^*, pid_{\mathcal{U}}^*, pid_{\mathcal{T}}^*, p^*, b^*, c_{\mathcal{T}}^{\text{in}*}, d_{\mathcal{T}}^{\text{in}*}, M_{\mathcal{T}}^{\text{in}*}, c_{\mathcal{R}}^{\text{in}*}, d_{\mathcal{R}}^{\text{in}*}, M_{\mathcal{R}}^{\text{in}*}, c_{\mathcal{T}}^{\text{out}*}, d_{\mathcal{T}}^{\text{out}*}, M_{\mathcal{T}}^{\text{out}*}, c_{\mathcal{R}}^{\text{out}*}, d_{\mathcal{R}}^{\text{out}*}, M_{\mathcal{R}}^{\text{out}*})$  with  $c_{\mathcal{R}}^{\text{out}*} = c_{\mathcal{R}}^{\text{prev}}$  and  $c_{\mathcal{T}}^{\text{out}*} = c_{\mathcal{T}}$  being used as keys; if no unique entry exists, give up simulation (event  $F1$ ).
  - (e) Give up simulation if any of these conditions meet:  $\Lambda \neq g_1^{\lambda^*}$  (event  $F2$ ),  $pk_{\mathcal{U}} \neq pk_{\mathcal{U}}^*$  with  $M_{\mathcal{T}}^{\text{out}*} = (\Lambda^*, pk_{\mathcal{U}}^*)$  (event  $F3$ ),  $s^{\text{prev}} \neq s^*$  (event  $F4$ ), or  $B^{\text{prev}} \neq g_1^{b^*}$  (event  $D1$ ).
  - (f) Retrieve  $pid_{\mathcal{U}}$  from  $\overline{\mathcal{G}}_{\text{bb}}$  for  $pk_{\mathcal{U}}$  and assert  $pid_{\mathcal{U}}^* = pid_{\mathcal{U}}$  else abort (event  $F5$ ).
  - (g) Call  $\mathcal{F}_{P4TC}$  with input  $(\text{clear\_debt}, s^{\text{prev}})$ ; upon being asked by  $\mathcal{F}_{P4TC}$  to provide an alternative PID, return  $pid_{\mathcal{U}}$  to  $\mathcal{F}_{P4TC}$ .
  - (h) If  $\mathcal{F}_{P4TC}$  aborts, give up simulation (event  $F6$ ).
  - (i) If being asked by  $\mathcal{F}_{P4TC}$  to provide  $\phi$ , return  $\phi := \text{PRF}(\lambda, x)$  with  $x := x^{\text{prev}} + 1$  to  $\mathcal{F}_{P4TC}$ .<sup>a</sup>
  - (j) Obtain the user's output ( $b^{\text{bill}}$ ) from  $\mathcal{F}_{P4TC}$ .
  - (k) Set  $c_{\mathcal{T}}^{\text{in}} := c_{\mathcal{T}}, d_{\mathcal{T}}^{\text{in}} := d_{\mathcal{T}}, M_{\mathcal{T}}^{\text{in}} := (\Lambda, pk_{\mathcal{U}}), c_{\mathcal{R}}^{\text{in}} := c_{\mathcal{R}}^{\text{prev}}, d_{\mathcal{R}}^{\text{in}} := d_{\mathcal{R}}^{\text{prev}}, M_{\mathcal{R}}^{\text{in}} := (\Lambda, g_1^{b^{\text{prev}}}, U_1, X, s^{\text{prev}}), c_{\mathcal{T}}^{\text{out}} := \perp, d_{\mathcal{T}}^{\text{out}} := \perp, M_{\mathcal{T}}^{\text{out}} := \perp, c_{\mathcal{R}}^{\text{out}} := \perp, d_{\mathcal{R}}^{\text{out}} := \perp, \text{ and } M_{\mathcal{R}}^{\text{out}} := \perp$ .
  - (l) Append  $(s^{\text{prev}}, s, \phi, x, \lambda, pid_{\mathcal{U}}, pid_{\mathcal{R}}, p, b, c_{\mathcal{T}}^{\text{in}}, d_{\mathcal{T}}^{\text{in}}, M_{\mathcal{T}}^{\text{in}}, c_{\mathcal{R}}^{\text{in}}, d_{\mathcal{R}}^{\text{in}}, M_{\mathcal{R}}^{\text{in}}, c_{\mathcal{T}}^{\text{out}}, d_{\mathcal{T}}^{\text{out}}, M_{\mathcal{T}}^{\text{out}}, c_{\mathcal{R}}^{\text{out}}, d_{\mathcal{R}}^{\text{out}}, M_{\mathcal{R}}^{\text{out}})$  to  $\overline{TRDB}$ .
  - (m) Append  $\overline{\omega}^{\text{dsp}} = (\phi, t, u_2)$  to  $\overline{\Omega}^{\text{dsp}}$
  - (n) Check if  $\overline{\omega}^{\text{dsp}\ddagger} = (\phi^{\ddagger}, t^{\ddagger}, u_2^{\ddagger}) \in \overline{\Omega}^{\text{dsp}}$  exists with  $\overline{\phi} = \overline{\phi}^{\ddagger}$  and  $u_2 \neq u_2^{\ddagger}$ ; in this case
    - (i)  $sk_{\mathcal{U}} := (t - t^{\ddagger}) \cdot (u_2 - u_2^{\ddagger})^{-1} \bmod p$
    - (ii) Record  $pid_{\mathcal{U}} \mapsto (pk_{\mathcal{U}}, sk_{\mathcal{U}})$  internally
  - (o) Output (OK) to  $\mathcal{Z}$  as the second message from  $\mathcal{T}$  to  $\mathcal{U}$ .

<sup>a</sup>N.b.:  $\mathcal{F}_{P4TC}$  does not always ask for the next serial number. If the corrupted user re-uses an old token, then  $\mathcal{F}_{P4TC}$  internally picks the next serial number which has already been determined in some earlier interaction. Hence, the simulator only needs to provide the next serial number, if the chain of transactions is extended.

Fig. 49. The simulator for System Security (cont. from Fig. 46)

- $\mathcal{Z}^{\text{sys-sec}}$  triggers an abort in the ideal experiment that does not exist in the real experiment. In other words  $\mathcal{S}_{\pi_{P4TC}}^{\text{sys-sec}}$  is not able to continue the simulation and gives up.
- $\mathcal{Z}^{\text{sys-sec}}$  makes the output of an honest party differ between the real and ideal experiment. As this events are directly observable by the simulator  $\mathcal{S}_{\pi_{P4TC}}^{\text{sys-sec}}$  as well,  $\mathcal{S}_{\pi_{P4TC}}^{\text{sys-sec}}$  proceeds as above and gives up simulation.

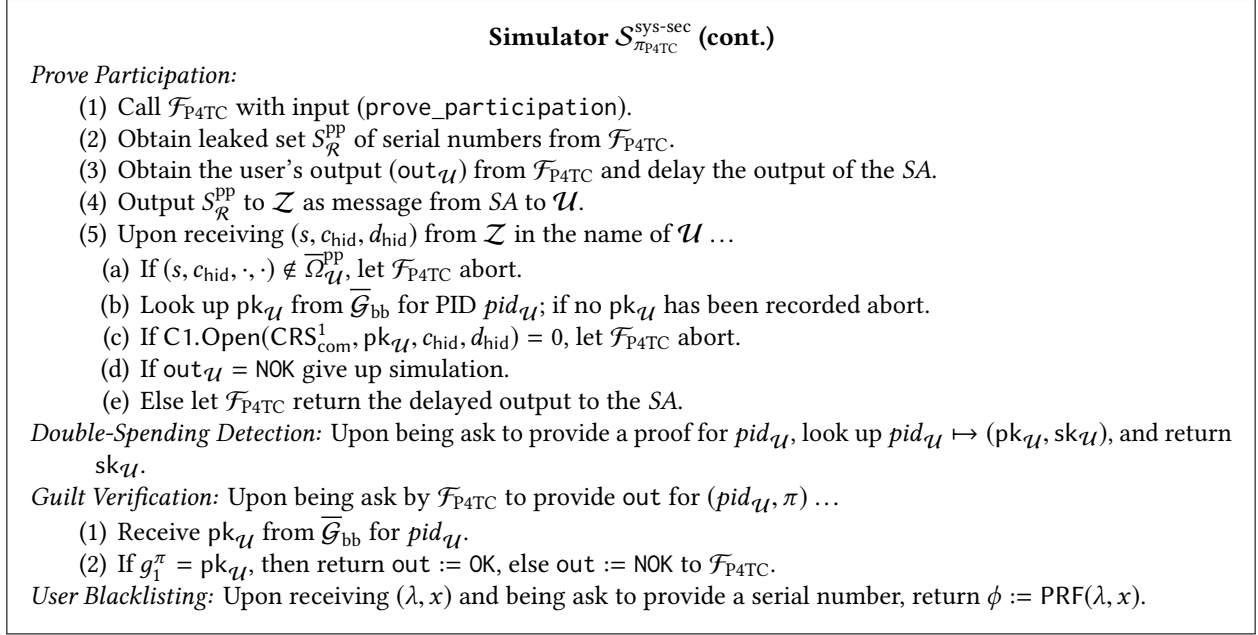


Fig. 50. The simulator for System Security (cont. from Fig. 46)

We call events of the first type *failure events* and events of the second type *discrepancy events*. Naturally, the proof of [Theorem F.22](#) proceeds in two parts according to these categories. Please note that discrepancy events do not immediately require  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$  to give up, we could also decide to let  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$  ignore this kind of event. But this way, the proof is simpler. Both parts heavily rely on the Augmented Transaction Graph.

In the first part we exploit that its graph structure is redundantly defined: Edges are defined by linking serial numbers to predecessor serial numbers and a second set of edges is defined by means of in- and out-commitments. As already shown in [Appendix F.1](#), both of these graph structures coincide if all parties are honest. We show that this still holds with overwhelming probability if the users are corrupted. More precisely, a *failure event* is a successful attempt of  $\mathcal{Z}^{\text{sys-sec}}$  to make both graph structures fall apart. In this case  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$  immediately gives up the simulation. We show that a failure event implies a successful attack on one of our cryptographic building blocks.

The second part of the proof considers executions of the UC-experiment in which  $\mathcal{Z}^{\text{sys-sec}}$  does not force  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$  to give up simulation prematurely, but in which  $\mathcal{Z}^{\text{sys-sec}}$  can distinguish due to different input/output distributions of honest parties.

The category of *failure events* comprises the following individual events:

*Definition F.25 (Failure Events).*  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$  gives up simulation because...

- (F1) The extracted  $c_{\mathcal{T}}$  and  $c_{\mathcal{R}}^{\text{prev}}$  do not point to a unique predecessor (cp. [Fig. 47, Step 5d](#) and [Fig. 48, Step 4d](#)).
- (F2) The extracted  $\Lambda$  does not correspond to the previously recorded  $\lambda^*$  in the Augmented Transaction Graph (cp. [Fig. 47, Step 5e](#) and [Fig. 48, Step 4e](#)).
- (F3) The extracted  $\text{pk}_{\mathcal{U}}$  does not match the previously recorded  $\text{pk}_{\mathcal{U}}^*$  in the Augmented Transaction Graph (cp. [Fig. 47, Step 5e](#) and [Fig. 48, Step 4e](#)).

- (F4) The extracted  $s^{\text{prev}}$  does not match the previously recorded  $s^*$  in the Augmented Transaction Graph (cp. Fig. 47, Step 5e and Fig. 48, Step 4e).
- (F5) The extracted  $\text{pk}_U$  has not been registered in  $\bar{\mathcal{G}}_{\text{bb}}$  (cp. Fig. 47, Step 5f) and Fig. 48, Step 4f).
- (F6)  $\mathcal{F}_{\text{P4TC}}$  aborts for some other reason than blacklisted (cp. Fig. 47, Step 5h and Fig. 48, Step 4h).

The category of *discrepancy events* comprises the following individual events:

*Definition F.26 (Discrepancy Events).*

- (D1) The dummy  $\mathcal{T}$  outputs a different  $b^{\text{bill}}$  to  $\mathcal{Z}^{\text{sys-sec}}$  in the scope of Debt Clearance than a real  $\mathcal{T}$  running  $\pi_{\text{P4TC}}$  (cp. Fig. 47, Step 5e and Fig. 48, Step 4e).
- (D2)  $\mathcal{Z}^{\text{sys-sec}}$  triggers a non-negligible discrepancy between the abort behavior of the dummy  $\mathcal{T}$  and the real  $\mathcal{T}$  within the scope of Debt Accumulation (cp. Fig. 47, Step 5e).
- (D3) The dummy SA outputs a different  $\text{out}_{\text{SA}}$  to  $\mathcal{Z}^{\text{sys-sec}}$  in the scope of Prove Participation than the real SA (cp. Fig. 49, Step 5d).

We start by ruling out failure events. As a preliminary step we need some kind of “loop invariant” that holds as long as no failure event has occurred.

LEMMA F.27. *For every execution of*

$$\text{EXEC}_{\mathcal{F}_{\text{P4TC}}, \bar{\mathcal{G}}_{\text{bb}}, \mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}, \mathcal{Z}^{\text{sys-sec}}(1^n)}$$

*that does not trigger a failure event the two graph structures in the Augmented Transaction Graph  $\overline{\text{TRDB}}$  maintained by  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$  are identical to the Ideal Transaction Graph  $\text{TRDB}$  maintained by  $\mathcal{F}_{\text{P4TC}}$  immediately after the first message from the user (cp. Fig. 46 Step 2, Fig. 47, Step 5 and Fig. 48, Step 4) and immediately before the last message to the user (cp. Fig. 46 Step 6h, Fig. 47, Step 5s and Fig. 48, Step 4o) within the scope of the tasks Wallet Issuing, Debt Accumulation and Debt Clearance.*

PROOF. (Informal) The statement is true for every execution of  $\text{EXEC}_{\mathcal{F}_{\text{P4TC}}, \bar{\mathcal{G}}_{\text{bb}}, \mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}, \mathcal{Z}^{\text{sys-sec}}(1^n)$  that never triggered Wallet Issuing, Debt Accumulation or Debt Clearance. Assume the statement holds at the end of  $n$  such interactions when the  $(n + 1)^{\text{th}}$  interaction starts. By assumption the Ideal Transaction Graph of  $\mathcal{F}_{\text{P4TC}}$  and the two graph structures of the Augmented Transaction graph of  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$  are in sync before the first message. After the the first message of the  $(n + 1)^{\text{th}}$  interaction has been received, the  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$  conducts a collection of consistency checks, that by assumption do not fail (otherwise a failure event is triggered). The ideal function is invoked and terminates without abort (again by assumption). During its execution  $\mathcal{F}_{\text{P4TC}}$  updates its internal Ideal Transaction Graph  $\text{TRDB}$ . After  $\mathcal{F}_{\text{P4TC}}$  returned its output to the  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$  and before  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$  sends its last message  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$  updates the Augmented Transaction Graph  $\overline{\text{TRDB}}$ . By definition of  $\mathcal{F}_{\text{P4TC}}$  and  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$  both updates are consistent.  $\square$

An immediate consequence of Lemma F.27 is the following

LEMMA F.28 (NO UNEXPECTED ABORT). *The failure event (F6) does not occur.*

PROOF. Assume for contradiction that (F6) has occurred. This implies that no other failure event has occurred before, in particular none of the events (F3) or (F4). Due to Lemma F.27 both transaction graphs  $\text{TRDB}$  and  $\overline{\text{TRDB}}$  are in sync. But  $\mathcal{F}_{\text{P4TC}}$  only aborts unexpectedly and raises (F6) if the predecessor serial number  $s^{\text{prev}}$  and the PID  $\text{pid}_U$  are not recorded in  $\text{TRDB}$ . In this case (F3) or (F4) must have occurred before.  $\square$

LEMMA F.29. *If P2 and P3 are perfectly sound, C1 is computationally binding and S is EUF-CMA secure, then in a run of*

$$\text{EXEC}_{\mathcal{F}_{\text{P4TC}}, \bar{\mathcal{G}}_{\text{bb}}, \mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}, \mathcal{Z}^{\text{sys-sec}}(1^n)}$$

the event (F1) (cp. Definition F.25) only occurs with negligible probability for any PPT environment  $\mathcal{Z}^{\text{sys-sec}}$  that is restricted to static corruption of users.

PROOF. Assume a  $\mathcal{Z}^{\text{sys-sec}}$  exists that triggers (F1) with non-negligible probability. Without loss of generality we only consider that (F1) occurs in Debt Accumulation; the case of Debt Clearance is analogous. Let  $(s', \pi, \phi, \mathbf{a}_U, \mathbf{a}_R^{\text{prev}}, c_{\text{hid}}, c'_R, t)$  denote the message that was sent by  $\mathcal{Z}^{\text{sys-sec}}$  and raised (F1). Let  $\text{Wit} \leftarrow \text{P2.ExtractW}(\text{CRS}, \text{td}_{\text{epok}}, \text{stmnt}, \pi)$  denote the extracted witness and parse it as  $(\cdot, \Lambda, \dots, s^{\text{prev}}, \dots, d_R^{\text{prev}}, d_{\mathcal{T}}, \text{pk}_R^{\text{prev}}, c_R^{\text{prev}}, c_{\mathcal{T}}, \sigma_R^{\text{prev}}, \sigma_{\mathcal{T}}^{\text{certprev}}, \sigma_{\mathcal{T}}) := \text{Wit}$ .

We observe the following facts:

- $\text{P2.Vfy}(\text{CRS}_{\text{pok}}, \text{stmnt}, \pi) = 1$  holds, otherwise  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$  would have let  $\mathcal{F}_{\text{P4TC}}$  aborted earlier.
- As P2 is perfectly sound, the equations

$$\begin{aligned} \text{C1.Open}(\text{CRS}_{\text{com}}^1, (\Lambda, \dots), c_{\mathcal{T}}, d_{\mathcal{T}}) &= 1 \\ \text{C1.Open}(\text{CRS}_{\text{com}}^1, (\Lambda, \dots), c_R^{\text{prev}}, d_R^{\text{prev}}) &= 1 \\ \text{S.Vfy}(\text{pk}_{\mathcal{T}}^{\mathcal{T}}, \sigma_{\mathcal{T}}, (c_{\mathcal{T}}, \mathbf{a}_U)) &= 1 \\ \text{S.Vfy}(\text{pk}_R^{\text{prev}}, \sigma_R^{\text{prev}}, (c_R^{\text{prev}}, s^{\text{prev}})) &= 1 \\ \text{S.Vfy}(\text{pk}_{\mathcal{T}}^{\text{cert}}, \sigma_{\mathcal{T}}^{\text{certprev}}, (\text{pk}_R^{\text{prev}}, \dots)) &= 1 \end{aligned}$$

hold.

First we deal with the cases that at least one of  $c_{\mathcal{T}}$  or  $c_R^{\text{prev}}$  does not exist as an out-commitment in  $\overline{\text{TRDB}}$ .

- (1) *The commitment  $c_{\mathcal{T}}$  does not exist as an out-commitment:* In other words, the commitment  $c_{\mathcal{T}}$  has never been issued by the TSP. We construct an efficient adversary  $\mathcal{B}$  against the EUF-CMA security of S. Internally,  $\mathcal{B}$  runs  $\mathcal{Z}^{\text{sys-sec}}$  in its head, and simulates  $\text{EXEC}_{\mathcal{F}_{\text{P4TC}}, \overline{\mathcal{G}}_{\text{bb}}, \mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}, \mathcal{Z}^{\text{sys-sec}}}(1^n)$  for  $\mathcal{Z}^{\text{sys-sec}}$ , i.e.,  $\mathcal{B}$  internally plays the role of  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$  and  $\mathcal{F}_{\text{P4TC}}$ . Externally,  $\mathcal{B}$  plays the EUF-CMA security experiment with a challenger  $\mathcal{C}$  and a signing oracle  $\text{O}_{\text{pk}, \text{sk}}^{\text{S}}$ . When  $\mathcal{B}$  must internally provide  $\text{pk}_{\mathcal{T}} = (\text{pk}_{\mathcal{T}}^{\text{cert}}, \text{pk}_{\mathcal{T}}^{\mathcal{R}}, \text{pk}_{\mathcal{T}}^{\mathcal{T}})$  playing the role of  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$  in the scope of the TSP Registration,  $\mathcal{B}$  embeds the external challenge as  $\text{pk}_{\mathcal{T}}^{\mathcal{T}} := \text{pk}_{\mathcal{C}}$ . Whenever  $\mathcal{B}$  playing the role of  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$  needs to issue signatures with respect to  $\text{pk}_{\mathcal{T}}$ , it does so using its external EUF-CMA oracle  $\text{O}_{\text{pk}, \text{sk}}^{\text{S}}$ . When the event (F1) occurs,  $\mathcal{B}$  extracts  $c_{\mathcal{T}}$  and  $\sigma_{\mathcal{T}}$  from the proof and outputs the forgery.
- (2) *The commitment  $c_R^{\text{prev}}$  does not exist as an out-commitment:* We need to distinguish two sub-cases. On an abstract level these sub-cases correspond to the following scenarios: Either the previous RSU exists. Then the signature on  $c_R^{\text{prev}}$  is a forgery. Or alternatively, the allegedly previous RSU does not exist but has been imagined by the user. Then  $c_R^{\text{prev}}$  may have a honest, valid signature (because the user feigned the RSU), but the certificate for the fake RSU is a forgery.
  - (a) *A record  $\text{pid}_R^{\text{prev}} \mapsto (\text{pk}_R^{\text{prev}}, \text{sk}_R^{\text{prev}})$  has been recorded:* Again, we construct an efficient adversary  $\mathcal{B}$  against the EUF-CMA security of S along the same lines as above. But in this case,  $\mathcal{B}$  needs to guess for which  $\text{pid}_R^{\text{prev}}$  the event (F1) eventually occurs. When the RSU with  $\text{pid}_R^{\text{prev}}$  registers itself, and  $\mathcal{B}$  playing  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$  needs to provide  $\text{pk}_R^{\text{prev}}$  it embeds the challenge as  $\text{pk}_R^{\text{prev}} := \text{pk}_{\mathcal{C}}$ . Whenever  $\mathcal{B}$  playing the role of  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$  needs to issue a signature with respect to  $\text{pk}_R^{\text{prev}}$ , it does so using its external EUF-CMA oracle  $\text{O}_{\text{pk}, \text{sk}}^{\text{S}}$ . When the event (F1) occurs,  $\mathcal{B}$  extracts  $(c_R^{\text{prev}}, s^{\text{prev}})$  and  $\sigma_R^{\text{prev}}$  from the proof and outputs the forgery. N.b., has never been signed wrpt.  $\text{pk}_R^{\text{prev}} = \text{pk}_{\mathcal{C}}$  by assumption.
  - (b) *No  $\text{pid}_R^{\text{prev}} \mapsto (\text{pk}_R^{\text{prev}}, \text{sk}_R^{\text{prev}}, \text{cert}_R)$  has been recorded:* We construct an efficient adversary  $\mathcal{B}$  against the EUF-CMA security of S. Internally,  $\mathcal{B}$  runs  $\mathcal{Z}^{\text{sys-sec}}$  in its head, and simulates  $\text{EXEC}_{\mathcal{F}_{\text{P4TC}}, \overline{\mathcal{G}}_{\text{bb}}, \mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}, \mathcal{Z}^{\text{sys-sec}}}(1^n)$

for  $\mathcal{Z}^{\text{sys-sec}}$ , i.e.,  $\mathcal{B}$  internally plays the role of  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$  and  $\mathcal{F}_{\text{P4TC}}$ . Externally,  $\mathcal{B}$  plays the EUF-CMA security experiment with a challenger  $\mathcal{C}$  and a signing oracle  $\text{O}_{\text{pk,sk}}^{\text{S}}$ . When the adversary  $\mathcal{B}$  has to internally provide  $\text{pk}_{\mathcal{T}} = (\text{pk}_{\mathcal{T}}^{\text{cert}}, \text{pk}_{\mathcal{T}}^{\mathcal{R}}, \text{pk}_{\mathcal{T}}^{\mathcal{T}})$  playing the role of  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$  in the scope of the TSP Registration,  $\mathcal{B}$  embeds the external challenge as  $\text{pk}_{\mathcal{T}}^{\text{cert}} := \text{pk}_{\mathcal{C}}$ . Whenever  $\mathcal{B}$  playing the role of  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$  in the scope of RSU Certification needs to issue signatures with respect to  $\text{pk}_{\mathcal{T}}^{\text{cert}}$ , it does so using its external EUF-CMA oracle  $\text{O}_{\text{pk,sk}}^{\text{S}}$ . When the event (F1) occurs,  $\mathcal{B}$  extracts  $\text{cert}_{\mathcal{R}}^{\text{prev}}$  and  $\sigma_{\mathcal{R}}^{\text{certprev}}$  from the proof and outputs the forgery. N.b.:  $\text{cert}_{\mathcal{R}}^{\text{prev}}$  has never been signed by the TSP with respect to  $\text{pk}_{\mathcal{T}}^{\text{cert}} = \text{pk}_{\mathcal{C}}$  as otherwise a mapping  $\text{pid}_{\mathcal{R}}^{\text{prev}} \mapsto (\text{pk}_{\mathcal{R}}^{\text{prev}}, \text{sk}_{\mathcal{R}}^{\text{prev}}, \text{cert}_{\mathcal{R}})$  would have been recorded.

At this point we know that both commitments  $c_{\mathcal{T}}$  and  $c_{\mathcal{R}}^{\text{prev}}$  exist as an out-commitment somewhere in  $\overline{\text{TRDB}}$ . By assumption there is no single  $\overline{\text{trdb}}^*$  with both commitments as out-commitments. Let  $\overline{\text{trdb}}^*$  denote the transaction entry with  $c_{\mathcal{T}}^{\text{out}*} = c_{\mathcal{T}}$  and Let  $\overline{\text{trdb}}^{**}$  denote the transaction entry with  $c_{\mathcal{R}}^{\text{out}*} = c_{\mathcal{R}}^{\text{prev}}$ . We have

$$\begin{aligned} \overline{\text{trdb}}^* &= (\dots, \lambda^*, \dots, \underbrace{c_{\mathcal{T}}^{\text{out}*}, d_{\mathcal{T}}^{\text{out}*}, M_{\mathcal{T}}^{\text{out}*}}_{=c_{\mathcal{T}}}, \dots) \\ \overline{\text{trdb}}^{**} &= (\dots, \lambda^{**}, \dots, \underbrace{c_{\mathcal{R}}^{\text{out}*}, d_{\mathcal{R}}^{\text{out}*}, M_{\mathcal{R}}^{\text{out}*}}_{=c_{\mathcal{R}}^{\text{prev}}}, \dots) \end{aligned}$$

and  $\overline{\text{trdb}}^* \neq \overline{\text{trdb}}^{**}$ . Because no failure event occurred before,  $\text{TRDB}$  and  $\overline{\text{TRDB}}$  are in sync by [Lemma F.27](#),  $\text{TRDB}$  is correct by definition and hence all  $c_{\mathcal{T}i}^{\text{in}*}, c_{\mathcal{T}i}^{\text{out}*}$  are constant and equal to  $c_{\mathcal{T}}$  within the tree to that the node  $\overline{\text{trdb}}^*$  belongs. This implies the nodes  $\overline{\text{trdb}}^*$  and  $\overline{\text{trdb}}^{**}$  and therefore the commitments  $c_{\mathcal{T}}$  and  $c_{\mathcal{R}}^{\text{prev}}$  belong to different trees. From [Lemma F.4](#) we conclude that  $\lambda^* \neq \lambda^{**}$  holds with overwhelming probability. On the one hand it follows that

$$\begin{aligned} M_{\mathcal{T}}^{\text{out}*} &= (\Lambda^*, \text{pk}_{\mathcal{U}}^*) \quad \text{and} \\ M_{\mathcal{R}}^{\text{out}*} &= (\Lambda^{**}, \dots) \end{aligned}$$

are implicit openings with different wallet IDs. On the other hand the perfect correctness and extractability of the proof system yield

$$\begin{aligned} M_{\mathcal{T}} &= (\Lambda, \text{pk}_{\mathcal{U}}) \quad \text{and} \\ M_{\mathcal{R}}^{\text{prev}} &= (\Lambda, \dots) \end{aligned}$$

to be valid implicit openings to the same wallet ID. In summary, at least one of the commitments  $c_{\mathcal{T}}^{\text{out}*} = c_{\mathcal{T}}$  or  $c_{\mathcal{R}}^{\text{out}*} = c_{\mathcal{R}}^{\text{prev}}$  has two different implicit openings. We construct an efficient adversary  $\mathcal{B}$  against the binding property of C1. Internally,  $\mathcal{B}$  runs  $\mathcal{Z}^{\text{sys-sec}}$  in its head, and simulates  $\text{EXEC}_{\mathcal{F}_{\text{P4TC}}, \overline{\mathcal{G}}_{\text{bb}}, \mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}, \mathcal{Z}^{\text{sys-sec}}}(1^n)$  for  $\mathcal{Z}^{\text{sys-sec}}$ , i.e.,  $\mathcal{B}$  internally plays the role of  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$  and  $\mathcal{F}_{\text{P4TC}}$ . When the event (F1) occurs,  $\mathcal{B}$  searches the forest  $\overline{\text{TRDB}}$  for the duplicate of the commitment and outputs both openings.  $\square$

The above lemma requires a remark. If the lemma is used to construct an efficient adversary  $\mathcal{B}$  out of an environment  $\mathcal{Z}^{\text{sys-sec}}$  that triggers the failure event (F1), the adversary  $\mathcal{B}$  is fixated on the security game it plays externally. As  $\mathcal{B}$  uses  $\mathcal{Z}^{\text{sys-sec}}$  in a black-box fashion it does not know how  $\mathcal{Z}^{\text{sys-sec}}$  eventually causes the event (F1). But there are only finite many options and if the probability that  $\mathcal{Z}^{\text{sys-sec}}$  distinguishes correctly is non-negligible, then at least one of the options is non-negligible, too. In other words,  $\mathcal{B}$  needs to push its luck

that  $\mathcal{Z}^{\text{sys-sec}}$  breaks the same underlying building block which  $\mathcal{B}$  attacks externally. If  $\mathcal{B}$  guesses wrong,  $\mathcal{B}$  loses its external game, but the advantage is still non-negligible.

LEMMA F.30. *If P2 and P3 are perfectly sound, C1 is computationally binding and S is EUF-CMA secure, then in a run of*

$$\text{EXEC}_{\mathcal{F}_{\text{P4TC}}, \overline{\mathcal{G}}_{\text{bb}}, \mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}, \mathcal{Z}^{\text{sys-sec}}(1^n)}$$

*the events (F2) or (F3) (cp. Definition F.25) only occur with negligible probability for any PPT environment  $\mathcal{Z}^{\text{sys-sec}}$  that is restricted to static corruption of users.*

PROOF. Assume a  $\mathcal{Z}^{\text{sys-sec}}$  exists that triggers (F2) or (F3) with non-negligible probability but not (F1) as otherwise  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$  would have aborted earlier. Without loss of generality we only consider Debt Accumulation; the case of Debt Clearance is analogous. Let  $(s', \pi, \phi, \mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}}, c_{\text{hid}}, c'_{\mathcal{R}}, t)$  denote the message that was sent by  $\mathcal{Z}^{\text{sys-sec}}$  and raised (F2) or (F3). Let  $\text{Wit} = (\dots, \Lambda, \text{pk}_{\mathcal{U}}, \dots, d_{\mathcal{T}}, \dots, c_{\mathcal{T}}, \dots) \leftarrow \text{P2.ExtractW}(\text{CRS}, \text{td}_{\text{epok}}, \text{stmnt}, \pi)$  denote the extracted witness. We observe the following facts:

- P2.Vfy(CRS<sub>pok</sub>, stmnt,  $\pi$ ) = 1 holds, otherwise  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$  would have let  $\mathcal{F}_{\text{P4TC}}$  aborted earlier.
- As P2 is perfectly sound, the equation C1.Open(CRS<sub>com</sub><sup>1</sup>,  $(\Lambda, \text{pk}_{\mathcal{U}}), c_{\mathcal{T}}, d_{\mathcal{T}}$ ) = 1, holds.

As (F1) has not occurred we know by Lemma F.29 that there is a unique  $\overrightarrow{\text{trdb}}^* = (\dots, \lambda^*, \dots, c_{\mathcal{T}}^{\text{out}*}, d_{\mathcal{T}}^{\text{out}*}, M_{\mathcal{T}}^{\text{out}*}, \dots)$  with  $c_{\mathcal{T}}^{\text{out}*} = c_{\mathcal{T}}$  and  $M_{\mathcal{T}}^{\text{out}*} = (\Lambda^*, \text{pk}_{\mathcal{U}}^*)$ . Hence

$$\text{C1.Open}(\text{CRS}_{\text{com}}^1, (\Lambda^*, \text{pk}_{\mathcal{U}}^*), c_{\mathcal{T}}, d_{\mathcal{T}}^{\text{out}*}) = 1$$

holds, too. By assumption at least one of  $\Lambda^* \neq \Lambda$  or  $\text{pk}_{\mathcal{U}}^* \neq \text{pk}_{\mathcal{U}}$  holds. This immediately yields an efficient adversary  $\mathcal{B}$  against the binding property of C1:  $\mathcal{B}$  internally executes the experiment  $\text{EXEC}_{\mathcal{F}_{\text{P4TC}}, \overline{\mathcal{G}}_{\text{bb}}, \mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}, \mathcal{Z}^{\text{sys-sec}}(1^n)}$  and plays  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$  and  $\mathcal{F}_{\text{P4TC}}$  for  $\mathcal{Z}^{\text{sys-sec}}$  until  $\mathcal{B}$  observes event (F2) or (F3). Then  $\mathcal{B}$  outputs  $c_{\mathcal{T}}$  together with two different openings  $M := (\Lambda, \text{pk}_{\mathcal{U}}), d_{\mathcal{T}} M_{\mathcal{T}}^{\text{out}*} = (\Lambda^*, \text{pk}_{\mathcal{U}}^*)$ , and  $d_{\mathcal{T}}^{\text{out}*}$ .  $\square$

LEMMA F.31. *If P2 and P3 are perfectly sound, C1 is computationally binding and S is EUF-CMA secure, then in a run of*

$$\text{EXEC}_{\mathcal{F}_{\text{P4TC}}, \overline{\mathcal{G}}_{\text{bb}}, \mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}, \mathcal{Z}^{\text{sys-sec}}(1^n)}$$

*the event (F4) (cp. Definition F.25) only occurs with negligible probability for any PPT environment  $\mathcal{Z}^{\text{sys-sec}}$  that is restricted to static corruption of users.*

PROOF. Assume a  $\mathcal{Z}^{\text{sys-sec}}$  exists that triggers (F4) with non-negligible probability but not (F1) as otherwise  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$  would have aborted earlier. Without loss of generality we only consider Debt Accumulation; the case of Debt Clearance is analogous. Let  $(s', \pi, \phi, \mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}}, c_{\text{hid}}, c'_{\mathcal{R}}, t)$  denote the message that was sent by  $\mathcal{Z}^{\text{sys-sec}}$  and raised (F4). Let  $\text{Wit} = (\dots, s^{\text{prev}}, \dots, c_{\mathcal{R}}^{\text{prev}}, \dots, \sigma_{\mathcal{R}}^{\text{prev}}, \dots) \leftarrow \text{P2.ExtractW}(\text{CRS}, \text{td}_{\text{epok}}, \text{stmnt}, \pi)$  denote the extracted witness. We observe the following facts:

- P2.Vfy(CRS<sub>pok</sub>, stmnt,  $\pi$ ) = 1 holds, otherwise  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$  would have let  $\mathcal{F}_{\text{P4TC}}$  aborted earlier.
- As P2 is perfectly sound, the equation S.Vfy(pk<sub>R</sub><sup>prev</sup>,  $\sigma_{\mathcal{R}}^{\text{prev}}, (c_{\mathcal{R}}^{\text{prev}}, s^{\text{prev}})) = 1$  holds.

As (F1) has not occurred we know by Lemma F.29 that there exists a unique  $\overrightarrow{\text{trdb}}^* = (\dots, s^*, \dots, \text{pid}_{\mathcal{R}}^{\text{prev}}, \dots, c_{\mathcal{R}}^{\text{out}*})$  with  $c_{\mathcal{R}}^{\text{out}*} = c_{\mathcal{R}}^{\text{prev}}$ , but  $s^* \neq s^{\text{prev}}$ . Due to the uniqueness of  $c_{\mathcal{R}}^{\text{prev}}$  the message  $(c_{\mathcal{R}}^{\text{prev}}, s^{\text{prev}})$  has never been signed and thus  $\sigma_{\mathcal{R}}^{\text{prev}}$  is a forgery. We show how to construct an efficient adversary  $\mathcal{B}$  against the EUF-CMA security of S. Externally,  $\mathcal{B}$  plays the EUF-CMA security experiment with a challenger  $\mathcal{C}$  and a signing oracle  $\text{O}_{\text{pk}, \text{sk}}^{\text{S}}$ . Internally,  $\mathcal{B}$  executes the experiment  $\text{EXEC}_{\mathcal{F}_{\text{P4TC}}, \overline{\mathcal{G}}_{\text{bb}}, \mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}, \mathcal{Z}^{\text{sys-sec}}(1^n)}$  in its head and plays  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$  and  $\mathcal{F}_{\text{P4TC}}$  for  $\mathcal{Z}^{\text{sys-sec}}$ .  $\mathcal{B}$  needs to guess for which  $\text{pid}_{\mathcal{R}}^{\text{prev}}$  the event (F1) eventually occurs. When the RSU with



$pid_{\mathcal{R}}^{\text{prev}}$  registers itself, and  $\mathcal{B}$  playing  $\mathcal{S}_{\pi_{\text{P4TC}}^{\text{sys-sec}}}$  needs to provide  $pk_{\mathcal{R}}^{\text{prev}}$  it embeds the challenge as  $pk_{\mathcal{R}}^{\text{prev}} := pk_{\mathcal{C}}$ . Whenever  $\mathcal{B}$  playing the role of  $\mathcal{S}_{\pi_{\text{P4TC}}^{\text{sys-sec}}}$  needs to issue a signature with respect to  $pk_{\mathcal{R}}^{\text{prev}}$ , it does so using its external EUF-CMA oracle  $\mathcal{O}_{pk,sk}^S$ . When the event (F4) occurs,  $\mathcal{B}$  extracts  $(c_{\mathcal{R}}^{\text{prev}}, s^{\text{prev}})$  and  $\sigma_{\mathcal{R}}^{\text{prev}}$  from the proof and outputs the forgery.  $\square$

LEMMA F.32. *If P2 and P3 are perfectly sound, C1 is computationally binding and S is EUF-CMA secure, then in a run of*

$$\text{EXEC}_{\mathcal{F}_{\text{P4TC}}, \overline{\mathcal{G}}_{\text{bb}}, \mathcal{S}_{\pi_{\text{P4TC}}^{\text{sys-sec}}}, \mathcal{Z}^{\text{sys-sec}}(1^n)}$$

the event (F5) (cp. Definition F.25) does not occur for any PPT environment  $\mathcal{Z}^{\text{sys-sec}}$  that is restricted to static corruption of users.

PROOF. Without loss of generality we only consider Debt Accumulation; the case of Debt Clearance is analogous. Let  $(s', \pi, \phi, a_{\mathcal{U}}, a_{\mathcal{R}}^{\text{prev}}, c_{\text{hid}}, c'_{\mathcal{R}}, t)$  denote the message that was sent by  $\mathcal{Z}^{\text{sys-sec}}$  and raised (F5). Let  $\text{Wit} \leftarrow \text{P2.ExtractW}(\text{CRS}, \text{td}_{\text{epok}}, \text{stmt}, \pi)$  denote the extracted witness and parse  $(\dots, pk_{\mathcal{U}}, \dots, c_{\mathcal{R}}^{\text{prev}}, \dots) := \text{Wit}$ . As (F1) has not occurred we know by Lemma F.29 that there is a unique  $\overline{\text{trdb}}^* = (\dots, pid_{\mathcal{U}}^*, \dots, c_{\mathcal{R}}^{\text{out}*}, \dots, M_{\mathcal{T}}^{\text{out}*}, \dots)$  with  $c_{\mathcal{R}}^{\text{out}*} = c_{\mathcal{R}}^{\text{prev}}$ . We distinguish two cases:

- (1)  $\overline{\text{trdb}}^*$  is the result of a run of Debt Accumulation or Debt Clearance: This immediately leads to a contradiction. The event (F3) has not occurred, as otherwise  $\mathcal{S}_{\pi_{\text{P4TC}}^{\text{sys-sec}}}$  would have aborted earlier. Hence  $pk_{\mathcal{U}} = pk_{\mathcal{U}}^*$  holds and this implies that  $pk_{\mathcal{U}}^*$  does not exist in the bulletin board neither. However,  $M_{\mathcal{T}}^{\text{out}*} = (\Lambda^*, pk_{\mathcal{U}}^*)$  is the result of a successful previous run of Debt Accumulation or Debt Clearance and hence  $pk_{\mathcal{U}}^*$  must exist in the bulletin board. Contradiction!
- (2)  $\overline{\text{trdb}}^*$  is the result of a run of Wallet Issuing: The event (F3) has not occurred, as otherwise  $\mathcal{S}_{\pi_{\text{P4TC}}^{\text{sys-sec}}}$  would have aborted earlier. Hence  $pk_{\mathcal{U}} = pk_{\mathcal{U}}^*$  holds and this implies that  $pk_{\mathcal{U}}^*$  does not exist in the bulletin board neither. However,  $M_{\mathcal{T}}^{\text{out}*} = (\Lambda^*, pk_{\mathcal{U}}^*)$  is the result of Wallet Issuing. But Wallet Issuing asserts that  $pk_{\mathcal{U}}^*$  exists or aborts (cp. Fig. 46, Step 2a). Contradiction!

$\square$

LEMMA F.33. *If P2 and P3 are perfectly sound, C1 is computationally binding and S is EUF-CMA secure, then in a run of*

$$\text{EXEC}_{\mathcal{F}_{\text{P4TC}}, \overline{\mathcal{G}}_{\text{bb}}, \mathcal{S}_{\pi_{\text{P4TC}}^{\text{sys-sec}}}, \mathcal{Z}^{\text{sys-sec}}(1^n)}$$

the event (D1) (cp. Definition F.26) only occurs with negligible probability for any PPT environment  $\mathcal{Z}^{\text{sys-sec}}$  that is restricted to static corruption of users.

PROOF. Assume a  $\mathcal{Z}^{\text{sys-sec}}$  exists that triggers (D1) with non-negligible probability but not (F1) as otherwise  $\mathcal{S}_{\pi_{\text{P4TC}}^{\text{sys-sec}}}$  would have aborted earlier. Unveiling a different, final balance  $b^{\text{bill}}$  in the scope of Debt Clearance implies that the balance  $b^*$  of the (real) wallet diverges from the expected balance of the (ideal) wallet at some point along the path of transactions from the root node with  $b^* = 0$  (Wallet Issue) to the leaf node with  $b^* = b^{\text{bill}}$  (Debt Clearance). Without loss of generality we only consider a transition between inner nodes (i.e., Debt Accumulation); the transition from the next-to-last to the leaf node (i.e., Debt Clearance) is analogous.

Let  $(s', \pi, \phi, a_{\mathcal{U}}, a_{\mathcal{R}}^{\text{prev}}, c_{\text{hid}}, c'_{\mathcal{R}}, t)$  denote the message that was sent by  $\mathcal{Z}^{\text{sys-sec}}$  and raised (D1). Let  $\text{Wit} = (X, \dots, \Lambda, \dots, B^{\text{prev}}, \dots, U_1, \dots, d_{\mathcal{R}}^{\text{prev}}, \dots, c_{\mathcal{R}}^{\text{prev}}, \dots) \leftarrow \text{P2.ExtractW}(\text{CRS}, \text{td}_{\text{epok}}, \text{stmt}, \pi)$  denote the extracted witness. We observe the following facts:

- $\text{P2.Vfy}(\text{CRS}_{\text{pok}}, \text{stmt}, \pi) = 1$  holds, otherwise  $\mathcal{S}_{\pi_{\text{P4TC}}^{\text{sys-sec}}}$  would have let  $\mathcal{F}_{\text{P4TC}}$  aborted earlier.
- As P2 is perfectly sound, the equation  $\text{C1.Open}(\text{CRS}_{\text{com}}^1, M_{\mathcal{R}}^{\text{prev}}, c_{\mathcal{R}}^{\text{prev}}, d_{\mathcal{R}}^{\text{prev}}) = 1$  with  $M_{\mathcal{R}}^{\text{prev}} := (\Lambda, B^{\text{prev}}, U_1, X)$ , holds.

As (F1) has not occurred we know by [Lemma F.29](#) that there is a unique  $\overline{trdb}^* = (\dots, \lambda^*, \dots, c_{\mathcal{R}}^{\text{out}*}, d_{\mathcal{R}}^{\text{out}*}, M_{\mathcal{R}}^{\text{out}*})$  with  $c_{\mathcal{R}}^{\text{out}*} = c_{\mathcal{R}}^{\text{prev}}$  and hence  $\text{C1.Open}(\text{CRS}_{\text{com}}^1, M_{\mathcal{R}}^{\text{out}*}, c_{\mathcal{R}}^{\text{prev}}, d_{\mathcal{R}}^{\text{out}*}) = 1$  holds, too. Let be  $M_{\mathcal{R}}^{\text{out}*} = (\Lambda^*, g_1^b, U_1, g_1^{x^*+1})$ . By assumption  $g_1^b \neq B^{\text{prev}}$  holds (raise condition for event (D1)) and so does  $M_{\mathcal{R}}^{\text{prev}} \neq M_{\mathcal{R}}^{\text{out}*}$ . This immediately yields an efficient adversary  $\mathcal{B}$  against the binding property of C1:  $\mathcal{B}$  internally executes the experiment  $\text{EXEC}_{\mathcal{F}_{\text{P4TC}}, \overline{\mathcal{G}}_{\text{bb}}, \mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}, \mathcal{Z}^{\text{sys-sec}}}(1^n)$  in its head and plays  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$  and  $\mathcal{F}_{\text{P4TC}}$  for  $\mathcal{Z}^{\text{sys-sec}}$  until  $\mathcal{B}$  observes event (D1). Then  $\mathcal{B}$  outputs  $c_{\mathcal{R}}^{\text{prev}}$  together with two different openings  $M_{\mathcal{R}}^{\text{prev}}, d_{\mathcal{R}}^{\text{prev}}$  and  $M_{\mathcal{R}}^{\text{out}*}, d_{\mathcal{R}}^{\text{out}*}$ .  $\square$

Please note that for an honest RSU  $\mathcal{R}$  and a malicious user  $\mathcal{U}$  the blacklist  $bl_{\mathcal{R}}$  (input of  $\mathcal{R}$ ) and fraud detection ID  $\phi$  (part of the message from  $\mathcal{U}$  to  $\mathcal{R}$ ) are both given by  $\mathcal{Z}^{\text{sys-sec}}$ . Moreover,  $\mathcal{Z}^{\text{sys-sec}}$  provides a proof that  $\phi = \text{PRF}(\lambda, x)$  has been evaluated correctly. Hence, the only way how  $\mathcal{Z}^{\text{sys-sec}}$  can achieve a discrepancy with respect to the abort behavior between the real and ideal experiment is to manipulate  $x$ .

**LEMMA F.34.** *If P2 and P3 are perfectly sound, C1 is computationally binding and S is EUF-CMA secure, then in a run of*

$$\text{EXEC}_{\mathcal{F}_{\text{P4TC}}, \overline{\mathcal{G}}_{\text{bb}}, \mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}, \mathcal{Z}^{\text{sys-sec}}}(1^n)$$

*the event (D2) (cp. [Definition F.26](#)) only occurs with negligible probability for any PPT environment  $\mathcal{Z}^{\text{sys-sec}}$  that is restricted to static corruption of users.*

**PROOF.** Identical to the proof of [Lemma F.33](#) except for a different raise condition.  $\square$

**LEMMA F.35.** *If P2 is perfectly sound and C1 is computationally binding, then in a run of*

$$\text{EXEC}_{\mathcal{F}_{\text{P4TC}}, \overline{\mathcal{G}}_{\text{bb}}, \mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}, \mathcal{Z}^{\text{sys-sec}}}(1^n)$$

*the event (D3) (cp. [Definition F.26](#)) only occurs with negligible probability for any PPT environment  $\mathcal{Z}^{\text{sys-sec}}$  that is restricted to static corruption of users.*

**PROOF.** Please note that in the event of (D3), we already know that  $\text{C1.Open}(\text{CRS}_{\text{com}}^1, \text{pk}_{\mathcal{U}}, c_{\text{hid}}, d_{\text{hid}}) = 1$  holds as otherwise  $\mathcal{F}_{\text{P4TC}}$  would have aborted earlier. Let  $\overline{\omega}_{\mathcal{U}}^{\text{pp}*} = (s^*, c_{\text{hid}}^*, d_{\text{hid}}^*, \text{pk}_{\mathcal{U}}^*) \in \overline{\Omega}_{\mathcal{U}}^{\text{pp}}$  denote the recorded entry with  $s^* = s$  and  $c_{\text{hid}}^* = c_{\text{hid}}$ . Obviously,  $\text{C1.Open}(\text{CRS}_{\text{com}}^1, \text{pk}_{\mathcal{U}}^*, c_{\text{hid}}^*, d_{\text{hid}}^*) = 1$  holds, as otherwise the entry never would have been recorded. On the other hand, the user with  $\text{pid}_{\mathcal{U}}$  and  $\text{pk}_{\mathcal{U}}$  has never participated in a transaction with serial number  $s$ , because  $\mathcal{F}_{\text{P4TC}}$  returned  $\text{out}_{\mathcal{U}} = 0$  to the simulator. Hence,  $\text{pk}_{\mathcal{U}} \neq \text{pk}_{\mathcal{U}}^*$  follows which is a contradiction to the binding property of C1. An reduction to an adversary  $\mathcal{B}$  can be constructed the usual way.  $\square$

Taking all the aforementioned statements together, [Theorem F.22](#) from the beginning of this section follows. For the sake of formal completeness we recall it again.

**THEOREM F.22 (SYSTEM SECURITY).** *Under the assumptions of [Theorem F.1](#) and static corruption of a subset of the users*

$$\pi_{\mathcal{P}_{\text{P4TC}}}^{\mathcal{F}_{\text{CRS}}, \overline{\mathcal{G}}_{\text{bb}}} \geq_{\text{UC}} \mathcal{F}^{\overline{\mathcal{G}}_{\text{bb}}}$$

*holds.*

**PROOF.** A direct consequence of [Lemmas F.23, F.24](#) and [F.27](#) to [F.35](#).  $\square$

### F.3 Proof of User Security and Privacy

In this section we show the following theorem.

**THEOREM F.36 (USER SECURITY AND PRIVACY).** *Under the assumptions of [Theorem F.1](#) and static corruption of a subset of RSUs, TSP and SA*

$$\pi_{\text{P4TC}}^{\mathcal{F}_{\text{CRS}}, \overline{\mathcal{G}}_{\text{bb}}} \geq_{\text{UC}} \mathcal{F}_{\text{P4TC}}$$

holds.

The definition of the UC-simulator  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{user-sec}}$  for [Theorem F.36](#) can be found in [Figs. 51 to 55](#). Please note that while the real protocol  $\pi_{\text{P4TC}}$  lives in the  $(\mathcal{F}_{\text{CRS}}, \overline{\mathcal{G}}_{\text{bb}})$ -model, the ideal functionality  $\mathcal{F}_{\text{P4TC}}$  has no CRS. Hence, the CRS (but not the bulletin board) is likewise simulated, providing  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{user-sec}}$  with a lever to simulate the ZK proofs P1, P2, and P3, to equivocate C1, and to extract C2.

The overall proof idea is to define a sequence of hybrid experiments  $H_i$  together with simulators  $\mathcal{S}_i$  and protocols  $\pi_i$  such that the first hybrid  $H_0$  is identical to the real experiment and the last hybrid  $H_{12}$  is identical to the ideal experiment. Each hybrid is of the form

$$H_i := \text{EXEC}_{\pi_i, \overline{\mathcal{G}}_{\text{bb}}, \mathcal{S}_i, \mathcal{Z}^{\text{user-sec}}}(1^n).$$

Instead of directly proving indistinguishability of the real and ideal experiment we can break the proof down into showing indistinguishability of each pair of consecutive hybrids. We achieve this by demonstrating that whenever  $\mathcal{Z}^{\text{user-sec}}$  can distinguish between two consecutive hybrids with non-negligible probability this yields an efficient adversary against one of the underlying cryptographic assumptions. The general idea is that the protocol  $\pi_i$  that honest parties perform gradually declines from the real protocol  $\pi_0 = \pi_{\text{P4TC}}$  to a dummy protocol  $\pi_{12}$ , which does nothing but relay in- and outputs. At the same time  $\mathcal{S}_i$  progresses from a dummy adversary  $\mathcal{S}_0$  to the final simulator  $\mathcal{S}_{12}$  which can be split up into the ideal functionality and  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{user-sec}}$ . We proceed by giving concrete (incremental) definitions of all hybrids  $H_i$ .

*Hybrid  $H_0$ .* The hybrid  $H_0$  is defined as

$$H_0 := \text{EXEC}_{\pi_0, \overline{\mathcal{G}}_{\text{bb}}, \mathcal{S}_0, \mathcal{Z}^{\text{user-sec}}}(1^n)$$

with  $\mathcal{S}_0 = \mathcal{A}$  being identical to the dummy adversary and  $\pi_0 = \pi_{\text{P4TC}}$ . Hence,  $H_0$  denotes the real experiment.

*Hybrid  $H_1$ .* In hybrid  $H_1$  we modify  $\mathcal{S}_1$  such that  $\text{CRS}_{\text{pok}}$  is generated by SetupSPoK,  $\text{CRS}_{\text{com}}^1$  is generated by C1.SimGen and  $\text{CRS}_{\text{com}}^2$  is generated by C2.ExtGen. Additionally,  $\mathcal{S}_1$  initializes the internal sets  $\overline{\Omega}^{\text{dsp}}$ ,  $\overline{\Omega}_{\mathcal{U}}^{\text{pp}}$  and  $\overline{HTD}$  as empty sets and records the respective entries as the final simulator  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{user-sec}}$  does.

*Hybrid  $H_2$ .* Hybrid  $H_2$  replaces the code in the tasks DR/TSP/RSU/User Registration of the protocol  $\pi_2$  such that the simulator  $\mathcal{S}_2$  is asked for the keys instead. This equals the method in which the keys are generated in the ideal experiment.

*Hybrid  $H_3$ .* In hybrid  $H_3$  the task RSU Certification is modified. For an honest TSP or an honest RSU the code of  $\pi_3$  is replaced by the code of a dummy party. The simulator  $\mathcal{S}_3$  behaves in this case as the final simulator  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{user-sec}}$  would.

*Hybrid  $H_4$ .*  $H_4$  modifies the tasks of Wallet Issuing and Debt Accumulation. The code of  $\pi_4$  for the user is modified such that it does not send  $s'$  but randomly picks  $s$  and sends it to  $\mathcal{S}_4$ . Then  $\mathcal{S}_4$  extracts  $s'' \leftarrow \text{C2.Extract}(\text{CRS}_{\text{com}}^2, c''_{\text{ser}})$ , calculates  $s' := s \cdot (s'')^{-1}$  and inserts  $s'$  into the message from the user to the TSP or RSU respectively.

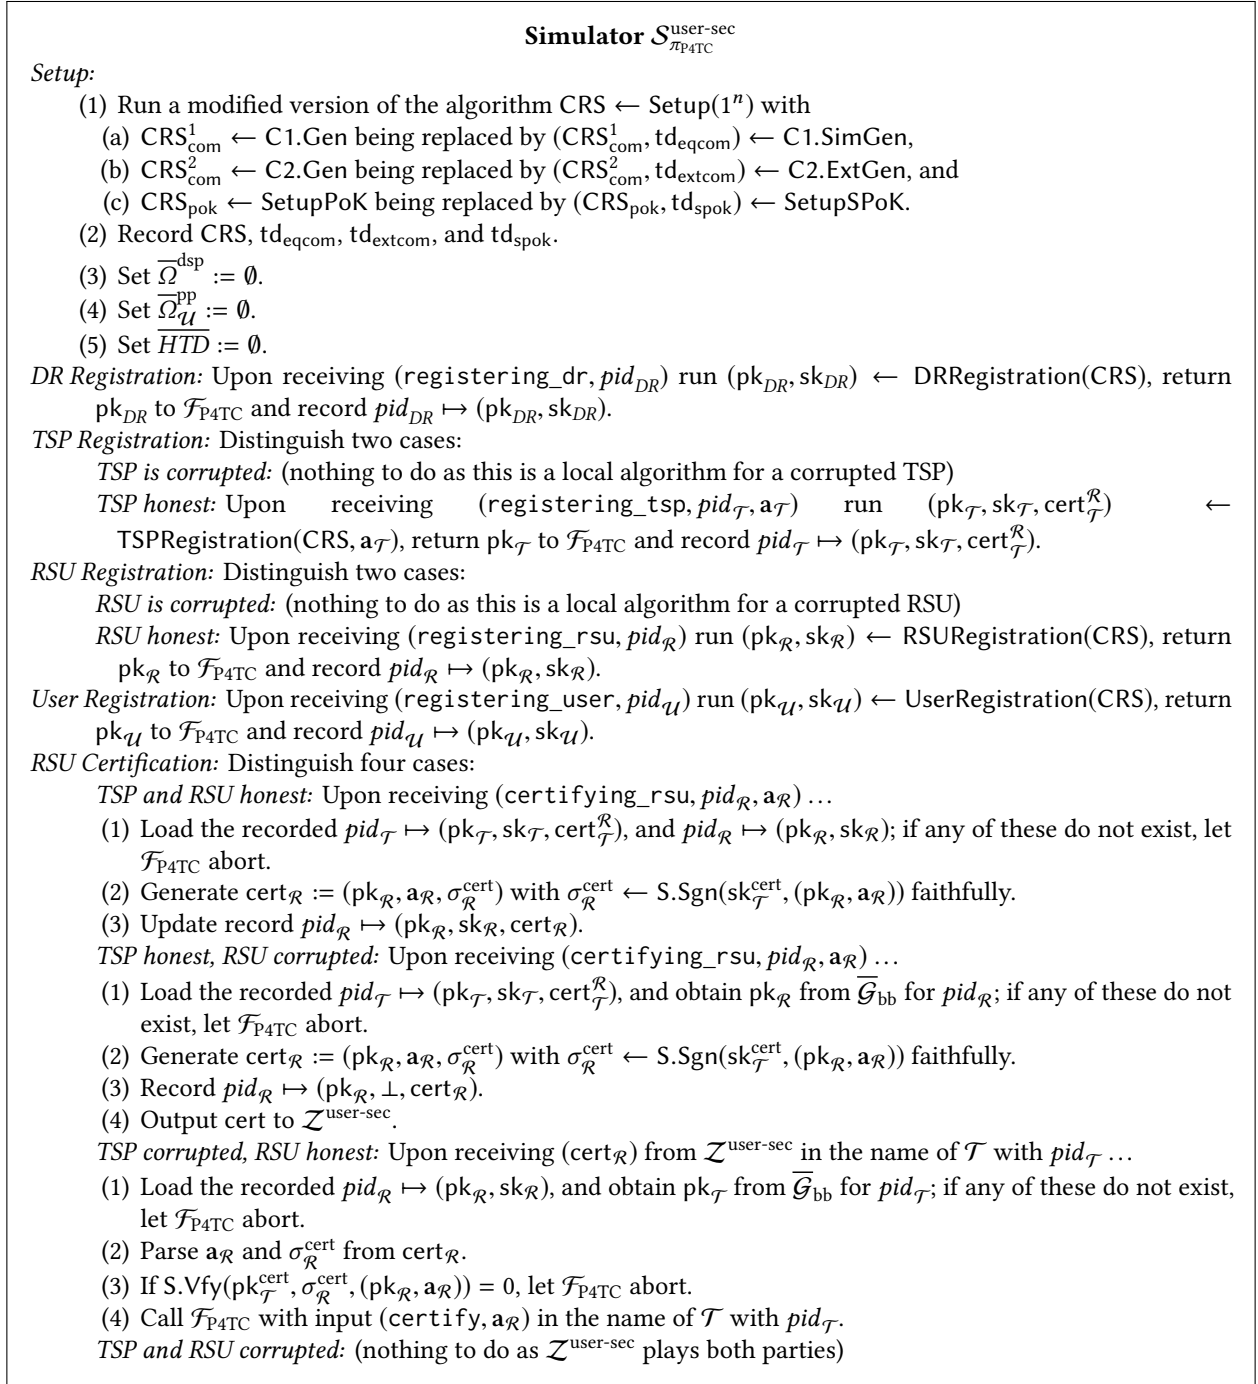


Fig. 51. The simulator for User Security and Privacy

**Simulator  $\mathcal{S}_{\pi_{P4TC}}^{\text{user-sec}}$  (cont.)**

*Wallet Issuing:* Distinguish two cases:

*TSP is honest:* (nothing to do)

*TSP is corrupted:*

- (1) Load the recorded  $pid_{\mathcal{U}} \mapsto (pk_{\mathcal{U}}, sk_{\mathcal{U}})$ , and obtain  $pk_{\mathcal{T}}$  from  $\overline{\mathcal{G}}_{bb}$  for  $pid_{\mathcal{T}}$ ; if any of these do not exist, let  $\mathcal{F}_{P4TC}$  abort.
- (2)  $(c'_{seed}, d'_{seed}) \leftarrow \text{C1.SimCom}(\text{CRS}_{com}^1)$
- (3) Send  $(pk_{\mathcal{U}}, c'_{seed})$  to  $\mathcal{Z}^{\text{user-sec}}$  as the 1<sup>st</sup> message from  $\mathcal{U}$  to  $\mathcal{T}$
- (4) Upon receiving  $(cert_{\mathcal{T}}^{\mathcal{R}}, a_{\mathcal{U}}, c''_{ser}, \lambda'')$  from  $\mathcal{Z}^{\text{user-sec}}$  in the name of  $\mathcal{T}$  with  $pid_{\mathcal{T}} \dots^a$ 
  - (a) Parse  $(pk_{\mathcal{T}}^{\mathcal{R}}, a_{\mathcal{T}}, \sigma_{\mathcal{T}}^{\text{cert}}) := cert_{\mathcal{T}}^{\mathcal{R}}$ .
  - (b) If  $S.Vfy(pk_{\mathcal{T}}^{\text{cert}}, \sigma_{\mathcal{T}}^{\text{cert}}, (pk_{\mathcal{T}}^{\mathcal{R}}, a_{\mathcal{T}})) = 0$  abort.
  - (c)  $\Lambda'' := g_1^{\lambda''}$
  - (d)  $s'' \leftarrow \text{C2.Extract}(\text{CRS}_{com}^2, \text{td}_{\text{extcom}}, c''_{ser})$ .
  - (e) Call  $\mathcal{F}_{P4TC}$  with input  $(\text{issue}, a_{\mathcal{U}}, \emptyset)$  in the name of  $\mathcal{T}$  with  $pid_{\mathcal{T}} \dots^b$
- (5) Upon receiving output  $(s)$  from  $\mathcal{F}_{P4TC}$  for  $\mathcal{T} \dots$ 
  - (a)  $s' := s \cdot s''^{-1}$ .
  - (b)  $r_1, r_2 \xleftarrow{\mathcal{R}} \mathbb{Z}_p$ .
  - (c)  $e^* \leftarrow \text{E1.Enc}(pk_{DR}, \underbrace{(1, \dots, 1)}_{\ell+2}; r_1, r_2)$
  - (d)  $(c_{\mathcal{T}}, d_{\mathcal{T}}) \leftarrow \text{C1.Com}(\text{CRS}_{com}^1, (0, 0))$ .
  - (e)  $(c_{\mathcal{R}}, d_{\mathcal{R}}) \leftarrow \text{C1.Com}(\text{CRS}_{com}^1, (0, 0, 0, 0))$ .
  - (f)  $stmnt := (pk_{\mathcal{U}}, pk_{DR}, e^*, c_{\mathcal{T}}, c_{\mathcal{R}}, c'_{seed}, \Lambda'', \lambda'')$ .
  - (g)  $\pi \leftarrow \text{P1.SimProof}(\text{CRS}_{pok}, \text{td}_{spok}, stmnt)$ .
  - (h) Send  $(s', e^*, c_{\mathcal{T}}, c_{\mathcal{R}}, \pi)$  to  $\mathcal{Z}^{\text{user-sec}}$  as the 2<sup>nd</sup> message from  $\mathcal{U}$  to  $\mathcal{T}$
- (6) Upon receiving  $(s'', d''_{ser}, \sigma_{\mathcal{T}}, \sigma_{\mathcal{R}})$  from  $\mathcal{Z}^{\text{user-sec}}$  in the name of  $\mathcal{T}$  with  $pid_{\mathcal{T}} \dots$ 
  - (a) If  $\text{C2.Open}(\text{CRS}_{com}^2, s'', c''_{ser}, d''_{ser}) = 0$ , let  $\mathcal{F}_{P4TC}$  abort.
  - (b) Set  $htd := (pk_{\mathcal{U}}, s, \lambda'', e^*)$  and insert  $htd$  into  $\overline{HTD}$ .
  - (c) Create real token  $\tau$  faithfully.
  - (d) If  $\text{WalletVerification}(pk_{\mathcal{T}}, pk_{\mathcal{U}}, \tau) = 0$ , let  $\mathcal{F}_{P4TC}$  abort.
  - (e) Let  $\mathcal{F}_{P4TC}$  return the delayed output to the User.

<sup>a</sup>If no message is received, let  $\mathcal{F}_{P4TC}$  abort; if blacklisted is received, override  $\mathcal{F}_{P4TC}$ 's delayed output for the User with blacklisted.

<sup>b</sup>Use empty set as blacklist. If the TSP intended to blacklist the user, the TSP would not have sent the previous message.

Fig. 52. The simulator for User Security and Privacy (cont. from Fig. 51)

*Hybrid H<sub>5</sub>.* This hybrid modifies  $\pi_5$  such that the honest parties do not send any proofs. Instead, the simulator  $\mathcal{S}_5$  appends a simulated proof to the message from a user to a TSP or RSU without knowing the witness.

*Hybrid H<sub>6</sub>.*  $H_6$  modifies  $\pi_6$  such that honest users do not send the commitments  $c'_{seed}$ ,  $c_{\mathcal{T}}$  and  $c_{\mathcal{R}}$  in the Wallet Issuing task and  $c'_{\mathcal{R}}$  in the Debt Accumulation task. Instead,  $\mathcal{S}_6$  injects suitable commitments to vectors of zeros. This equals the behavior of the final simulator  $\mathcal{S}_{\pi_{P4TC}}^{\text{user-sec}}$ .

**Simulator  $\mathcal{S}_{\pi_{P4TC}}^{\text{user-sec}}$  (cont.)**

*Debt Accumulation:* Distinguish two cases:

*RSU is honest:* (nothing to do)

*RSU is corrupted:*

- (1) Obtain  $\text{pk}_{\mathcal{T}}$  from  $\overline{\mathcal{G}}_{\text{bb}}$  for  $\text{pid}_{\mathcal{T}}$ ; if it does not exist, let  $\mathcal{F}_{P4TC}$  abort.
- (2) Upon receiving  $u_2, c''_{\text{ser}}, \text{cert}_{\mathcal{R}}$  from  $\mathcal{Z}^{\text{user-sec}}$  in the name of  $\mathcal{R}$  with  $\text{pid}_{\mathcal{R}}$ , do ...
  - (a) Parse  $(\text{pk}_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}}, \sigma_{\mathcal{R}}^{\text{cert}}) := \text{cert}_{\mathcal{R}}$ .
  - (b) If  $\text{S.Vfy}(\text{pk}_{\mathcal{T}}^{\text{cert}}, \sigma_{\mathcal{R}}^{\text{cert}}, (\text{pk}_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}})) = 0$  abort.
  - (c)  $s'' \leftarrow \text{C2.Extract}(\text{CRS}_{\text{com}}^2, c''_{\text{ser}})$ .
  - (d) Call  $\mathcal{F}_{P4TC}$  with input  $(\text{pay\_toll}, \emptyset)^a$  in the name of  $\mathcal{R}$  with  $\text{pid}_{\mathcal{R}}$ .
  - (e) Obtain RSU's output  $(s, \phi, \mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}})$  from  $\mathcal{F}_{P4TC}$ , and delay the output of the User.
  - (f)  $s' := s \cdot s''^{-1}$ .
- (3) Upon being requested by  $\mathcal{Z}^{\text{user-sec}}$  to provide the second message ...
  - (a) Run  $(c_{\text{hid}}, \overline{d}_{\text{hid}}) \leftarrow \text{C1.SimCom}(\text{CRS}_{\text{com}}^1)$  and append  $(s, c_{\text{hid}}, \overline{d}_{\text{hid}})$  to  $\overline{\Omega}_{\mathcal{U}}^{\text{pp}}$ .
  - (b)  $(c'_{\mathcal{R}}, d'_{\mathcal{R}}) \leftarrow \text{C1.Com}(\text{CRS}_{\text{com}}^1, (0, 0, 0, 0))$ .
  - (c) Check if any  $(\phi, t', u'_2) \in \overline{\Omega}^{\text{dsp}}$  has been recorded previously with  $(\phi)$  being used as key. If no, pick  $t \xleftarrow{\text{R}} \mathbb{Z}_p$ . If yes, load the recorded  $\text{pid}_{\mathcal{U}} \mapsto (\text{pk}_{\mathcal{U}}, \text{sk}_{\mathcal{U}})$  and set  $t := t' + \text{sk}_{\mathcal{U}}(u_2 - u'_2)$ . Insert  $(\phi, t, u_2)$  into  $\overline{\Omega}^{\text{dsp}}$ .
  - (d)  $\text{stmnt} := (\text{pk}_{\mathcal{T}}, \text{pk}_{\mathcal{T}}^{\text{cert}}, \phi, \mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}}, c_{\text{hid}}, c'_{\mathcal{R}}, t, u_2)$ .
  - (e)  $\pi \leftarrow \text{P2.SimProof}(\text{CRS}_{\text{pok}}, \text{td}_{\text{spok}}, \text{stmnt})$ .
  - (f) Output  $(s', \pi, \phi, \mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}}, c_{\text{hid}}, c'_{\mathcal{R}}, t)$  to  $\mathcal{Z}^{\text{user-sec}}$ .
- (4) Upon being ask to provide a price for  $(\mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}})$  return a price  $p$  as the dummy adversary would do.
- (5) Upon receiving  $(s'', d''_{\text{ser}}, c_{\mathcal{R}}, d''_{\mathcal{R}}, \sigma_{\mathcal{R}}, p)$  from  $\mathcal{Z}^{\text{user-sec}}$  ... <sup>b</sup>
  - (a)  $d_{\mathcal{R}} := d'_{\mathcal{R}} \cdot d''_{\mathcal{R}}$ .
  - (b) If  $\text{C1.Open}(\text{CRS}_{\text{com}}^1, (1, g_1^p, 1, g_1), c_{\mathcal{R}}, d_{\mathcal{R}}) = 0$  let  $\mathcal{F}_{P4TC}$  abort.
  - (c) If  $\text{S.Vfy}(\text{pk}_{\mathcal{R}}, \sigma_{\mathcal{R}}, (c_{\mathcal{R}}, s)) = 0$  let  $\mathcal{F}_{P4TC}$  abort.
  - (d) Let  $\mathcal{F}_{P4TC}$  return the delayed output to the User.

<sup>a</sup>Use empty set as blacklist.

<sup>b</sup>If no message is received, let  $\mathcal{F}_{P4TC}$  abort; if `blacklisted` is received, override  $\mathcal{F}_{P4TC}$ 's delayed output for the User with `blacklisted`.

Fig. 53. The simulator for User Security and Privacy (cont. from Fig. 51)

*Hybrid H<sub>7</sub>.* This hybrid introduces a lookup table that links hidden user trapdoors to their origin. More precisely, if the task Wallet Issuing is executed,  $\mathcal{S}_7$  records  $\overline{htd} = (\text{pk}_{\mathcal{U}}, s, \lambda'', e^*)$  in  $\overline{HTD}$ . Please note that  $\mathcal{S}_7$  knows  $s$  due to the change in  $H_4$ .

Moreover, if the task User Blacklisting is invoked,  $\mathcal{S}_7$  partitions the set of hidden user trapdoors  $\overline{HTD}_{\mathcal{U}}$  provided by the environment into two “subsets”<sup>34</sup>  $\overline{HTD}_{\mathcal{U}}^{\text{genuine}}$  and  $\overline{HTD}_{\mathcal{U}}^{\text{fake}}$  (cp. Fig. 55, Steps 2a to 2c). If for a hidden user trapdoor  $\overline{htd} = (\text{pk}_{\mathcal{U}}, s, \lambda'', e^*)$  a corresponding entry  $(\cdot, \cdot, \lambda'', e^*)$  is recorded in  $\overline{HTD}$ , then we call it a *genuine*

<sup>34</sup>The sets  $\overline{HTD}_{\mathcal{U}}^{\text{genuine}}$  and  $\overline{HTD}_{\mathcal{U}}^{\text{fake}}$  might be no actual “subsets”. The hidden user trapdoors are classified with respect to  $(\cdot, \cdot, \lambda'', e^*)$  which are left unmodified, but the first two components  $(\text{pk}_{\mathcal{U}}, s, \cdot)$  are sanitized.

**Simulator  $\mathcal{S}_{\pi_{P4TC}}^{\text{user-sec}}$  (cont.)**

*Debt Clearance:* Distinguish two cases:  
*TSP is honest:* (nothing to do)  
*TSP is corrupted:*

- (1) Load the recorded  $pid_{\mathcal{U}} \mapsto (pk_{\mathcal{U}}, sk_{\mathcal{U}})$ , and obtain  $pk_{\mathcal{T}}$  from  $\overline{\mathcal{G}}_{bb}$  for  $pid_{\mathcal{T}}$ ; if any of these do not exist, let  $\mathcal{F}_{P4TC}$  abort.
- (2) Upon receiving  $u_2$  from  $\mathcal{Z}^{\text{user-sec}}$  in the name of  $\mathcal{T}$  with  $pid_{\mathcal{T}}$ , do ...
  - (a) Call  $\mathcal{F}_{P4TC}$  with input (clear\_debt) in the name of  $\mathcal{T}$  with  $pid_{\mathcal{T}}$ .
  - (b) Obtain leaked  $a_{\mathcal{R}}^{\text{prev}}$ .
  - (c) Obtain TSP's output  $(pid_{\mathcal{U}}, \phi, b^{\text{bill}})$  from  $\mathcal{F}_{P4TC}$ , and delay the output of the User.
- (3) Upon being requested by  $\mathcal{Z}^{\text{user-sec}}$  to provide the second message ...
  - (a) Check if any  $(\phi, t', u'_2) \in \overline{\Omega}^{\text{dsp}}$  has been recorded previously with  $(\phi)$  being used as key. If no, pick  $t \xleftarrow{R} \mathbb{Z}_p$ . If yes, load the recorded  $pid_{\mathcal{U}} \mapsto (pk_{\mathcal{U}}, sk_{\mathcal{U}})$  and set  $t := t' + sk_{\mathcal{U}}(u_2 - u'_2)$ . Insert  $(\phi, t, u_2)$  into  $\overline{\Omega}^{\text{dsp}}$ .
  - (b)  $stmnt := (pk_{\mathcal{U}}, pk_{\mathcal{T}}, pk_{\mathcal{T}}^{\text{cert}}, \phi, a_{\mathcal{U}}, a_{\mathcal{R}}^{\text{prev}}, g_1^{b^{\text{bill}}}, t, u_2)$ .
  - (c)  $\pi \leftarrow \text{P3.SimProof}(\text{CRS}_{\text{pok}}, \text{td}_{\text{spok}}, stmnt)$ .
  - (d) Output  $(pk_{\mathcal{U}}, \pi, \phi, a_{\mathcal{U}}, a_{\mathcal{R}}^{\text{prev}}, b^{\text{bill}}, t)$  to  $\mathcal{Z}^{\text{user-sec}}$ .
- (4) Upon receiving (OK) from  $\mathcal{Z}^{\text{user-sec}}$ ,<sup>a</sup> let  $\mathcal{F}_{P4TC}$  return the delayed output to the User.

<sup>a</sup>If no message is received, let  $\mathcal{F}_{P4TC}$  abort.

Fig. 54. The simulator for User Security and Privacy (cont. from Fig. 51)

hidden user trapdoor and the public key  $pk_{\mathcal{U}}$  and the serial number  $s$  are set to the originally recorded value. Genuine hidden user trapdoors are those which have legitimately been created in the scope of Wallet Issue. Else we call it a *fake* hidden user trapdoor. In this case, the provided serial number  $s$  is left as is and the public key  $pk_{\mathcal{U}}$  is set to the decrypted value from  $e^*$ .<sup>35</sup>  $\mathcal{S}_7$  additionally checks if the hidden user trapdoors of both sets  $HTD_{\mathcal{U}}^{\text{genuine}}$  and  $HTD_{\mathcal{U}}^{\text{fake}}$  belong to the same user public key  $pk_{\mathcal{U}}^{\mathcal{T}}$  or else aborts. This equals the behavior of the final simulator.

$\mathcal{S}_7$  runs the code of an honest DR on  $HTD_{\mathcal{U}}^{\text{genuine}} \cup HTD_{\mathcal{U}}^{\text{fake}}$  to recover  $\Phi_{\mathcal{U}}$ , i.e., it decrypts every hidden user trapdoor and evaluates the PRF itself. For the DR the code of  $\pi_7$  is changed such that it simply signals its consent by forwarding its input  $pk_{\mathcal{U}}^{\text{DR}}$  to  $\mathcal{S}_7$ .

Hybrid  $H_7$  is a preparative step to eventually free the simulator from having to actually decrypt genuine hidden user trapdoors in case a user is blacklisted. Decryption of genuine hidden user trapdoors becomes impossible for the final simulator  $\mathcal{S}_{\pi_{P4TC}}^{\text{user-sec}}$  as hybrid  $H_9$  replaces all hidden user trapdoors with an encryption of zeros. However, a hidden user trapdoor is not bound to any user secrets. This allows a (malicious) TSP to pick an arbitrary chosen wallet ID  $\lambda^{\text{fake}}$  and create a syntactically valid hidden user trapdoor for  $\lambda^{\text{fake}}$  and any public user key  $pk_{\mathcal{U}}$ . If the DR and the (malicious) TSP agree to blacklist a user, the TSP may send these fake hidden user trapdoors in addition to the genuine hidden user trapdoors to the DR. In the real protocol the DR simply decrypts both types of hidden user trapdoors and returns a list of pseudo-random fraud detection IDs for each of them.<sup>36</sup> The final

<sup>35</sup>We assume, that Dec returns  $\perp$  if  $e^*$  cannot be decrypted.

<sup>36</sup>Skipping ahead, please note that this is not a "real" attack but only a very cumbersome way for the TSP to evaluate the PRF on self-chosen seeds.

**Simulator  $\mathcal{S}_{\pi_{P4TC}}^{\text{user-sec}}$  (cont.)**

*Prove Participation:*

- (1) Load the recorded  $pid_{\mathcal{U}} \mapsto (sk_{\mathcal{U}}, pk_{\mathcal{U}})$ ; if this does not exist let  $\mathcal{F}_{P4TC}$  abort.
- (2) Upon receiving  $S_{\mathcal{R}}^{\text{PP}}$  from  $\mathcal{Z}^{\text{user-sec}}$  in the name of SA ...
  - (a) Call  $\mathcal{F}_{P4TC}$  with input (prove\_participation,  $pid_{\mathcal{U}}, S_{\mathcal{R}}^{\text{PP}}$ ).
  - (b) Obtain the SA's output (out) from  $\mathcal{F}_{P4TC}$ .
  - (c) If out = NOK, abort.
  - (d) Pick  $\bar{w}_{\mathcal{U}}^{\text{PP}} = (s, c_{\text{hid}}, \bar{d}_{\text{hid}})$  from  $\bar{\mathcal{D}}_{\mathcal{U}}^{\text{PP}}$  such that  $s \in S_{\mathcal{R}}^{\text{PP}}$ ; if this does not exist abort.
  - (e) Equivoke  $d_{\text{hid}} \leftarrow \text{C1.Equiv}(\text{CRS}_{\text{com}}^1, sk_{\mathcal{U}}, c_{\text{hid}}, \bar{d}_{\text{hid}})$ .
  - (f) Output  $(s, c_{\text{hid}}, d_{\text{hid}})$  to  $\mathcal{Z}^{\text{user-sec}}$  as message from  $\mathcal{U}$  to SA.

*Double-Spending Detection:* Upon being ask to provide a proof for  $pid_{\mathcal{U}}$ , look up  $pid_{\mathcal{U}} \mapsto (pk_{\mathcal{U}}, sk_{\mathcal{U}})$ , and return  $sk_{\mathcal{U}}$ .

*Guilt Verification:* Upon being ask by  $\mathcal{F}_{P4TC}$  to provide out for  $(pid_{\mathcal{U}}, \pi)$  ...

- (1) Receive  $pk_{\mathcal{U}}$  from  $\bar{\mathcal{G}}_{\text{bb}}$  for  $pid_{\mathcal{U}}$ .
- (2) If  $g_1^\pi = pk_{\mathcal{U}}$ , then return out := OK, else out := NOK to  $\mathcal{F}_{P4TC}$ .

*User Blacklisting:* Distinguish two cases:

*TSP honest:* (nothing to do)

*TSP corrupted:*

- (1) Load the recorded  $pid_{DR} \mapsto (pk_{DR}, sk_{DR})$ .
- (2) Upon receiving  $HTD_{\mathcal{U}}$  from  $\mathcal{Z}^{\text{user-sec}}$  ...
  - (a)  $HTD_{\mathcal{U}}^{\text{genuine}} = \left\{ (pk_{\mathcal{U}}, s, \lambda'', e^*) \mid (\cdot, \cdot, \lambda'', e^*) \in HTD_{\mathcal{U}} \wedge (pk_{\mathcal{U}}, s, \lambda'', e^*) \in \overline{HTD} \right\}$
  - (b)  $HTD_{\mathcal{U}}^{\text{fake}} = \left\{ (pk_{\mathcal{U}}, s, \lambda'', e^*) \mid (\cdot, \cdot, \lambda'', e^*) \in HTD_{\mathcal{U}} \wedge (\cdot, \cdot, \lambda'', e^*) \notin \overline{HTD} \wedge (\dots, pk_{\mathcal{U}}) \leftarrow \text{E1.Dec}(sk_{DR}, e^*) \right\}$
  - (c) Assert  $pk_{\mathcal{U}}^{(1)} = pk_{\mathcal{U}}^{(2)}$  for all  $(pk_{\mathcal{U}}^{(1)}, \cdot, \cdot, \cdot), (pk_{\mathcal{U}}^{(2)}, \cdot, \cdot, \cdot) \in HTD_{\mathcal{U}}^{\text{genuine}} \cup HTD_{\mathcal{U}}^{\text{fake}}$  and set  $pk_{\mathcal{U}}^{\mathcal{T}} := pk_{\mathcal{U}}^{(1)}$ , else abort
  - (d) Call  $\mathcal{F}_{P4TC}$  with (blacklist\_user,  $pk_{\mathcal{U}}^{\mathcal{T}}$ ).
- (3) Upon being ask by  $\mathcal{F}_{P4TC}$  to provide  $S_{\text{root}}$  ...
  - (a)  $S_{\text{root}} := \{s \mid (\cdot, s, \cdot, \cdot) \in HTD_{\mathcal{U}}^{\text{genuine}}\}$
  - (b) Provide  $S_{\text{root}}$  to  $\mathcal{F}_{P4TC}$
- (4) Upon receiving TSP's output  $(b^{\text{bill}}, \Phi_{\text{bl}})$  from  $\mathcal{F}_{P4TC}$  ...
  - (a) For  $HTD_{\mathcal{U}}^{\text{fake}}$  recover  $\Phi^{\text{fake}}$  as the real DR would do
  - (b) Send  $\Phi_{\mathcal{U}} := \Phi_{\text{bl}} \cup \Phi^{\text{fake}}$  to  $\mathcal{F}_{P4TC}$  as the message from  $DR$  to  $\mathcal{T}$

Fig. 55. The simulator for User Security and Privacy (cont. from Fig. 51)

simulator needs to mimic this behavior. The ideal functionality always returns a list of uniformly drawn fraud detection IDs of the correct length *only* for legitimately issued wallets. Hence, the final simulator extends this list by fraud detection IDs for each of the fake hidden user trapdoors.

*Hybrid  $H_8$ .* This hybrid introduces a new incorruptible entity  $\mathcal{F}_{\phi\text{-rand}}$  into the experiment that is only accessible by honest users and the simulator through subroutine input/output tapes.<sup>37</sup>

<sup>37</sup>I.e., communication is confidential, reliable and trustworthy. One might think of this entity as a preliminary version of the eventual ideal functionality.



$\mathcal{F}_{\phi\text{-rand}}$  provides the following functionality: Internally,  $\mathcal{F}_{\phi\text{-rand}}$  manages a partial map  $f_{\Phi}$ , mapping pairs of wallet IDs  $\lambda$  and counters  $x$  to fraud detection IDs. Whenever an as yet undefined value  $f_{\Phi}(\lambda, x)$  is required,  $\mathcal{F}_{\phi\text{-rand}}$  defines  $f_{\Phi}(\lambda, x) := \text{PRF}(\lambda, x)$ . If a user requests a fraud detection ID  $\phi$  for  $(\lambda, x)$ ,  $\mathcal{F}_{\phi\text{-rand}}$  returns  $f_{\Phi}(\lambda, x)$  to the user. If an honest user inquires  $\mathcal{F}_{\phi\text{-rand}}$  for the first time for a fresh  $\lambda$ , the user has to also provide the corresponding serial number  $s$  of the current transaction.  $\mathcal{F}_{\phi\text{-rand}}$  internally records that  $s$  is associated with  $\lambda$ . If  $\mathcal{F}_{\phi\text{-rand}}$  is invoked by the simulator with input  $s$ ,  $\mathcal{F}_{\phi\text{-rand}}$  looks up the associated wallet ID  $\lambda$  and returns the set  $\{f_{\Phi}(\lambda, 0), \dots, f_{\Phi}(\lambda, x_{\text{bl}_R})\}$ .

An honest user running  $\pi_8$  does not evaluate the PRF to obtain a fraud detection ID  $\phi$ , but requests  $\mathcal{F}_{\phi\text{-rand}}$  to provide one.

If User Blacklisting is invoked, the simulator  $\mathcal{S}_8$  proceeds as follows: For hidden user trapdoors in the set  $HTD_{\mathcal{U}}^{\text{fake}}$  it still decrypts the seed and evaluates the PRF. However, for hidden user trapdoors  $(\text{pk}_{\mathcal{U}}, s, \lambda'', e^*)$  in the set  $HTD_{\mathcal{U}}^{\text{genuine}}$ ,  $\mathcal{S}_8$  it does not decrypt  $e^*$ , but requests  $\mathcal{F}_{\phi\text{-rand}}$  for the corresponding set of fraud detection IDs using  $s$ .

*Hybrid H<sub>9</sub>*. In the scope of Wallet Issuing  $\pi_9$  is modified such that honest users do not send  $e^*$ . Instead,  $\mathcal{S}_6$  injects an encrypted zero-vector as  $e^*$ .

*Hybrid H<sub>10</sub>*. Hybrid H<sub>10</sub> replaces the PRF inside  $\mathcal{F}_{\phi\text{-rand}}$  by truly random values. Whenever an as yet undefined value  $f_{\Phi}(\lambda, x)$  is required,  $\mathcal{F}_{\phi\text{-rand}}$  independently and uniformly draws a fresh random fraud detection id  $\phi$  and sets  $f_{\Phi}(\lambda, x) := \phi$ .

*Hybrid H<sub>11</sub>*. In H<sub>11</sub> Debt Accumulation and Debt Clearance are modified such that the simulator  $\mathcal{S}_{11}$  replaces  $t$  in the message from the user. If no  $(\phi, t', u'_2) \in \overline{\Omega}^{\text{dsp}}$  has been recorded previously,  $\mathcal{S}_{11}$  picks  $t \xleftarrow{R} \mathbb{Z}_p$ , else  $\mathcal{S}_{11}$  sets  $t := t' + \text{sk}_{\mathcal{U}}(u_2 - u'_2)$ . Finally,  $\mathcal{S}_{11}$  inserts  $(\phi, t, u_2)$  into  $\overline{\Omega}^{\text{dsp}}$ . This equals the behavior of the final simulator  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{user-sec}}$ .

*Hybrid H<sub>12</sub>*. The hybrid H<sub>12</sub> modifies Debt Accumulation and Prove Participation. In Debt Accumulation the simulator  $\mathcal{S}_{12}$  runs  $(c_{\text{hid}}, \bar{d}_{\text{hid}}) \leftarrow \text{C1.SimCom}(\text{CRS}_{\text{com}}^1)$  and appends  $(s, c_{\text{hid}}, \bar{d}_{\text{hid}})$  to  $\overline{\Omega}_{\mathcal{U}}^{\text{pp}}$ . In Prove Participation the code of  $\mathcal{S}_{12}$  for the honest user is replaced by a code that just checks if the user has a matching and correct  $(s^*, c_{\text{hid}}^*, \bar{d}_{\text{hid}}^*)$  and respectively sends OK or NOK to  $\mathcal{S}_{12}$ . If  $\mathcal{S}_{12}$  receives OK from the user, then it picks  $\omega_{\mathcal{U}}^{\text{pp}} = (s, c_{\text{hid}}, \bar{d}_{\text{hid}})$  from  $\overline{\Omega}_{\mathcal{U}}^{\text{pp}}$  such that  $s \in S_{\mathcal{R}}^{\text{pp}}$ . Furthermore, it runs  $d_{\text{hid}} \leftarrow \text{C1.Equiv}(\text{CRS}_{\text{com}}^1, \text{sk}_{\mathcal{U}}, c_{\text{hid}}, \bar{d}_{\text{hid}})$  and sends  $(s, c_{\text{hid}}, d_{\text{hid}})$  to  $\mathcal{T}$ . Again, this equals the behavior of the final simulator  $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{user-sec}}$ .

The combinations of all modifications from H<sub>0</sub> to H<sub>12</sub> yields

$$\begin{aligned} \text{H}_{12} &= \text{EXEC}_{\pi_{12}, \bar{\mathcal{G}}_{\text{bb}}, \mathcal{S}_{12}, \mathcal{Z}^{\text{user-sec}}}(\mathbf{1}^n) \\ &= \text{EXEC}_{\mathcal{F}_{\text{P4TC}}, \bar{\mathcal{G}}_{\text{bb}}, \mathcal{S}_{\pi_{\text{P4TC}}}^{\text{user-sec}}, \mathcal{Z}^{\text{user-sec}}}(\mathbf{1}^n). \end{aligned}$$

With these hybrids we can now give the proof of [Theorem F.36](#). We do not spell out all steps of the proofs in full detail, but rather sketch the necessary reductions.

#### PROOF OF [THEOREM F.36](#).

*From H<sub>0</sub> to H<sub>1</sub>*: This hop solely changes how the CRS is created during the setup phase. This is indistinguishable for  $\text{CRS}_{\text{pok}}$ ,  $\text{CRS}_{\text{com}}^1$ , and  $\text{CRS}_{\text{com}}^2$  (see the composable zero-knowledge property of [Definition C.5](#), the equivocality property and the extractability property of [Definition C.7](#), resp., condition (a) each).

*From H<sub>1</sub> to H<sub>2</sub>*: This hop does not change anything in the view of  $\mathcal{Z}^{\text{user-sec}}$  as  $\mathcal{S}_2$  runs the same key generation algorithm as the real protocol does for honest parties.

*From H<sub>2</sub> to H<sub>3</sub>*: Again, this hop only changes which party runs which part of the code, but has no effect on the view of  $\mathcal{Z}^{\text{user-sec}}$ .

*From  $H_3$  to  $H_4$ :* This hop does not change anything from the perspective of  $\mathcal{Z}^{\text{user-sec}}$  as C2 is perfectly extractable. The change is a purely syntactical one to push the simulator closer to  $\mathcal{S}_{\mathcal{T}P4TC}^{\text{user-sec}}$ .

*From  $H_4$  to  $H_5$ :* This game hop replaces the real proofs by simulated proofs. To show indistinguishability despite this change, we actually have to consider a sequence of sub-hybrids—one for each of the different ZK proof systems P1, P2 and P3. In the first sub-hybrid all proofs for P1 are replaced by simulated proofs, in the second sub-hybrid all proofs for P2 are replaced and finally all proofs for P3. Assume there exists  $\mathcal{Z}^{\text{user-sec}}$  that notices a difference between  $H_4$  and the first sub-hybrid. Then we can construct an adversary  $\mathcal{B}$  that has a non-negligible advantage  $\text{Adv}_{\text{POK}, \mathcal{B}}^{\text{pok-zk}}(n)$ . Internally,  $\mathcal{B}$  runs  $\mathcal{Z}^{\text{user-sec}}$  and plays the protocol and simulator for  $\mathcal{Z}^{\text{user-sec}}$ . All calls of the simulator to P1.Prove are forwarded by  $\mathcal{B}$  to its own oracle in the external challenge game which is either P1.Prove or P1.SimProof.  $\mathcal{B}$  outputs whatever  $\mathcal{Z}^{\text{user-sec}}$  outputs. The second and third sub-hybrid follow the same line, but this time  $\mathcal{B}$  internally needs to generate simulated proofs for the proof system that has already been replaced in the previous sub-hybrid. As  $\mathcal{B}$  gets the simulation trapdoor as part of its input in the external challenge game,  $\mathcal{B}$  can do so.

*From  $H_5$  to  $H_6$ :* In this hop the commitments  $c'_{\text{seed}}$ ,  $c_{\mathcal{T}}$ ,  $c_{\mathcal{R}}$  and  $c'_{\mathcal{R}}$  are replaced with commitments to zero-messages for every honest user. Again, the hop from  $H_5$  to  $H_6$  is further split into a sequence of sub-hybrids with each sub-hybrid replacing a single commitment in reverse order of appearance. Assume  $\mathcal{Z}^{\text{user-sec}}$  can distinguish between  $H_5$  and  $H_6$  with non-negligible advantage. This yields an efficient adversary  $\mathcal{B}$  against the hiding property of C1. Please note that none of the commitments are ever opened, hence in each sub-hybrid only a single message is replaced. Internally,  $\mathcal{B}$  runs  $\mathcal{Z}^{\text{user-sec}}$  and plays the role of all parties and the simulator for  $\mathcal{Z}^{\text{user-sec}}$ . Externally,  $\mathcal{B}$  plays the hiding game. First,  $\mathcal{B}$  guesses the index  $i$  of the sub-hybrid which lets  $\mathcal{Z}^{\text{user-sec}}$  distinguish. For the first  $(i - 1)$  commitments,  $\mathcal{B}$  commits to the true message. For the  $i^{\text{th}}$  commitment,  $\mathcal{B}$  sends the actual message and an all-zero message to the external challenger.  $\mathcal{B}$  embeds the external challenge commitment (either to the actual message or the all-zero message) as the  $i^{\text{th}}$  commitment. All remaining commitments are replaced by commitments to zeros.  $\mathcal{B}$  outputs whatever  $\mathcal{Z}^{\text{user-sec}}$  outputs.

*From  $H_6$  to  $H_7$ :* This hop is perfectly indistinguishable from the environment's perspective as the additional code executed by  $\mathcal{S}_7$  does not change the output. Note that the hidden user trapdoors are still recovered in the same way the real DR would. For hidden user trapdoors  $htd = (\text{pk}_{\mathcal{U}}, s, \lambda'', e^*)$  in  $HTD_{\mathcal{U}}^{\text{genuine}}$  the (outer) public key  $\text{pk}_{\mathcal{U}}$  is replaced by the public key that has originally been recorded for  $(\cdot, \cdot, \lambda'', e^*)$ . However, due to the correctness of E1 the ciphertext  $e^*$  determines a unique message (for a fix key pair  $\text{pk}_{DR}, \text{sk}_{DR}$ ) and thus the originally recorded  $\text{pk}_{\mathcal{U}}$  equals the one that  $e^*$  decrypts to. The additional, pairwise equality check for all public keys triggers an abort if and only if the real DR aborts as well.

*From  $H_7$  to  $H_8$ :* Again, this hop is purely syntactical. The inserted entity  $\mathcal{F}_{\phi\text{-rand}}$  is invisible for  $\mathcal{Z}^{\text{user-sec}}$ . Moreover,  $\mathcal{F}_{\phi\text{-rand}}$  still uses the real PRF to generate fraud detection IDs. However, this hop frees  $\mathcal{S}_8$  from the decryption of genuine hidden user trapdoors. Instead,  $\mathcal{S}_8$  uses the originally recorded serial number  $s$  of the associated Wallet Issuing task to look up the set  $\{\text{PRF}(\lambda, 0), \dots, \text{PRF}(\lambda, x_{\text{bl}_R})\}$ , if required. Again, this is possible due to the correctness of E1, i.e.,  $e^*$  uniquely determines  $\lambda$  and thus maps to a unique  $s$ .

*From  $H_8$  to  $H_9$ :* In this hop every encryption  $e^*$  of a wallet ID  $\lambda$  is replaced by an encryption of a 1-vector for every honest user. We further split this hop into a sequence of sub-hybrids, with each sub-hybrid replacing a single encryption in reverse order of appearance. Assume  $\mathcal{Z}^{\text{user-sec}}$  can distinguish between  $H_8$  and  $H_9$  with non-negligible advantage. This yields an efficient adversary  $\mathcal{B}$  against the IND-CCA security of the encryption scheme E1. Internally,  $\mathcal{B}$  runs  $\mathcal{Z}^{\text{user-sec}}$  and plays the role of all parties and the simulator for  $\mathcal{Z}^{\text{user-sec}}$ . Externally,  $\mathcal{B}$  plays the IND-CCA game. When  $\mathcal{B}$ —playing the role of the simulator—needs to provide the public key in the scope of DR Registration, it embeds the challenge key  $\text{pk}_{DR} := \text{pk}^C$ .  $\mathcal{B}$  needs to guess the index of the sub-hybrid that causes a non-negligible difference, i.e.,  $\mathcal{B}$  needs to guess which (user) wallet causes distinguishability. For the first  $(i - 1)$  invocations of Wallet Issuing,  $\mathcal{B}$  encrypts the true seed, in the  $i^{\text{th}}$  invocation  $\mathcal{B}$  embeds the external

challenge and  $\mathcal{B}$  encrypts a 1-vector for the remaining invocations of Wallet Issuing. If  $\mathcal{Z}^{\text{user-sec}}$  invokes the task Blacklist User and  $\mathcal{B}$  needs to restore the wallet ID, the following two cases may occur: a) The presented hidden user trapdoor is a genuine trapdoor. In this case  $\mathcal{B}$  uses its lookup table to recover the correct set of fraud detection IDs. b) The presented hidden user trapdoor is a fake trapdoor. In this case  $\mathcal{B}$  uses its decryption oracle of the external CCA-game to restore the wallet ID  $\lambda$  and to create a set of fraud detection IDs.  $\mathcal{B}$  outputs whatever  $\mathcal{Z}^{\text{user-sec}}$  outputs.

*From  $H_9$  to  $H_{10}$ :* In this hop the pseudo-random fraud detection IDs for honest users are replaced by uniformly drawn random IDs. Again, we proceed by introducing a sequence of sub-hybrids. In each sub-hybrid the fraud detection IDs for one particular wallet ID  $\lambda$  are replaced. If  $\mathcal{Z}^{\text{user-sec}}$  can distinguish between two of the sub-hybrids, this immediately yields an efficient adversary against the pseudo-random game as defined in [Definition C.13](#). Internally,  $\mathcal{B}$  runs  $\mathcal{Z}^{\text{user-sec}}$  and plays the protocol and simulator for  $\mathcal{Z}^{\text{user-sec}}$ . Externally,  $\mathcal{B}$  interacts with an oracle that is either a true random function  $R(\cdot)$  or a pseudo-random function  $\text{PRF}(\hat{\lambda}, \cdot)$  for an unknown seed  $\hat{\lambda}$ . Whenever  $\mathcal{B}$  playing  $\mathcal{F}_{\phi\text{-rand}}$  internally needs to draw a fraud detection ID for the particular wallet  $\lambda$ ,  $\mathcal{B}$  uses its external oracle.  $\mathcal{B}$  outputs whatever  $\mathcal{Z}^{\text{user-sec}}$  outputs. Please note, this argument crucially uses the fact that  $\mathcal{Z}^{\text{user-sec}}$  is information-theoretically independent of  $\lambda$ . The hidden user trapdoors  $e^*$  have already been replaced by encryptions of 1-vectors in the previous hybrid  $H_9$ . This enables the external challenger to pick any seed  $\hat{\lambda}$ .

*From  $H_{10}$  to  $H_{11}$ :* This hop is statistically identical. As long as no double-spending occurs, the user chooses a fresh  $u_1$  in every transaction and thus a single point  $(u_2, t)$  is information-theoretically independent from  $sk_{\mathcal{U}}$ .

*From  $H_{11}$  to  $H_{12}$ :* In this hop the simulator  $\mathcal{S}_{12}$  sends simulated commitments  $c_{\text{hid}}$  for the hidden user ID instead of commitments to the true values. Later,  $\mathcal{S}_{12}$  equivocates these commitments on demand to the correct  $pk_{\mathcal{U}}$ , if  $\mathcal{Z}^{\text{user-sec}}$  triggers Prove Participation. Again, if  $\mathcal{Z}^{\text{user-sec}}$  has a non-negligible advantage to distinguish between  $H_{11}$  and  $H_{12}$ , then an efficient adversary  $\mathcal{B}$  can be constructed against the hiding property and equivocality of C1. The reduction follows the same lines as in the hop from hybrid  $H_5$  to  $H_6$ .  $\square$