

Valerie Fetzter*, Max Hoffmann, Matthias Nagel, Andy Rupp*, and Rebecca Schwerdt*

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

Abstract: 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. 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. Additionally, we provide a prototypical implementation on realistic hardware which already features fairly practical performance figures. 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.

Keywords: Toll Collection, Privacy, Cyber-Physical Systems, Provable Security, Real-World Crypto, Universal Composability

DOI 10.2478/popets-2020-0046

Received 2019-11-30; revised 2020-03-15; accepted 2020-03-16.

*Corresponding Author: **Valerie Fetzter:** Karlsruhe Institute of Technology, E-mail: valerie.fetzter@kit.edu

Max Hoffmann: Ruhr University Bochum, E-mail: max.hoffmann@rub.de

Matthias Nagel: Karlsruhe Institute of Technology, E-mail: matthias.nagel@kit.edu

*Corresponding Author: **Andy Rupp:** University of Luxembourg, E-mail: andy.rupp@rub.de

1 Introduction

Electronic toll collection (ETC) is already deployed in many countries world-wide. A recent study [58] predicts a Compound Annual Growth Rate (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) by 2027 [37]. 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 are inherently violating user privacy. This encourages abuse: EZ-Pass records have been used as evidence in divorce lawsuits [1], EZ-pass transponders are abused to track drivers throughout New York City [6], and the Norwegian AutoPASS system allowed anyone to obtain a transcript of visited toll booths [30]. The only legitimate reason for a toll service provider (TSP) to store personalized location records is to bill customers. Personalized location records may also be needed for violation enforcement, but for this the SA (cf. Section 2) is responsible and not the TSP. 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 TSPs. 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].

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 [36, 39]. Concerning what the user is charged for, we distinguish between two major types of charging schemes:

*Corresponding Author: **Rebecca Schwerdt:** Karlsruhe Institute of Technology, E-mail: rebecca.schwerdt@kit.edu

Distance-based: 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 dynamically be adapted by other parameters. This charging scheme is typically used in urban areas—not only to finance road infrastructure but also for congestion and pollution management through dynamically adapted tolls.

There are two 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): This is the most widely used ETC technology worldwide and the de facto standard in Europe [39]. It is based on bidirectional radio communication between roadside units (RSUs) and on-board units (OBUs) installed in the vehicles. In conventional 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): Video tolling inherently violates privacy.

Global Navigation Satellite System (GNSS): With GNSS the OBU keeps track of the vehicle’s location 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.

Drawbacks of Existing Systems and Proposals

Independent of the technology used, systems deployed in practice build on identifying the user to charge him. Previous academic work on privacy-preserving toll collection mainly considers the GNSS setting and comes—apart from a few exceptions [8, 28, 29]—without any formal security analysis. Although DSRC is currently the most dominant setting in practice, and according to recent market forecasts [41], predicting an exponential growth, it will stay like this at least for the near future, 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. See [Section 1.1](#) for details.

P4TC

We propose a comprehensive security model as well as provably-secure and efficient protocols for privacy-friendly ETC in the DSRC setting. Note that our instantiation is not post-quantum secure as it relies on pairing-based cryptography. 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.

1.1 Related Work

In this section, we review the academic literature on 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. See [Table 1](#) for a comparison of the contemplated systems.

DSRC (OBU/RSU) Setting

Previous work [31, 47–49] 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. None of them comes with a formal security model and proof.

General problems involved with an e-cash approach are the assurance of sufficient funds and exact payments. The first means that the user or some mechanism needs to ensure that during a trip a user never runs out of e-coins to pay his toll. The problem of “exact payments” is well-known in the e-cash community and led to research on divisible e-cash [17] (s.t. coins can be split-up appropriately) and transferable e-cash [9] (s.t. the payee can return change). However, instantiations of these variants still are significantly more heavyweight than traditional e-cash which hampers the deployment in time-critical settings such as toll collection. Moreover, in the case of transferable e-cash there also is an impossibility result [21] negating “perfect anonymity”, which translated to the ETC setting means: if the entity issuing e-coins

Table 1. Comparison of Electronic Toll Collection Systems

System	Setting	Post-payments	Dynamic Pricing	Offline RSUs	No TPM in OBU	Fraud Detection	Blacklisting Mechanism	Formal Model and Proof	Prototype Implementation
[47–49]	DSRC	X	(✓) ²	X	X	✓	X	X	X
SPEcTRe [31]	DSRC	✓	X	n/a ⁴	✓	(✓) ⁶	X	X	✓
[12]	DSRC	X ¹	✓	X	✓	✓	✓	X	✓
P4TC	DSRC	✓	✓	(✓) ⁵	✓	✓	✓	✓	✓
VPriv [29, 61]	GNSS	✓	✓	n/a ⁴	✓	(✓) ⁶	X	✓	✓
PrETP [8]	GNSS	✓	(✓) ³	n/a ⁴	✓	(✓) ⁶	X	✓	✓
Milo [59]	GNSS	✓	✓	n/a ⁴	✓	(✓) ⁶	X	X	✓
[27, 28]	GNSS	✓	✓	n/a ⁴	✓	(✓) ⁶	X	✓	X
[40]	GNSS	✓	X	n/a ⁴	X	(✓) ⁶	X	X ⁷	X
[52]	GNSS	✓	✓	(✓) ⁵	✓	✓	X	X	X ⁸

¹ Pre-payment scheme with refunds

² In [47] the pricing is not dynamic, in [48, 49] it is dynamic

³ PrETP supports dynamic prices, but for practical reasons the set of possible prices must be kept “small”

⁴ Not applicable, since these systems do not have RSUs

⁵ RSUs are offline most of the time, they only sometimes need to be online to transmit data to a central server

⁶ Fraud detection is probabilistic, so fraud will not always be detected

⁷ Informal proof

⁸ No implementation is given, but performance is estimated

(usually the TSP) colludes with the RSUs, then e-coins given out as change by an RSU to a user can be linked to their spendings at other RSUs. Certainly, this reduces the suitability for an ETC setting.¹

In [47–49] multiple electronic road pricing systems specifically tailored for Low Emission Zones (LEZ) are proposed. In [47] a user drives by an RSU and directly pays some price depending on this RSU. In [48, 49] 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.

SPEcTRe [31] is a simple access-based toll collection system, where RSA-based tokens/coins are used to pay toll while driving. Some ideas are presented to keep the double-spending database small.

In [12], the authors propose an efficient 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 a reimbursement is done. Unfortunately, the description of their system misses some important details. For instance, it is unclear how the zero-knowledge proof generated by the user in their token issuance protocol can be efficiently instantiated as it mixes statements and variables over a Paillier ring and a discrete logarithm group. As opposed to our scheme, their system relies on online “over-spending” detection and lacks a formal model and security proof.²

GNSS Setting

A variety of GNSS-based system proposals can be found in the literature. Here, the OBU equipped with GPS and GSM typically collects location-time data or road segment prices and sends this data to the TSP. To ensure that the user behaves honestly and, e.g., does not

¹ Note that relaxations of “perfect anonymity” can be achieved [9]. It would be interesting to see if those allow for practical ETC constructions with reasonable privacy guarantees.

² For the proofs the authors refer to the full version, which, however, does not seem to be publicly available.

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 his calculation to be correct.

In VPriv [61], 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 downloads the database of all tagged location-time tuples, calculates his total fee, and proves that for this purpose, he correctly used the tuples belonging to his tags without leaking those. In [29], ProVerif is used to verify VPriv’s privacy guarantees for honest-but-curious adversaries.

In the PrETP [8] 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.

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. Choosing an appropriate group size is difficult (the larger the size, the better the anonymity, but the computation overhead rises), as well as choosing a suited group division policy (users in a group should have similar driving regions and simi-

lar driving patterns). In [28], the system’s security and privacy properties are verified using ProVerif.

Another cell-based road pricing scheme is presented in [40]. Here, a roadpricing area is divided into cells and certain cells 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 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 key issue with all of the above systems is that their security relies on a strong assumption: invisible spot checks at unpredictable locations 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 anymore in terms of infrastructure requirements compared to the DSRC setting.

1.2 Our Contribution

The contribution of this work is threefold:

Protocols

A major challenge of this work was to select, adapt, and combine the numerous cryptographic building blocks and techniques to design a system satisfying simulation-based security and practicality *at the same time*.

To this end, we started from a payment system building-block called black-box accumulation (BBA+) [44]. BBA+³ 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.⁴ 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 (P9) partially) of the desired properties identified in Section 2, it has not been designed to cope with the real-world issues of an ETC scenario. Thus, we have to enhance BBA+ in several ways to enable its usage in our scenario in the first place.

For instance, consider the case an OBU, which contains the BBA+ wallet, is (claimed to) be broken. As a user collects debt with this wallet and those transactions are perfectly anonymous, the TSP is not able to recalculate what the user owes (P7) nor blacklist his wallet (P6). As this could result in a serious loss of revenue for the TSP, it renders the original BBA+ scheme impractical for an ETC scenario. We solve this by having an individual anonymity revocation trapdoor for each wallet. In case of an incident (like a broken OBU or possibly court-ordered assistance to law enforcement) this makes its transactions forward and backward linkable with the help of a trusted dispute resolver (DR). The trapdoor does not allow to link transactions of other wallets. To realize this, we adopt an idea from the e-cash literature [18]. More precisely, we make use of a PRF applied to a counter value 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 deposits the key encrypted under the DR's public key. The correctness of this deposit is ensured by 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 .

As a minor but very useful modification we added (user and RSU) attributes, which get signed along with

the wallet to protect from forgery. This allows to bind wallets to a billing period encoded in the attribute. Having RSUs only accept wallets from the current period reduces the size of the blacklist it has to check, which enables faster 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 arbitrarily create new wallets, forge user attributes and balances, etc. In order to mitigate this problem (P8), 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 for a user to prove its participation in certain transactions (P10) and to enable simulation-based security.

System Definition, Security Model and Proof

Having the scenario from Section 2 at the back of our mind, we propose a system definition and security model for post-payment toll collection systems based on the UC framework [24]. Our work is one of very few combining a complex, yet practical crypto system with a thorough UC security analysis.

The security of BBA+ has been modeled by individually formalizing and proving security properties from a list. This approach is common in the game-based setting but bears the intrinsic risk that important but non-obvious security aspects are overlooked, i.e., the list is incomplete. This danger is eliminated by our UC approach where we do not formalize a list of individual properties but rather what an ideal system looks like.

³ Note that we cannot build on the faster version of BBA+ [46] as we aim for a proof in the UC model which collides with their need to rewind adversaries.

⁴ Unfortunately, *online* double-spending detection is not feasible in this scenario (even if we required permanent online connection from RSUs) due to the strong time constraints of open road tolling and potentially simultaneously conducted interactions.

A challenging task was to find a formalization of such an ideal functionality which accomplishes a reasonable trade-off between various aspects: On one hand, it needs to be sufficiently abstract to represent the semantics of the goal “toll collection” while still admitting 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 security proofs 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 deduce the security of the system from the composition theorem. We refrain from this approach, as it entails a slew of technical subtleties due to the shared state between 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, is very complex, and technically demanding.

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 [62]. 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. More details on the implementation and performance figures can be found in [Section 6](#).

2 Considered Scenario

This sketch of the scenario shows the flexibility and complexity required of our system. We target post-payment toll collection in a DSRC setting which allows access- as well as distance-based charging and can be deployed in a toll-plaza as well as an open-road environment.

Parties

Our scenario involves the following entities:

TSP: A Toll Collection Service Provider (TSP) which may be a private company.

SA: A State Authority (SA), e.g., Department of Transportation, which outsourced toll collection to the TSP but is responsible for violation enforcement.

DR: A Dispute Resolver (DR), e.g., a privacy-protecting NGO. It is involved in case of incidents or disputes.⁵

User: Users who participate in the system by means of a (portable or mounted) On-Board Unit (OBU). OBUs are used for on-road transactions and in 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.

RSU: Road-Side Units (RSUs) which interact with OBUs, 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 only assume them to be periodically online for a short duration (e.g., at night when there is not much traffic) to exchange data with the TSP.

Camera: 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.

Main Protocols

In the following, we sketch the main protocols of the system involving the parties from above. [Figure 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 monthly billing periods.

User Registration: To participate in the system, a user needs to register with the TSP using some physi-

⁵ Note that we assume the DR to be trusted by all other parties as it implements a key escrow mechanism.




Fig. 1. The P4TC System Model

cal ID (e.g., passport, 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 bills.

RSU Certification: RSU gets a certificate from the TSP.

Wallet Issuing: An (empty) wallet is issued to a user by the TSP. The wallet is bound to the user’s ID and attributes (see next paragraph) in a privacy-preserving manner.

Debt Accumulation: Every time a user passes an RSU, his OBU and the RSU execute the Debt Accumulation protocol. The due toll is determined and added to the user’s wallet—possibly along with public attributes of the RSU. The toll may be dynamic and depend on different factors like the current time of day, congestion level, 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 or a car in the range of the RSU refuses to communicate at all. Due to technical limitations, it might be impossible to determine which car triggered the camera, especially in open-road tolling environments [50]. In this case, photos of more than one car 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, users participate in an (asynchronous) clearance phase with the TSP. As this protocol is not anonymous, users can be penalized if they refuse to run the protocol within a certain time frame. In a protocol execution, a user presents his wallet and the accumulated debt to the TSP. Then he may clear his debt immediately or within a certain grace period. A successful protocol execution invalidates his wallet. He can get a new one by rerunning Wallet Issuing.⁶

Double-spending Detection: As the system is largely offline, a malicious user might re-use an old state of his wallet (e.g., with lower debt) and thus commit double-spending. To mitigate this problem, Wallet Issuing, Debt Accumulation and Debt Clearance generate double-spending information that is eventually collected by the TSP. The TSP periodically runs Double-Spending Detection on its database to find pairs of matching double-spending information and to identify fraudulent users.

Blacklisting and Recalculation: With the help of the trusted DR, the TSP is able to blacklist fraudulent users and to recalculate what they owe.

Attributes, Pricing Function and Privacy Leakage

Our scenario involves two types of attribute vectors: user attributes and (previous) RSU attributes. To keep our framework flexible, we do not stipulate which kind of attributes are used. Those details depend on the concrete pricing model. We expect, however, that for most

⁶ If a user plans to be inactive the upcoming billing periods, he would request a new wallet not until he plans to become active again. But this requires to plan inactivity periods in advance. Alternatively, the TSP could be allowed to request the user’s trapdoor for such a period from the DR, which would reveal an empty transaction history.

scenarios very little information needs to be encoded in these attributes. For instance, access-based toll collection can be realized with prices primarily depending on auxiliary input like location, time and possibly the congestion level—without any previous RSU attributes and only the current billing period as user attribute. Including the previous RSU’s attributes into the toll calculation allows for 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.⁷ 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 RSU attribute. Likewise, to enforce that a user eventually pays his debt, the user attributes should encode the billing period.⁸

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 (cp. [Appendices A.3](#) and [B](#)).

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, a (corrupted) RSU could try to break the privacy of a user by charging a peculiar price that differs from the prescribed pricing function. We assume that the user verifies the validity of the price after each transaction and files a claim if the price was computed incorrectly.

⁷ This links entry and exit point. Our system still ensures anonymity and that multiple entry/exit pairs are unlinkable.

⁸ Clearly, for privacy reasons, unique expiration dates in attributes need to be avoided.

Desired Properties

The following list informally summarizes some desirable high-level properties that one would reasonably expect from an electronic toll collection system. These properties inspire the eventual definition of the ideal functionality in [Section 3](#). Note that the ideal functionality (and not this list) formally determines the security of our proposed protocol. In [Appendix A.3](#) we show how these high-level goals are reflected in the ideal functionality.

- (P1) **Owner-binding:** A user may only use a wallet legitimately issued to him.
- (P2) **Attribute-binding:** In Debt Accumulation, a user cannot pretend that he owes less by forging the attributes attached to his wallet.
- (P3) **Balance-binding:** In Debt Clearance, a user cannot claim to owe less than the amount added to his wallet unless he has committed double-spending.
- (P4) **Double-spending Detection:** If a user reuses an old copy of his wallet, he will be identified.⁹
- (P5) **Participation Enforcement:** If a user fails to participate in Debt Accumulation, he will be identified.
- (P6) **Blacklisting:** The TSP is able—with a hint from the DR—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.
- (P7) **Debt Recalculation:** The TSP is able—with a hint from the DR—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.
- (P8) **Renegade Expulsion:** As an RSU’s secrets enable tampering with user debt, there is a mechanism to mitigate financial loss due to compromised RSUs.
- (P9) **Unlinkability:** If user attributes and (previous) RSU attributes are ignored, a collusion of TSP and

⁹ Note that this also applies to multiple instances of double-spending. Since the TSP only periodically checks for double-spending, a user could reuse an old copy of his wallet several times, say 100 times, before the TSP runs the Double-spending Detection task. In that case all 100 double-spending will be identified and the user can be punished accordingly. Note that hardware cloning is also captured by this property. An adversary cloning his own/stolen hardware (including secrets) may coordinate simultaneous double-spending with many vehicles. All these double-spending will be detected and penalized.

RSUs may not be able to link a set of Wallet Issue, Debt Accumulation and Debt Clearance transactions of an honest user given that he is not blacklisted nor committed double-spending. More precisely, Wallet Issuing, Debt Accumulation and Debt Clearance do not reveal any information (except for user attributes, RSU attributes and the final balance in case of Debt Clearance) that may help in linking transactions.

- (P10) **Participation Provability:** Prove Participation enables the SA to deanonymize a single transaction of an honest user in case of an incident.
- (P11) **Protection Against False Accusation:** The user is protected against false accusation of having committed double-spending (cp. (P4)).

3 System Definition

We model our system P4TC within the UC-framework by Canetti [22], a simulation-based security notion. In the UC-framework, an ideal functionality \mathcal{F} , which is acting as a trusted third party (TTP), is defined which plainly solves the problem at hand in a secure and privacy-preserving manner. A protocol π is a *secure realization* of this ideal functionality \mathcal{F} if no PPT-machine \mathcal{Z} —called the *environment*—can distinguish between two experiments: the *real experiment* (running π) and the *ideal experiment* (using \mathcal{F}).

In the *real experiment*, \mathcal{Z} interacts with parties running the actual protocol π 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. The adversary \mathcal{A} is instructed by \mathcal{Z} and represents \mathcal{Z} 's interface to the network, e.g., \mathcal{A} reports all messages generated by any party to \mathcal{Z} and can manipulate, reroute, inject and/or suppress messages on \mathcal{Z} 's order. Moreover, \mathcal{Z} may instruct \mathcal{A} to corrupt parties. 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 π .

In the *ideal experiment* the protocol parties are mere dummies that pass their input to the TTP \mathcal{F} and hand over \mathcal{F} 's output as their own output. The ideal functionality \mathcal{F} is incorruptible and executes the task in a trustworthy manner. 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).

If no environment \mathcal{Z} can distinguish executions of the real and ideal experiment, any attack on the real experiment is also possible in the ideal one. Therefore the protocol π guarantees the same level of security as the (inherently secure) ideal functionality \mathcal{F} .

Regarding privacy, note that all parties (incl. the simulator) use the ideal functionality as a black-box and only know what it explicitly allows them to know as part of their output. 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 output of the ideal functionality.

The Ideal Functionality $\mathcal{F}_{\text{P4TC}}$

Due to space limitations, we only give a brief sketch in this section. More detailed information on the ideal functionality $\mathcal{F}_{\text{P4TC}}$ can be found in [Appendix A](#).

The key idea behind $\mathcal{F}_{\text{P4TC}}$ is to keep track of all transactions in a pervasive transaction database $TRDB$. Each entry $trdb$ 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,¹⁰ the price p (toll) of this particular transaction and the user wallet's total balance b , which is the accumulated sum of all transaction prices up to and including p . In other words, $\mathcal{F}_{\text{P4TC}}$ implements a trustworthy global bookkeeping service managing the wallets of all users. Each transaction entry is uniquely identified by a serial number s . Additionally, each entry contains a pointer s^{prev} to the logically previous transaction, a unique wallet ID λ , a counter x indicating the number of previous transactions with this particular wallet, and a fraud detection ID ϕ .

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

¹⁰ 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., this 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 definition of the UC framework.

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 and commit double-spending without being noticed *instantly* and double-spending is eventually detected after-the-fact.

In order to accurately define security, these technicalities have to be incorporated into $\mathcal{F}_{\text{P4TC}}$, which causes the bookkeeping to be more involved. The transaction database is best depicted as a directed graph. Nodes are identified by serial numbers s and additionally labeled with $(\phi, x, \lambda, \text{pid}_{\mathcal{U}}, b)$. Edges are given by (s^{prev}, s) and additionally labeled with $(\text{pid}_{\mathcal{R}}, p)$. A user’s wallet is represented by the subgraph of all nodes with the same wallet ID λ and forms a connected component. Unless a user commits double-spending the particular subgraph is a linked, linear list. In this case, each transaction entry has a globally unique fraud detection ID ϕ . If a user misbehaves and reuses an old wallet state (i.e., there are edges (s^{prev}, s) and (s^{prev}, s')), the corresponding subgraph becomes a directed tree. In this case, all transaction entries that constitute double-spending, i.e., all nodes with the same predecessor, share the same fraud detection ID ϕ . To consistently manage fraud detection IDs, $\mathcal{F}_{\text{P4TC}}$ uses a counter x and an injective map $f_{\Phi} : \mathcal{L} \times \mathbb{N}_0 \rightarrow \Phi, (\lambda, x) \mapsto \phi$. For any newly issued wallet with ID λ , 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_{Φ} maps a transaction to its fraud detection ID ϕ , i.e., $f_{\Phi}(\lambda, x) = \phi$.

Besides storing transaction data, $\mathcal{F}_{\text{P4TC}}$ also keeps track of parties’ attributes by internally storing RSU attributes $\mathbf{a}_{\mathcal{R}}$ upon certification and user (or rather wallet) attributes $\mathbf{a}_{\mathcal{U}}$ when the wallet is issued.

4 System Instantiation

This section briefly introduces the basic idea of P4TC and provides some details of the main protocols. For the full protocols see [Appendix D](#).

We start with some remarks on cryptographic building blocks and the algebraic setting. For formal definitions we refer to [Appendix C](#).

Cryptographic Building Blocks: Our construction makes use of 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 proof (which is Groth-Sahai [42] in our case).

Algebraic Setting: Our system instantiation is based on an asymmetric bilinear group setting $(G_1, G_2, G_T, e, \mathfrak{p}, g_1, g_2)$. Here, G_1, G_2, G_T are cyclic groups of prime order \mathfrak{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 as well as on the security of the building blocks in this setting.

Secure (Authenticated) Channels: All messages are encrypted using CCA-secure encryption. To this end a new session key is chosen by the user and encrypted under the public key of the RSU/TSP for each interaction. We omit these encryptions in the following. Apart from Debt Accumulation, we furthermore assume all channels to be authenticated.

Basic Idea

The basic technique underlying P4TC is a commit-sign-rerandomize-prove approach: a wallet is essentially a commitment c to the balance b (and some important auxiliary information) which is initially certified, i.e., signed, by the TSP during Wallet Issuing. To update b during Debt Accumulation, c cannot be sent over to the RSU as is since this would allow to link transactions. Instead it has to be rerandomized, i.e., the user generates a fresh commitment c' on b to send it over for updating. However, c' is not certified (directly) but only the original commitment c . Here a NIZK proof comes into play to show that c' is certified indirectly, i.e., that c' is just a new commitment to the same balance b as contained in c for which the user knows a signature. If the RSU is convinced that this proof is correct, the RSU updates c' by the price p the user has to pay using the homomorphic property of the commitment and signs c' . In this way the user obtains a new certified wallet state. Debt Clearance essentially works the same as Debt Accumulation except that the user reveals b to the TSP.

To incorporate blacklisting & recalculation as well as double-spending detection capabilities, c has to be augmented by additional information. For blacklisting & recalculation, the user must additionally commit to a PRF key λ and counter x as well as deposit λ with the DR during Wallet Issuing and prove that this has been done correctly. In the x -th transaction, the value $\text{PRF}(\lambda, x)$ now serves as the wallet’s fraud detection ID ϕ which is revealed to the RSU/TSP. In the same transaction, the new wallet state c' for the upcoming transac-

tion $x + 1$ is generated in which b is increased by p and the old counter value x by one. This is done in a similar manner as for the simpler case before. The NIZK proof now additionally enforces that the correct PRF key and counter value is used. Computing the fraud detection ID using a PRF has the advantage that the different (previous and future) states of a wallet are traceable given λ but untraceable if λ is unknown.

To encourage the user to only use his most recent wallet state by means of double-spending detection, each wallet state is bound to a random straight line encoding the user's secret key as slope: $t = \text{sk}_{\mathcal{U}}u_2 + u_1$. To this end, the user chooses some random u_1 for his wallet state and the parameters determining the straight line, i.e., $\text{sk}_{\mathcal{U}}$ and u_1 , are added to the commitment c in Wallet Issuing (resp. c' in Debt Accumulation). Now in each transaction using this wallet state, the user is enforced (by means of the NIZK) to reveal a point (u_2, t) for random u_2 on this straight line. Clearly, if only one such point is revealed, this leaks nothing about the slope. However, if two (different) points are revealed, so the wallet state is used twice, then the slope $\text{sk}_{\mathcal{U}}$ can be determined. This allows to compute $\text{pk}_{\mathcal{U}}$ identifying the cheating user. Note that transactions involving the same wallet state can be recognized by exhibiting the same fraud detection ID ϕ (see before).

Some Details on Wallet Structure and Protocols

Fig. 1 gives an overview of the protocols and parties.

System Setup: In the setup phase, a TTP generates the bilinear group and a public common reference string.¹¹ The latter contains parameters for the NIZK and commitment scheme(s) we use.

DR/TSP/RSU Key Generation: The DR generates a key pair $(\text{pk}_{DR}, \text{sk}_{DR})$ for an IND-CCA secure encryption scheme, where pk_{DR} is used to deposit a user-specific trapdoor (PRF-key) which allows to link the user's transactions in case of disputes. An individual signature key pair $(\text{pk}_{\mathcal{R}}, \text{sk}_{\mathcal{R}})$ is generated for each RSU to sign the updatable part (see below) of a user's wallet. Moreover, the TSP generates key pairs $(\text{pk}_{\mathcal{T}}^{\mathcal{T}}, \text{sk}_{\mathcal{T}}^{\mathcal{T}})$, $(\text{pk}_{\mathcal{T}}^{\text{cert}}, \text{sk}_{\mathcal{T}}^{\text{cert}})$ and $(\text{pk}_{\mathcal{T}}^{\mathcal{R}}, \text{sk}_{\mathcal{T}}^{\mathcal{R}})$ for an EUF-CMA secure signature scheme. The key $\text{sk}_{\mathcal{T}}^{\mathcal{T}}$ is used to sign fixed user-specific information (see

below) when a new wallet is issued. Using $\text{sk}_{\mathcal{T}}^{\mathcal{R}}$ the TSP plays the role of the initial RSU signing the updatable part of a new wallet.

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 $\text{pk}_{\mathcal{R}}$, 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 $\text{sk}_{\mathcal{T}}^{\text{cert}}$.

User Registration: To participate in the system, a user needs to generate a key pair $(\text{pk}_{\mathcal{U}} = g_1^{\text{sk}_{\mathcal{U}}}, \text{sk}_{\mathcal{U}}) \in G_1 \times \mathbb{Z}_p$. We assume that the TSP out-of-band binds $\text{pk}_{\mathcal{U}}$ to a verified physical user ID, in order to hold the user liable in case of a misuse. The key pair $(\text{pk}_{\mathcal{U}}, \text{sk}_{\mathcal{U}})$ is used to bind a wallet to a user.

Wallet Structure: A user's wallet essentially consists of two signed commitments $(c_{\mathcal{T}}, c_{\mathcal{R}})$, where $c_{\mathcal{T}}$ represents the fixed part and $c_{\mathcal{R}}$ the updatable part. Accordingly, the fixed part is signed along with the user's attributes $\mathbf{a}_{\mathcal{U}}$ by the TSP using $\text{sk}_{\mathcal{T}}^{\mathcal{T}}$ 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 $\text{sk}_{\mathcal{R}}$. The fixed part $c_{\mathcal{T}} = \text{Com}(\lambda, \text{sk}_{\mathcal{U}})$ is a commitment on the PRF key λ (used as wallet ID) and the user's secret key $\text{sk}_{\mathcal{U}}$. The updatable part $c_{\mathcal{R}} = \text{Com}(\lambda, b, u_1, x)$ also contains λ (to link both parts), the current balance b (debt), some user-chosen randomness u_1 to generate double-spending tags for the current wallet state, and a counter value x being the input to the PRF.

Wallet Issuing: This protocol is executed between user and TSP to create a new wallet with a fresh wallet ID λ and balance 0. Additionally, the PRF key λ is deposited under the DR's public encryption key. To generate the wallet, the user encodes λ together with his other secrets into the wallet. The PRF key λ 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 $\lambda := \lambda' + \lambda''$ is fixed and the user knows his own share λ' as well as the TSP's share λ'' . Then the user computes the commitments $c_{\mathcal{T}} = \text{Com}(\lambda, \text{sk}_{\mathcal{U}})$ and $c_{\mathcal{R}} := \text{Com}(\lambda, b := 0, u_1, x := 0)$.

Additionally, he prepares the deposit of λ . This is a bit tricky, as the user needs to prove that the ciphertext he gives to the TSP is actually an encryption of λ under pk_{DR} . For practical reasons, we use

¹¹ If one does not want to assume the existence of a TTP, one can let distrusting parties perform a secure multi-party computation once to generate the common reference string (CRS). We assume that the CRS (~ 7 KB) is distributed along with the OBU software, certified by a trusted certification authority.

Groth-Sahai NIZKs [42] and the Dodis-Yampolskiy PRF [33]. 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. Instead, we use a CCA-secure structure-preserving encryption scheme for vectors of G_1 -elements [20] and the following workaround: The user splits up its share λ' 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_U are jointly encrypted as e^* under pk_{DR} . The CCA-secure ciphertext e^* unambiguously binds λ to the user's key pk_U and rules out malleability attacks. Otherwise, a malicious TSP could potentially trick the DR into recovering the trapdoor of a different (innocent) user. The user then sends over e^* , $c_{\mathcal{T}}$, $c_{\mathcal{R}}$ along with a NIZK π proving that everything has been generated honestly and the wallet is bound to the user owning sk_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}_U and $c_{\mathcal{R}}$ along with s using $\text{sk}_{\mathcal{T}}^{\mathcal{T}}$.¹² 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 token $\tau := (\mathbf{a}_U, c_{\mathcal{R}}, d_{\mathcal{R}}, \sigma_{\mathcal{R}}, c_{\mathcal{T}}, d_{\mathcal{T}}, \sigma_{\mathcal{T}}, \lambda := \lambda' + \lambda'', b := 0, u_1, x := 1, s)$, 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 $\text{htd} := (\text{pk}_U, s, \lambda'', e^*)$ as hidden user trapdoor to recover λ with help of DR.

Blacklisting and Recalculation: 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 pk_U^{DR} should be blacklisted before the protocol starts. Given e^* and λ'' , the DR recovers the contained PRF key but only if it is bound to pk_U^{DR} . To this end, DR decrypts e^* using sk_{DR} to obtain Λ'_i , Λ'' , and pk_U . If $\text{pk}_U \neq \text{pk}_U^{DR}$ or $\Lambda'' \neq g_1^{\lambda''}$ it aborts. Otherwise, it computes the (small) discrete logarithms of the Λ'_i to the base g_1 to recover the chunks λ'_i of the user's share of the PRF key. The key is computed as $\lambda := \lambda'' + \sum_{i=0}^{\ell-1} \lambda'_i \cdot \mathcal{B}^i$. Using λ , the fraud detection IDs belonging to the previous and upcoming states 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 blacklisted.

Debt Accumulation: In this protocol (Fig. 2) 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 user's main input is the token $\tau^{\text{prev}} := (\mathbf{a}_U, c_{\mathcal{R}}^{\text{prev}}, d_{\mathcal{R}}^{\text{prev}}, \sigma_{\mathcal{R}}^{\text{prev}}, c_{\mathcal{T}}, d_{\mathcal{T}}, \sigma_{\mathcal{T}}, \lambda, b^{\text{prev}}, u_1^{\text{prev}}, x, s^{\text{prev}})$ 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^{next} the same values are contained in the previous wallet state. The randomness u_1^{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 $t := u_2 \cdot \text{sk}_U + u_1^{\text{prev}} \bmod \mathfrak{p}$ (a linear equation in the unknowns u_1^{prev} and sk_U), the fraud detection ID $\phi := \text{PRF}(\lambda, x)$, and a hidden user ID $c_{\text{hid}} := \text{Com}(\text{pk}_U)$. The latter is used in Prove Participation to associate this interaction with the user. As response, the user sends over c_{hid} , $c'_{\mathcal{R}}$, ϕ , t , \mathbf{a}_U , $\mathbf{a}_{\mathcal{R}}^{\text{prev}}$, π , and s' , where π is a NIZK proving that everything has been computed honestly, and s' is the user's share of the serial number. In particular, π 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 , ϕ , and c_{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 ϕ is not on the blacklist and π is correct. Then it calculates the price p , adds it to the user's balance b^{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 RSU sends $c_{\mathcal{R}}$ to the

¹² A uniformly random serial number s for this transaction is jointly generated by user and TSP by means of a Blum-like coin toss (see also Debt Accumulation for details).




Fig. 2. A Simplified Version of Debt Accumulation

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 d''_{ser} for c''_{ser} .

The user checks if the received data is valid and ends up with an updated token $\tau := (\mathbf{a}_{\mathcal{U}}, c_{\mathcal{R}}, d_{\mathcal{R}}, \sigma_{\mathcal{R}}, c_{\mathcal{T}}, d_{\mathcal{T}}, \sigma_{\mathcal{T}}, \lambda, b^{\text{prev}} + p, u_1^{\text{next}}, x + 1, s)$ 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_{hid} allows to open the hidden ID. The RSU stores $(\phi, \mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}}, p, t, u_2, c_{\text{hid}}, s)$.

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 time frame. If the user owns some c_{hid} contained in $\Omega_{\mathcal{R}}^{\text{PP}}$, he simply sends over c_{hid} along with the corresponding opening d_{hid} he stored and the serial number s of the transaction. The SA accepts if the commitment d_{hid} indeed opens to $\text{pk}_{\mathcal{U}}$ and s is part of the corresponding prove participation information the RSU stored. No other user may claim to have sent c_{hid} , as this would imply to open c_{hid} with a different public key.

Debt Clearance: In this protocol conducted by a user and the TSP, the user identifies himself using $\text{sk}_{\mathcal{U}}$ and reveals the balance of his current wallet state. The

wallet is terminated 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 ϕ . This is the same as in Debt Accumulation. Then the user sends over $\text{pk}_{\mathcal{U}}$, b^{prev} , ϕ , t and a NIZK proof π . This NIZK proof asserts that the user knows a certified wallet state with balance b^{prev} , fraud detection ID ϕ , 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.

Double-Spending Detection: This algorithm is applied by the TSP to its database containing all double-spending information (ϕ, t, u_2) collected by the RSUs. The fraud detection ID ϕ is thereby used as the database index. 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 information as follows: Let us assume there exist two records (ϕ, t, u_2) , (ϕ, t', u'_2) . In this case, $\text{sk}_{\mathcal{U}}$ can be recovered as $\text{sk}_{\mathcal{U}} = (t - t')(u_2 - u'_2)^{-1} \bmod \mathfrak{p}$ 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 $\pi := \text{sk}_{\mathcal{U}}$.

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 π , 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 [Appendix D](#) (combined as one comprehensive toll collection protocol π_{P4TC}) UC-realize the ideal model $\mathcal{F}_{\text{P4TC}}$ from [Appendix A.1](#) in the $(\mathcal{F}_{\text{CRS}}, \bar{\mathcal{G}}_{\text{bb}})$ -hybrid model, i.e.,

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

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 a 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. While full versions of all proofs can be found in [Appendix E](#), the following proof sketch explains the ideas behind them.

Proof Outline

For our security statement, we separately prove correctness, system security and user security and privacy.

Although proofs of correctness are often neglected for smaller UC-secure protocols, they are highly non-trivial for extensive ideal models as are required for a complex real-world task like toll collection. Since it is not only instructive to understand the underlying functionality but also a helpful basis for our proofs of system and user security, we briefly sketch the idea behind our proof of correctness here.

First, note that the entries $\text{trdb} = (s^{\text{prev}}, s, \phi, \lambda, \text{pid}_{\mathcal{U}}, \text{pid}_{\mathcal{R}}, p, b)$ of the ideal transaction database TRDB define a graph structure where serial numbers s are considered vertices and predecessors s^{prev} define edges (s^{prev}, s) . Assigning a label $(\text{pid}_{\mathcal{R}}, p)$ to this edge and $(\phi, \lambda, \text{pid}_{\mathcal{U}}, b)$ 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 λ , double-spending corresponds to branching and different wallet states have the same fraud detection ID ϕ 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 updatable 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.

6 Performance Evaluation

We implemented our system for a realistic target platform. Measurements for the user side were done 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 [62]. 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 took the ECONOLITE Connected Vehicle CoProcessor Module as a reference system, which was specifically “designed to enable third-party-developed or processor-intensive applications” [34] and measured on comparable hardware. The TSP’s runtime was measured on a standard laptop featuring an i7-6600U processor, although the TSP is typically equipped with more powerful hardware.

Building Block Instantiation

We implemented P4TC in C++17 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 [7]. We developed our own library for Groth-Sahai NIZK proofs [35, 42] and employed the method in [19] to realize the range proofs required in Wallet Issuing. We make use of two types of commitments: the shrinking commitment scheme from [4], as well as a (dual-mode) extractable commitment scheme from [42]. Moreover, we implemented the structure-preserving signature scheme from [2]. We adapted the structure-preserving IND-CCA-secure encryption scheme from [20] to the asymmetric setting for encrypting PRF key shares in the Wallet Issuing protocol, as well as implemented the IND-CCA-secure encryption scheme from [26] (in combination with AES-CBC and HMAC-SHA256) to establish secure channels. The PRF is instantiated with the Dodis-Yampolskiy construction [33].

Parameter Choice

As bilinear group setting we use the Barreto-Naehrig curves Fp254BNb and Fp254n2BNb [13, 51] together with the optimal Ate pairing, which yields shortest execution times [60]). This results in a security level of about 100 bit [11]. 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 254-bit curves, each vector component can encode 253 bits of information. In practice, it should be possible to encode multiple attributes in one such component.

Implementation Results

Table 2 shows the results of our measurements 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. The row entitled “Session Key Generation” in Table 2 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. 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 bracketed entries in Table 2) are likely an upper bound, since TSP hardware will be more powerful in practice.

Debt Accumulation: While Wallet Issuing and Debt Clearance are *non-time-critical* as they only need to be executed once a month (offroad), 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. During the actual interaction with the RSU (*online* phase), the remaining part of the NIZK is computed and all data is transmitted. Regarding $|\mathbf{a}_U| = |\mathbf{a}_R| = 1$, 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-

Table 2. P4TC performance results. Runtime t is averaged over 1 000 executions. Transmitted data n is rounded up to full bytes.

Protocol	$ \mathbf{a}_{\mathcal{U}} = \mathbf{a}_{\mathcal{R}} = 1$				$ \mathbf{a}_{\mathcal{U}} = \mathbf{a}_{\mathcal{R}} = 4$			
	t_{user} [ms]	$t_{\text{RSU/TSP}}$ [ms]	n_{user} [byte]	$n_{\text{RSU/TSP}}$ [byte]	t_{user} [ms]	$t_{\text{RSU/TSP}}$ [ms]	n_{user} [byte]	$n_{\text{RSU/TSP}}$ [byte]
Session Key Generation	15	≈ 0	131	—	15	≈ 0	131	—
Wallet Issuing	27 064	(8 490)	87 951	944	27 183	(8 545)	88 107	1 152
Debt Accumulation								
– Precomp (offline)	2 715	—	—	—	2 716	—	—	—
– Online	348	475	8 128	976	456	5 256	8 336	1 088
– Online (cached certificate)	41	475	8 128	976	40	526	8 336	1 088
– Postprocessing (offline)	34	—	—	—	34	—	—	—
Debt Clearance	2 435	(745)	7 071	96	2 436	(767)	7 280	96

Sahai verification by Herold et al. [45]. 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 OBU collect data in order to, e.g., prevent double-spending and 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 the OBU performs 10000 transactions in one billing period, 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 stores 246 bytes of transaction information for each run of Debt Accumulation

(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 [43], 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 ID requires 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 , requiring 35MB of memory (neglecting overhead due to the data structure). Hence, blacklisting does not affect the execution time of the RSU.

Computing DLOGs: To blacklist a user, the DR has to compute a number of discrete logarithms to recover λ . With our choice of parameters, λ 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. [16] can be used 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.

Acknowledgements

This research received funding from the German Research Foundation (DFG) within grants RU 1664/3-1 and PA 587/10-1 and as part of the Research Training Group GRK 2153.

References

- [1] ABC News. Toll records catch unfaithful spouses, 2008. URL <https://abcnews.go.com/Technology/story?id=3468712&page=1>.
- [2] M. Abe, J. Groth, K. Haralambiev, and M. Ohkubo. Optimal structure-preserving signatures in asymmetric bilinear groups. In P. Rogaway, editor, *CRYPTO 2011*, volume 6841 of *LNCS*, pages 649–666. Springer, Heidelberg, Aug. 2011.
- [3] M. Abe, J. Groth, M. Ohkubo, and T. Tango. Converting cryptographic schemes from symmetric to asymmetric bilinear groups. In J. A. Garay and R. Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 241–260. Springer, Heidelberg, Aug. 2014. [10.1007/978-3-662-44371-2_14](https://doi.org/10.1007/978-3-662-44371-2_14).
- [4] M. Abe, M. Kohlweiss, M. Ohkubo, and M. Tibouchi. Fully structure-preserving signatures and shrinking commitments. In E. Oswald and M. Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 35–65. Springer, Heidelberg, Apr. 2015. [10.1007/978-3-662-46803-6_2](https://doi.org/10.1007/978-3-662-46803-6_2).
- [5] C. Albrecht, F. Gurski, J. Rethmann, and E. Yilmaz. Knapsack problems: A parameterized point of view. *Theoretical Computer Science*, 2018.
- [6] American Civil Liberties Union. Newly obtained records reveal extensive monitoring of e-zpass tags throughout new york, 2015. URL <https://www.aclu.org/blog/privacy-technology/location-tracking/newly-obtained-records-reveal-extensive-monitoring-e-zpass>.
- [7] D. F. Aranha and C. P. L. Gouvêa. Relic is an efficient library for cryptography, 2016. URL <https://github.com/relic-toolkit/relic>.
- [8] J. Balasch, A. Rial, C. Troncoso, B. Preneel, I. Verbauwhede, and C. Gueuns. PrETP: Privacy-preserving electronic toll pricing. In *USENIX Security 2010*, pages 63–78. USENIX Association, Aug. 2010.
- [9] F. Baldimtsi, M. Chase, G. Fuchsbauer, and M. Kohlweiss. Anonymous transferable E-cash. In J. Katz, editor, *PKC 2015*, volume 9020 of *LNCS*, pages 101–124. Springer, Heidelberg, Mar. / Apr. 2015. [10.1007/978-3-662-46447-2_5](https://doi.org/10.1007/978-3-662-46447-2_5).
- [10] B. Barak, R. Canetti, J. B. Nielsen, and R. Pass. Universally composable protocols with relaxed set-up assumptions. In *45th FOCS*, pages 186–195. IEEE Computer Society Press, Oct. 2004.
- [11] R. Barbulescu and S. Duquesne. Updating key size estimations for pairings. Cryptology ePrint Archive, Report 2017/334, 2017. <https://eprint.iacr.org/2017/334>.
- [12] A. Barki, S. Brunet, N. Desmoulins, S. Gams, S. Gharout, and J. Traoré. Private eCash in practice (short paper). In J. Grossklags and B. Preneel, editors, *FC 2016*, volume 9603 of *LNCS*, pages 99–109. Springer, Heidelberg, Feb. 2016.
- [13] P. S. L. M. Barreto and M. Naehrig. Pairing-friendly elliptic curves of prime order. In B. Preneel and S. Tavares, editors, *SAC 2005*, volume 3897 of *LNCS*, pages 319–331. Springer, Heidelberg, Aug. 2006.
- [14] M. Bellare. New proofs for NMAC and HMAC: Security without collision resistance. *Journal of Cryptology*, 28(4): 844–878, Oct. 2015. [10.1007/s00145-014-9185-x](https://doi.org/10.1007/s00145-014-9185-x).
- [15] M. Bellare, A. Desai, D. Pointcheval, and P. Rogaway. Relations among notions of security for public-key encryption schemes. In H. Krawczyk, editor, *CRYPTO'98*, volume 1462 of *LNCS*, pages 26–45. Springer, Heidelberg, Aug. 1998.
- [16] D. J. Bernstein and T. Lange. Computing small discrete logarithms faster. Cryptology ePrint Archive, Report 2012/458, 2012. <https://eprint.iacr.org/2012/458>.
- [17] F. Bourse, D. Pointcheval, and O. Sanders. Divisible e-cash from constrained pseudo-random functions. In S. D. Galbraith and S. Moriai, editors, *Advances in Cryptology - ASIACRYPT 2019 - 25th International Conference on the Theory and Application of Cryptology and Information Security, Kobe, Japan, December 8-12, 2019, Proceedings, Part I*, volume 11921 of *Lecture Notes in Computer Science*, pages 679–708. Springer, 2019. [10.1007/978-3-030-34578-5_24](https://doi.org/10.1007/978-3-030-34578-5_24).
- [18] J. Camenisch, S. Hohenberger, and A. Lysyanskaya. Compact e-cash. In R. Cramer, editor, *EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 302–321. Springer, Heidelberg, May 2005.
- [19] J. Camenisch, R. Chaabouni, and a. shelat. Efficient protocols for set membership and range proofs. In J. Pieprzyk, editor, *ASIACRYPT 2008*, volume 5350 of *LNCS*, pages 234–252. Springer, Heidelberg, Dec. 2008.
- [20] J. Camenisch, K. Haralambiev, M. Kohlweiss, J. Lapon, and V. Naessens. Structure preserving CCA secure encryption and applications. In D. H. Lee and X. Wang, editors, *ASIACRYPT 2011*, volume 7073 of *LNCS*, pages 89–106. Springer, Heidelberg, Dec. 2011.
- [21] S. Canard and A. Gouget. Anonymity in transferable e-cash. In S. M. Bellovin, R. Gennaro, A. D. Keromytis, and M. Yung, editors, *ACNS 08*, volume 5037 of *LNCS*, pages 207–223. Springer, Heidelberg, June 2008.
- [22] R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd FOCS*, pages 136–145. IEEE Computer Society Press, Oct. 2001.
- [23] R. Canetti. Obtaining universally composable security: Towards the bare bones of trust. Cryptology ePrint Archive, Report 2007/475, 2007. <https://eprint.iacr.org/2007/475>.
- [24] R. Canetti, Y. Dodis, R. Pass, and S. Walfish. Universally composable security with global setup. In S. P. Vadhan, editor, *TCC 2007*, volume 4392 of *LNCS*, pages 61–85. Springer, Heidelberg, Feb. 2007.
- [25] R. Canetti, D. Shahaf, and M. Vald. Universally composable authentication and key-exchange with global PKI. In C.-M. Cheng, K.-M. Chung, G. Persiano, and B.-Y. Yang, editors, *PKC 2016, Part II*, volume 9615 of *LNCS*, pages 265–296. Springer, Heidelberg, Mar. 2016. [10.1007/978-3-662-49387-8_11](https://doi.org/10.1007/978-3-662-49387-8_11).
- [26] D. Cash, E. Kiltz, and V. Shoup. The twin Diffie-Hellman problem and applications. In N. P. Smart, editor, *EUROCRYPT 2008*, volume 4965 of *LNCS*, pages 127–145. Springer, Heidelberg, Apr. 2008.
- [27] X. Chen, G. Lenzini, S. Mauw, and J. Pang. A group signature based electronic toll pricing system. In *2012 Seventh International Conference on Availability, Reliability and Security, ARES 2012*, pages 85–93, 2012.
- [28] X. Chen, G. Lenzini, S. Mauw, and J. Pang. Design and formal analysis of A group signature based electronic toll pricing system. *JoWUA*, 4(1):55–75, 2013. URL <http://isyou.info/jowua/papers/jowua-v4n1-3.pdf>.

- [29] M. Dahl, S. Delaune, and G. Steel. Formal analysis of privacy for anonymous location based services. *Theory of Security and Applications*, pages 98–112, 2012.
- [30] Datatilsynet. Statens vegvesen holdt tilbake viktig autopass-informasjon (press release), 2007. URL <http://www.datatilsynet.no/>.
- [31] J. Day, Y. Huang, E. Knapp, and I. Goldberg. Spectre: Spot-checked private ecash tolling at roadside. In *Proceedings of the 10th Annual ACM Workshop on Privacy in the Electronic Society*, WPES '11, pages 61–68, 2011.
- [32] Deutsches Bundesamt für Güterverkehr. Monatliche mautstatistik für januar 2018, 2018. URL https://www.bag.bund.de/SharedDocs/Downloads/DE/Statistik/Lkw-Maut/18_Monatstab_01.html.
- [33] Y. Dodis and A. Yampolskiy. A verifiable random function with short proofs and keys. *Cryptology ePrint Archive*, Report 2004/310, 2004. <https://eprint.iacr.org/2004/310>.
- [34] ECONOLITE Group. Connected vehicle coprocessor module, 2018. URL <http://www.econolitegroup.com/wp-content/uploads/2017/05/controllers-connectedvehicle-datasheet.pdf>.
- [35] A. Escala and J. Groth. Fine-tuning Groth-Sahai proofs. In H. Krawczyk, editor, *PKC 2014*, volume 8383 of *LNCS*, pages 630–649. Springer, Heidelberg, Mar. 2014. [10.1007/978-3-642-54631-0_36](https://doi.org/10.1007/978-3-642-54631-0_36).
- [36] European Commission. Study on state of the art of electronic road tolling, 2015. URL https://ec.europa.eu/transport/sites/transport/files/modes/road/road_charging/doc/study-electronic-road-tolling.pdf.
- [37] European Commission. 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), 2017. URL <https://ec.europa.eu/transport/sites/transport/files/com20170280-eets-directive.pdf>.
- [38] European Commission. The eu general data protection regulation (gdpr), 2018. URL <https://www.eugdpr.org/>.
- [39] European Parliament. Technology options for the european electronic toll service, 2014. URL [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).
- [40] F. D. Garcia, E. R. Verheul, and B. Jacobs. Cell-based roadpricing. In *European Public Key Infrastructure Workshop – EuroPKI 2011*, volume 7163 of *Lecture Notes in Computer Science*, pages 106–122, 2011.
- [41] Global Markets Insights. ETC Market Report 2019, 2019. URL <https://www.marketwatch.com/press-release/electronic-toll-collection-market-2019-in-depth-industry-analysis-by-product-technology-application-opportunities-and-growth-forecast-by-2025-2019-11-12>.
- [42] J. Groth and A. Sahai. Efficient non-interactive proof systems for bilinear groups. In N. P. Smart, editor, *EUROCRYPT 2008*, volume 4965 of *LNCS*, pages 415–432. Springer, Heidelberg, Apr. 2008.
- [43] E.-Z. Group. E-ZPass Statistics: 2005 - 2016, 2017. URL <https://e-zpassag.com/about-us/statistics>.
- [44] G. Hartung, M. Hoffmann, M. Nagel, and A. Rupp. BBA+: Improving the security and applicability of privacy-preserving point collection. In B. M. Thuraisingham, D. Evans, T. Malkin, and D. Xu, editors, *ACM CCS 2017*, pages 1925–1942. ACM Press, Oct. / Nov. 2017.
- [45] G. Herold, M. Hoffmann, M. Kloß, C. Ràfols, and A. Rupp. New techniques for structural batch verification in bilinear groups with applications to groth-sahai proofs. In B. M. Thuraisingham, D. Evans, T. Malkin, and D. Xu, editors, *ACM CCS 2017*, pages 1547–1564. ACM Press, Oct. / Nov. 2017.
- [46] M. Hoffmann, M. Kloß, M. Raiber, and A. Rupp. Black-box wallets: Fast anonymous two-way payments for constrained devices. *Proceedings on Privacy Enhancing Technologies*, 2020(1):165–194, 2020. [10.2478/popets-2020-0010](https://doi.org/10.2478/popets-2020-0010).
- [47] R. Jardí-Cedó, J. Castellà-Roca, and A. Viejo. 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*, volume 8872 of *Lecture Notes in Computer Science*, pages 327–334, 2014.
- [48] R. Jardí-Cedó, M. Mut-Puigserver, M. M. Payeras-Capellà, J. Castellà-Roca, and A. Viejo. Electronic road pricing system for low emission zones to preserve driver privacy. In *Modeling Decisions for Artificial Intelligence – MDAI 2014. Proceedings*, pages 1–13, 2014.
- [49] R. Jardí-Cedó, M. Mut-Puigserver, M. M. Payeras-Capellà, J. Castellà-Roca, and A. Viejo. 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, pages 158–165, 2016.
- [50] Kapsch (Toll Collection System Integrator and Supplier). Personal Communication, 2018.
- [51] Y. Kawahara, T. Kobayashi, M. Scott, and A. Kato. Barreto-naehrig curves. Internet draft, Internet Engineering Task Force, Mar. 2016. Work in Progress.
- [52] F. Kerschbaum and H. W. Lim. Privacy-preserving observation in public spaces. In G. Pernul, P. Y. A. Ryan, and E. R. Weippl, editors, *ESORICS 2015, Part II*, volume 9327 of *LNCS*, pages 81–100. Springer, Heidelberg, Sept. 2015. [10.1007/978-3-319-24177-7_5](https://doi.org/10.1007/978-3-319-24177-7_5).
- [53] E. Kiltz, J. Pan, and H. Wee. Structure-preserving signatures from standard assumptions, revisited. In R. Gennaro and M. J. B. Robshaw, editors, *CRYPTO 2015, Part II*, volume 9216 of *LNCS*, pages 275–295. Springer, Heidelberg, Aug. 2015. [10.1007/978-3-662-48000-7_14](https://doi.org/10.1007/978-3-662-48000-7_14).
- [54] H. Koelmeyer and S. Kandeepan. Tagless tolling using dsrc for intelligent transport system: An interference study. In *Asia Modelling Symposium*. IEEE, 2017.
- [55] S. Kummer, M. Dieplinger, and M. Dobrovnik. Endbericht der studie “flächendeckende schwerverkehrs-maut in Österreich”, 2015. URL https://www.wko.at/branchen/stmk/transport-verkehr/gueterbefoerderungsgewerbe/Endbericht_FlaechendeckendeMaut_final.pdf.
- [56] J. Lapon, M. Kohlweiss, B. D. Decker, and V. Naessens. 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*, pages 289–301, 2010. [10.1007/978-3-642-15257-3_26](https://doi.org/10.1007/978-3-642-15257-3_26).

- [57] Y. J. Li. An overview of the dsrc/wave technology. In *International Conference on Heterogeneous Networking for Quality, Reliability, Security and Robustness*, pages 544–558. Springer, 2010.
- [58] Markets and Markets. Electronic Toll Collection Market Study, 2017. URL <https://www.marketsandmarkets.com/Market-Reports/electronic-toll-collection-system-market-224492059.html>.
- [59] S. Meiklejohn, K. Mowery, S. Checkoway, and H. Shacham. The phantom tollbooth: Privacy-preserving electronic toll collection in the presence of driver collusion. In *USENIX Security 2011*. USENIX Association, Aug. 2011.
- [60] D. Moody, R. C. Peralta, R. A. Perlner, A. R. Regenscheid, A. L. Roginsky, and L. Chen. Report on pairing-based cryptography. In *Journal of Research of the National Institute of Standards and Technology*, volume 120, pages 11–27, Gaithersburg, MD, USA, Feb. 2015. National Institute of Standards and Technology.
- [61] R. A. Popa, H. Balakrishnan, and A. J. Blumberg. VPriv: Protecting privacy in location-based vehicular services. In F. Monrose, editor, *USENIX Security 2009*, pages 335–350. USENIX Association, Aug. 2009.
- [62] Savari.net. Mobiwave on-board-unit (obu), 2017. URL http://savari.net/wp-content/uploads/2017/05/MW-1000_April2017.pdf.
- [63] L. Sweeney. k-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 10(5):557–570, 2002. [10.1142/S0218488502001648](https://doi.org/10.1142/S0218488502001648).

A Full System Definition

In this appendix we give a detailed description and explanation of our ideal privacy-preserving electronic toll collection functionality \mathcal{F}_{P4TC} . We first mention a few preliminary remarks, then we describe the different tasks in a high-level fashion before presenting the full-fledged ideal functionality. Last, we discuss why this ideal functionality precisely models a secure and privacy-preserving electronic toll collection scheme.

Setup Assumptions and Implicit Writing Conventions

As commonly found in the UC setting, we also draw from setup assumptions. Setup assumptions are ideal functionalities that still remain ideal in the real experiment, e.g., their secure realization is left unspecified. Especially, the security of their realization has to be either proven outside of the (UC-)model or simply assumed. In our scenario, we assume a common reference string (CRS), denoted \mathcal{F}_{CRS} , and a globally available bulletin board, denoted $\bar{\mathcal{G}}_{bb}$ [25, Fig. 3], which is sometimes also referred to as a key registration service in the literature [10, 23].

A CRS is a short piece of information that is shared between all parties. The trust assumptions are that the CRS has been generated honestly and that all parties possess the same CRS.

A bulletin board can be depicted as a key registration service which associates (physical) party identifiers (PIDs) with (cryptographic) public keys. The assumptions about $\bar{\mathcal{G}}_{bb}$ are that upon registration the operator of the bulletin board checks the identity of the registering party in a trustworthy way and that every party can retrieve information from $\bar{\mathcal{G}}_{bb}$ trustworthily. As the PID establishes an entity’s physical identity it cannot be captured by cryptographic means. 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.

Lastly, we assume our functionality also uses the implicit writing conventions for ideal functionalities [22]. 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.

A.1 Tasks provided by \mathcal{F}_{P4TC}

Before describing the ideal functionality in full detail, we first give a condensed description. We explain Debt Accumulation in some detail and give only short sketches of the remaining tasks.

Individual tasks (e.g., Wallet Issuing, Debt Accumulation etc.) are not formalized as separate functionalities. Instead the whole system is given as one monolithic, highly reactive ideal functionality \mathcal{F}_{P4TC} with polynomially many parties as users and RSUs. This allows for a shared state between individual interactions. An excerpt of \mathcal{F}_{P4TC} is depicted in Fig. 3. For the ease of presentation, all instructions which are typically executed are printed in normal font, while some conditional side tracks are grayed out. The conditional branches deal with corrupted, misbehaving,¹³ or blacklisted parties.

Task Debt Accumulation

In this task, the user provides a serial number s^{prev} as input to \mathcal{F}_{P4TC} , indicating which past wallet state he wishes to use for this transaction. Of course, an honest user always uses the wallet state resulting from the pre-

¹³ Please note, that users do not need to be formally corrupted in order to commit double-spending. We call these users honest, but misbehaving.

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

I. State

The following information is stored by $\mathcal{F}_{\text{P4TC}}$:

- Set $TRDB$ of transaction entries $trdb$. Each entry has the form $trdb = (s^{\text{prev}}, s, \phi, x, \lambda, pid_{\mathcal{U}}, pid_{\mathcal{R}}, p, b)$.
- Mapping $f_{\Phi} : \mathcal{L} \times \mathbb{N}_0 \rightarrow \Phi, (\lambda, x) \mapsto \phi$ assigning a fraud detection ID ϕ to a wallet ID λ and counter x .
- Mapping $f_{\mathcal{A}_{\mathcal{U}}} : \mathcal{L} \rightarrow \mathcal{A}_{\mathcal{U}}, \lambda \mapsto \mathbf{a}_{\mathcal{U}}$ assigning user attributes to a given wallet ID λ .
- Mapping $f_{\mathcal{A}_{\mathcal{R}}} : \mathcal{PID}_{\mathcal{R}} \rightarrow \mathcal{A}_{\mathcal{R}}, pid_{\mathcal{R}} \mapsto \mathbf{a}_{\mathcal{R}}$ assigning RSU attributes to a given RSU PID $pid_{\mathcal{R}}$.

I. Behavior—Task Debt Accumulation

User input $(\text{pay_toll}, s^{\text{prev}})$

RSU input $(\text{pay_toll}, bl_{\mathcal{R}})$

- (1) Pick serial number $s \xleftarrow{\mathbb{R}} S$ that has not previously been used.
- (2) *User corrupted*: Ask the adversary if the PID $pid_{\mathcal{U}}$ of another corrupted user should be used.^a
- (3) Look up $(\cdot, s^{\text{prev}}, \phi^{\text{prev}}, x^{\text{prev}}, \lambda^{\text{prev}}, pid_{\mathcal{U}}, pid_{\mathcal{R}}^{\text{prev}}, \cdot, b^{\text{prev}}) \in TRDB$, with $(s^{\text{prev}}, pid_{\mathcal{U}})$ being the unique key.
- (4) Adopt previous wallet ID $\lambda := \lambda^{\text{prev}}$ and increase counter $x := x^{\text{prev}} + 1$.
- (5) *Double-spending/Blacklisted*: In this case $\phi := f_{\Phi}(\lambda, x)$ has already been defined; continue with (6).
Pick fraud detection ID $\phi \xleftarrow{\mathbb{R}} \Phi$ that has not previously been used.
User corrupted: Allow the adversary to choose a previously unused fraud detection ID ϕ .^b
Append assignment $(\lambda, x) \mapsto \phi$ to f_{Φ} .
- (6) If $\phi \in bl_{\mathcal{R}}$, output blacklisted to both parties and abort.
- (7) Look up attributes $(\mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}}) := (f_{\mathcal{A}_{\mathcal{U}}}(\lambda), f_{\mathcal{A}_{\mathcal{R}}}(pid_{\mathcal{R}}), f_{\mathcal{A}_{\mathcal{R}}}(pid_{\mathcal{R}}^{\text{prev}}))$.
- (8) Calculate price $p := \mathcal{O}_{\text{pricing}}(\mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}})$.
RSU corrupted: Leak $(\mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}})$ to the adversary and obtain a price p .
- (9) Calculate new balance $b := b^{\text{prev}} + p$.
- (10) Append new transaction $(s^{\text{prev}}, s, \phi, x, \lambda, pid_{\mathcal{U}}, pid_{\mathcal{R}}, p, b)$ to $TRDB$.

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

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

^a If corrupted users collude, they might share their credentials and use each other's wallet.

^b The ideal model only guarantees privacy for honest users. For corrupted users the fraud detection ID might be chosen adversarially.

Fig. 3. An excerpt of the ideal functionality $\mathcal{F}_{\text{P4TC}}$.

vious transaction. The participating RSU inputs a list of fraud detection IDs that are blacklisted. Firstly, \mathcal{F}_{P4TC} randomly picks a fresh serial number s for the upcoming transaction. If the user is corrupted, \mathcal{F}_{P4TC} allows the simulator to provide a different value for $pid_{\mathcal{U}}$ that belongs to another corrupted user.¹⁴ \mathcal{F}_{P4TC} looks up the previous wallet state $trdb^{\text{prev}}$ in $TRDB$. The ideal functionality extracts the wallet ID from the previous record ($\lambda := \lambda^{\text{prev}}$) and increases the counter for this particular wallet ($x := x^{\text{prev}} + 1$). Then, it checks if a fraud detection ID ϕ has already been defined for the wallet ID λ and counter x (n.b.: tree and depth of node). If so, the current transaction record will be assigned the same fraud detection ID ϕ . 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 λ . If this fraud detection ID is blacklisted, the task aborts.¹⁵ If and only if the user is corrupted and ϕ has not previously been defined, the adversary is allowed to overwrite the fraud detection ID with another value.¹⁶ Moreover, \mathcal{F}_{P4TC} looks up the user's attributes bound to this particular wallet ($\mathbf{a}_{\mathcal{U}} := f_{\mathcal{A}_{\mathcal{U}}}(\lambda)$) and the attributes of the current and previous RSU ($\mathbf{a}_{\mathcal{R}} := f_{\mathcal{A}_{\mathcal{R}}}(pid_{\mathcal{R}})$, $\mathbf{a}_{\mathcal{R}}^{\text{prev}} := f_{\mathcal{A}_{\mathcal{R}}}(pid_{\mathcal{R}}^{\text{prev}})$). Finally, the ideal functionality queries the pricing oracle O_{pricing} for the price p of this transaction, calculates the new balance of the wallet ($b := b^{\text{prev}} + p$) 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 dis-

pute) files an out-of-band claim later. The user's output are the serial number s of the current transaction, the current RSU's attributes $\mathbf{a}_{\mathcal{R}}$, 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 ϕ and the attributes $\mathbf{a}_{\mathcal{U}}$ and $\mathbf{a}_{\mathcal{R}}^{\text{prev}}$ of the user and the previous RSU, respectively. Learning their mutual attributes is necessary, because RSU and user must evaluate the pricing function themselves in the real protocol without the help of a third party.

Remaining Tasks

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

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

The task *Debt Clearance* is very similar to Debt Accumulation described above except that the TSP additionally learns the user's party ID $pid_{\mathcal{U}}$ 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 transaction entry becomes a leaf node of the wallet tree.

The task *Blacklisting and Recalculation* 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 past and upcoming serial numbers. For the latter, \mathcal{F}_{P4TC} draws a sequence of fresh, uniform and independent fraud detection IDs ϕ_i and pre-fills the mapping f_{ϕ} for all wallets the user owns. This ensures that upcoming transactions use predetermined fraud detection IDs that are actually blacklisted.

The task *Prove Participation* checks if a TRDB record exists for the particular user and serial number.

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}_{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}_{P4TC} outputs both—the user's identity and the proof—to the TSP and records the proof as valid for the user.

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

¹⁴ This is a required technicality as corrupted users might share their credentials and thus might use each other's wallet. Please note that this does not affect honest users.

¹⁵ Note, that the probability to blacklist a freshly drawn fraud detection ID is negligible. Only if $f_{\phi}(\lambda, x)$ has already been defined by a past task, this yields a chance to successfully blacklist a user.

¹⁶ 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.

A.2 Full Ideal Functionality

As explained before we define $\mathcal{F}_{\text{P4TC}}$ 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 task¹⁷ 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

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

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 λ 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^{prev} to the previous transaction $trdb^{\text{prev}}$ (which corresponds to the wallet state before $trdb$). Lastly, a fraud detection ID ϕ 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 λ if and only if no double-spending has been committed with this wallet. The fraud detection ID is constant for each pair (λ, 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 map-

¹⁷ 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_{corrupt}	set of corrupted party identifiers
f_{Φ}	(partial) mapping assigning a fraud detection ID ϕ to given wallet ID λ and counter x
$f_{\mathcal{A}_{\mathcal{U}}}$	(partial) mapping assigning user attributes $\mathbf{a}_{\mathcal{U}}$ to a given wallet ID λ
$f_{\mathcal{A}_{\mathcal{R}}}$	(partial) mapping assigning RSU attributes $\mathbf{a}_{\mathcal{R}}$ to a given RSU PID $pid_{\mathcal{R}}$

ping to fraud detection IDs contains most of the information our toll collection scheme needs, $\mathcal{F}_{\text{P4TC}}$ stores three more partially defined mappings: $f_{\mathcal{A}_{\mathcal{U}}} : \mathcal{L} \rightarrow \mathcal{A}_{\mathcal{U}}$ and $f_{\mathcal{A}_{\mathcal{R}}} : PID_{\mathcal{R}} \rightarrow \mathcal{A}_{\mathcal{R}}$ of user and RSU attribute vectors as well as f_{Π} 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.

A.2.1 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.

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 provide a public

Functionality \mathcal{F}_{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_{Φ} giving the fraud detection ID ϕ corresponding to given wallet ID λ 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 λ :

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

- 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_{Π} assigning a user PID $pid_{\mathcal{U}}$ and a proof of guilt π to a validity bit:

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

II. Behavior

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

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

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}).^a
- (3) Call $\overline{\mathcal{G}}_{bb}$ with input (register, pk_{DR}).

DR output: (pk_{DR})

^a 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)

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

Identifier	Abstract	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
$\mathbf{a}_{\mathcal{U}}$	$\mathcal{A}_{\mathcal{U}}$	G_2^j	user attributes
$\mathbf{a}_{\mathcal{R}}$	$\mathcal{A}_{\mathcal{R}}$	G_1^y	RSU attributes
$\mathbf{a}_{\mathcal{T}}$	$\mathcal{A}_{\mathcal{T}}$	G_1^y	TSP attributes
b	\mathbb{Z}_p	\mathbb{Z}_p	balance
p	\mathbb{Z}	\mathbb{Z}	price to pay at an RSU
ϕ	Φ	G_1	fraud detection ID
s	S	G_1	serial number
λ	\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 Φ 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 $\mathcal{F}_{\text{P4TC}}$ (cont.) – Task *TSP Registration*

TSP input: (register, $\mathbf{a}_{\mathcal{T}}$)

- (1) If this task has been run before, output \perp and abort.
- (2) $TRDB := \emptyset$
- (3) Send (registering_tsp, $pid_{\mathcal{T}}$, $\mathbf{a}_{\mathcal{T}}$) to the adversary and obtain the key ($pk_{\mathcal{T}}$).^a
- (4) Call $\bar{\mathcal{G}}_{bb}$ with input (register, $pk_{\mathcal{T}}$).

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

^a 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 $\mathcal{F}_{\text{P4TC}}$ (cont. from Fig. 4)

Functionality $\mathcal{F}_{\text{P4TC}}$ (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 $(pk_{\mathcal{R}})$.^a
- (3) Call $\bar{\mathcal{G}}_{\text{bb}}$ with input (register, $pk_{\mathcal{R}}$).

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

^a 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 $\mathcal{F}_{\text{P4TC}}$ (cont. from Fig. 4)

Functionality $\mathcal{F}_{\text{P4TC}}$ (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 $(pk_{\mathcal{U}})$.^a
- (3) Call $\bar{\mathcal{G}}_{\text{bb}}$ with input (register, $pk_{\mathcal{U}}$).

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

^a 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 $\mathcal{F}_{\text{P4TC}}$ (cont. from Fig. 4)

key pk for the respective party which is then registered with the bulletin board $\bar{\mathcal{G}}_{bb}$ 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 $\mathbf{a}_{\mathcal{T}}$ as input, which—after a check if this task has been run before—is leaked to the adversary together with $pid_{\mathcal{T}}$ when the public key $pk_{\mathcal{T}}$ is obtained. In addition, all data structures of \mathcal{F}_{P4TC} are initialized as empty sets and empty (partial) mappings respectively. Again, the public key $pk_{\mathcal{T}}$ is output to the TSP.

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 $\mathbf{a}_{\mathcal{R}}$.¹⁸ The content of attribute vectors is in no way restricted by \mathcal{F}_{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}_{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.

A.2.2 Basic Tasks

Now 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.

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}_{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 $\bar{\mathcal{G}}_{bb}$ and checked against the TSP’s blacklist $bl_{\mathcal{T}}$. If this does not lead to an abort, a new wallet ID λ and fraud detection ID ϕ are uniquely and randomly picked, unless the user is corrupted in which case the adversary chooses ϕ . 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 the attribute vector $\mathbf{a}_{\mathcal{U}}$ to check this has been assigned correctly and more importantly does not contain any identifying information.

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^{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}_{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 transaction. \mathcal{F}_{P4TC} looks up if a wallet state $trdb^{\text{prev}}$ in $TRDB$ corresponds to the user input s^{prev} and belongs to the users $pid_{\mathcal{U}}$. 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. The ideal functionality uses part of

¹⁸ 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}_{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}_{P4TC} (cont. from Fig. 4)

Functionality \mathcal{F}_{P4TC} (cont.) – Task *Wallet Issuing*

User input: (issue)

TSP input: (issue, $\mathbf{a}_{\mathcal{U}}$, $bl_{\mathcal{T}}$)

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

User output: ($s, \mathbf{a}_{\mathcal{U}}$)

TSP output: (s)

[†]If this does not exist, output \perp and abort.

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

^b 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}_{P4TC} (cont. from Fig. 4)

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

User input: $(\text{pay_toll}, s^{\text{prev}})$

RSU input: $(\text{pay_toll}, bl_{\mathcal{R}})$

- (1) Pick serial number $s \xleftarrow{\mathbb{R}} S$ 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.^a
- (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)[⊥].
- (4) Set $\lambda := \lambda^{\text{prev}}$ and $x := x^{\text{prev}} + 1$.
- (5) 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 $\phi \xleftarrow{\mathbb{R}} \Phi$ that has not previously been used, otherwise ask the adversary for a fraud detection ID ϕ that has not previously been used.^b Append $f_{\Phi}(\lambda, x) := \phi$ to f_{Φ} .
- (6) If $\phi \in bl_{\mathcal{R}}$, output blacklisted to both parties and abort.
- (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 $\mathbf{a}_{\mathcal{R}}^{\text{prev}} := f_{\mathcal{A}_{\mathcal{R}}}(pid_{\mathcal{R}}^{\text{prev}})$.[⊥]
- (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: $(s, \phi, \mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}})$

[⊥]If this does not exist, output \perp and abort.

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

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

the information from the previous wallet state

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

to determine the content of the new transaction entry $trdb$. The user ID $pid_{\mathcal{U}}$ and wallet ID λ stay the same, $pid_{\mathcal{R}}$ is set to the identity of the participating RSU, and the counter x^{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 (λ, 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 ϕ 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 fraud detection ID ϕ is checked against $bl_{\mathcal{R}}$.¹⁹ The attributes $\mathbf{a}_{\mathcal{U}}$, $\mathbf{a}_{\mathcal{R}}$ and $\mathbf{a}_{\mathcal{R}}^{\text{prev}}$ 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 $\mathcal{F}_{\text{P4TC}}$ 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 $\mathcal{F}_{\text{P4TC}}$ 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.

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

¹⁹ Note, that the probability to blacklist a freshly drawn fraud detection ID is negligible. Only if $f_{\Phi}(\lambda, x)$ has already been defined by a past task, this yields a chance to successfully blacklist a user.

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^{bill} which in turn is taken to be the previous balance b^{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^{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 ϕ to enable double-spending detection.

A.2.3 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.

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 $S_{\mathcal{R}}^{\text{DP}}$ 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 $S_{\mathcal{R}}^{\text{DP}}$. Hence, the ideal

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

User input: (clear_debt, s^{prev})

TSP input: (clear_debt)

- (1) Pick serial number $s \xleftarrow{\text{R}} S$ 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.^a
- (3) Receive $pk_{\mathcal{U}}$ from the bulletin-board $\overline{\mathcal{G}}_{\text{bb}}$ for PID $pid_{\mathcal{U}}$.[⊥]
- (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).[⊥]
- (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 $\phi \xleftarrow{\text{R}} \Phi$ that has not previously been used, otherwise ask the adversary for a fraud detection ID ϕ that has not previously been used.^b Append $f_{\Phi}(\lambda, x) := \phi$ to f_{Φ} .
- (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^{bill})

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

[⊥]If this does not exist, output \perp and abort.

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

^b 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)

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

User input: (prove_participation)

SA input: (prove_participation, $pk_{\mathcal{U}}, S_{\mathcal{R}}^{\text{PP}}$)

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

User output: ($out_{\mathcal{U}}$)

SA output: (out_{SA})

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

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

functionality ensures that either the SA receives the correct answer—even for corrupted users—or aborts.

Some remarks are in order to the abort condition—which is only a technical concession to what can be practically realized. If $pid'_{\mathcal{U}} = pid_{\mathcal{U}}$ holds (i.e., the SA communicates with the expected user), everything is fine. If $pid'_{\mathcal{U}} \neq pid_{\mathcal{U}}$ 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 $pid'_{\mathcal{U}} \neq pid_{\mathcal{U}}$ 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 $pid'_{\mathcal{U}}$). This models the limitation that two corrupted users can share their credentials. A corrupted user (with $pid'_{\mathcal{U}}$) can pass its transaction information to another corrupted user (with $pid_{\mathcal{U}}$) 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 bisectional 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.

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 ϕ as input and checks the transaction database $TRDB$ for two distinct entries containing this same fraud detection ID. In case such entries are present the adversary

is asked for a proof π to be issued for this instance of double-spending. The user ID and proof $(pid_{\mathcal{U}}, \pi)$ are appended to f_{Π} 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 $pid_{\mathcal{U}}$ and a double-spending proof π as input. First, it checks if this particular pair $(pid_{\mathcal{U}}, \pi)$ has already been defined and 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.

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^{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^{bill} . Note that although this sum may contain the prices of transactions 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.

Functionality $\mathcal{F}_{\text{P4TC}}$ (cont.) – Task *Double-Spending Detection*

TSP input: $(\text{scan_for_fraud}, \phi)$

- (1) Pick $\text{trdb} \neq \text{trdb}'$ in TRDB such that $\text{trdb} = (\cdot, \cdot, \phi, \cdot, \cdot, \text{pid}_{\mathcal{U}}, \cdot, \cdot, \cdot)$ and $\text{trdb}' = (\cdot, \cdot, \phi, \cdot, \cdot, \text{pid}_{\mathcal{U}}, \cdot, \cdot, \cdot)$.[⊥]
- (2) Ask the adversary for a proof $\pi \in \Pi$ corresponding to $\text{pid}_{\mathcal{U}}$ and append $(\text{pid}_{\mathcal{U}}, \pi) \mapsto \text{OK}$ to f_{Π} .
- (3) Receive $\text{pk}_{\mathcal{U}}$ from the bulletin-board $\bar{\mathcal{G}}_{\text{bb}}$ for $\text{pid}_{\mathcal{U}}$.[⊥]

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

[⊥]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: $(\text{verify_guilt}, \text{pk}_{\mathcal{U}}, \pi)$

- (1) Receive $\text{pid}_{\mathcal{U}}$ from the bulletin-board $\bar{\mathcal{G}}_{\text{bb}}$ for key $\text{pk}_{\mathcal{U}}$.[⊥]
- (2) If $f_{\Pi}(\text{pid}_{\mathcal{U}}, \pi)$ is defined, then set $\text{out} := f_{\Pi}(\text{pid}_{\mathcal{U}}, \pi)$ and output (out).
- (3) If $\text{pid}_{\mathcal{U}} \in \mathcal{PID}_{\text{corrupt}}$, then leak $(\text{pid}_{\mathcal{U}}, \pi)$ to the adversary and obtain result out, else set $\text{out} := \text{NOK}$.
- (4) Append $(\text{pid}_{\mathcal{U}}, \pi) \mapsto \text{out}$ to f_{Π} .

Party output: (out)

[⊥]If this does not exist, output \perp and abort.

Fig. 15. The functionality $\mathcal{F}_{\text{P4TC}}$ (cont. from Fig. 4)

For the actual blacklisting the set of all wallet IDs belonging to $\text{pid}_{\mathcal{U}}$ is looked up and the remainder of the task is conducted for every wallet λ separately. $\mathcal{F}_{\text{P4TC}}$ checks how many values of $f_{\Phi}(\lambda, \cdot)$ are already defined and extends them to the first $x_{\text{bl}_{\mathcal{R}}}$ fraud detection IDs, where $x_{\text{bl}_{\mathcal{R}}}$ 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 $x \leq x_{\text{bl}_{\mathcal{R}}}$ 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 $x \leq x_{\text{bl}_{\mathcal{R}}}$ and all wallets λ 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.

A.3 Properties of $\mathcal{F}_{\text{P4TC}}$

In this appendix we discuss why the previously defined ideal functionality $\mathcal{F}_{\text{P4TC}}$ captures an ideal model of a secure and privacy-preserving ETC scheme. Especially, we illustrate how the high-level objectives of a toll collection scheme (cp. Section 2) are reflected in $\mathcal{F}_{\text{P4TC}}$. The properties (P1) to (P8) are consolidated under the term Operator security, while properties (P9) to (P11) are summed up under User Security and Privacy.

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

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

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

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

DR output: (OK)

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

[⊥]If this does not exist, output \perp and abort.

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

Operator Security

At the bottom line, operator security, especially correctness of billing, follows from the fact that \mathcal{F}_{P4TC} represents an incorruptible accountant which manages all wallets and their associated transactions in a single, pervasive database. In Debt Accumulation and Debt Clearance a (possibly malicious) user only inputs a serial number to indicate which previous wallet state should be used. All relevant information is then looked up by \mathcal{F}_{P4TC} internally.

- (P1) **Owner-binding:** Given the serial number of a previous wallet state, \mathcal{F}_{P4TC} checks that the associated wallet belongs to the calling user and thus that it has legitimately been issued to him. If the user is malicious, the adversary is allowed to indicate another corrupted user instead. Only corrupted users are able to swap wallets, they cannot use wallets of honest users. Please note that this does not change the total amount due to the operator. It only changes the party liable, which is unavoidable if corrupted users collude and mutually share their credentials.
- (P2) **Attribute-binding:** As the attributes \mathbf{a}_U and $\mathbf{a}_R^{\text{prev}}$ attached to the user's wallet are internally managed by \mathcal{F}_{P4TC} , the user is unable to claim his wallet contains any other information than it actually does.
- (P3) **Balance-binding:** By the same argument as in (P2) a user is unable to claim an incorrect balance b^{bill} in Debt Clearance. For each transaction, given the previous serial number, \mathcal{F}_{P4TC} looks up the previous balance b^{prev} , calculates the price p and creates a new entry in the transaction database with balance $b := b^{\text{prev}} + p$. The user only learns the price and the new balance, but cannot tamper with it. Assuming that no double-spending occurred, the set of transactions for a particular wallet forms a linked, linear list and hence b^{bill} equals the sum of all prices.
- (P4) **Double-spending Detection:** The same fraud detection ID ϕ occurs in multiple transaction entries if and only if double-spending was committed. In this case the task Double-Spending Detection provides the TSP with the identity pid_U of the respective user and a publicly verifiable proof π that this user has committed double-spending. Again, a fraudulent user cannot elude detection as \mathcal{F}_{P4TC} internally asserts that all transactions which have the same predecessor share the same fraud detection ID ϕ .
- (P5) **Participation Enforcement:** This is handled outside the scope of \mathcal{F}_{P4TC} . As discussed in Section 2 we assume users to be physically identified by cameras if they do not properly participate in Debt Accumulation.

- (P6) **Blacklisting:** Given a user ID pid_U , the task Blacklisting and Recalculation provides the TSP with a set of past and upcoming fraud detection IDs ϕ of all wallets of this user. In order to lock down the upcoming fraud detection IDs \mathcal{F}_{P4TC} pre-fills the mapping f_ϕ with fresh, uniform and independent fraud detection IDs ϕ_i . In Debt Accumulation \mathcal{F}_{P4TC} uses these already determined fraud detection IDs from f_ϕ for the new wallet state. Then, \mathcal{F}_{P4TC} checks whether the fraud detection ID of the new wallet state is contained in the blacklist provided by the RSU. Hence, the user is successfully blacklisted, if the RSU inputs the “correct” blacklist, i.e., a list containing fraud detection IDs from the TSP.
- (P7) **Debt Recalculation:** Within the Blacklisting and Recalculation task, \mathcal{F}_{P4TC} sums up all prices of all past transactions of all wallets of the user in question and outputs the resulting amount b^{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. Again, the user is not able to tamper with the resulting balance as the database of all transaction entries is internally controlled by \mathcal{F}_{P4TC} .
- (P8) **Renegade Expulsion:** This is handled outside the scope of \mathcal{F}_{P4TC} by encoding a limited time of validity into the RSU's attributes (cp. Section 2).

User Security and Privacy

The information leakage needed to assess the level of user privacy is directly determined by the in- and output of \mathcal{F}_{P4TC} . We stress that we only care about privacy for honest, well-behaving, non-blacklisted²⁰ users. Hence, the grayed out steps in Fig. 3 can be ignored.

- (P9) **Unlinkability:** First note that the serial number of the previous transaction s^{prev} is a private input of the user and never output to any party. After Debt Accumulation the RSU only learns the serial number s and fraud detection ID ϕ of the current transaction which are both freshly, uniformly and independently drawn by \mathcal{F}_{P4TC} . Wallet Issue only outputs s and Debt Clearance additionally outputs the final balance b^{bill} . Hence, it is *information-theoretically impossible* to track an honest and well-behaving user across any pair of transactions using any of these

²⁰ Note that the TSP cannot blacklist users alone and the incorruptible DR only cooperates if the user agreed or misbehaved.

numbers. The only “real” information leakage in Debt Accumulation is determined by the user’s and the previous RSU’s attributes \mathbf{a}_U and $\mathbf{a}_R^{\text{prev}}$ which need to be assessed separately (see below).

(P10) Participation Provability: As discussed in Section 2 we assume the user to be physically identified (out-of-scope) if he does not properly participate in Debt Accumulation. In Prove Participation the SA inputs a user ID and a set $S_{\mathcal{R}}^{\text{PP}}$ of serial numbers in question. $\mathcal{F}_{\text{P4TC}}$ checks whether the user participated in any of these transactions. If so, $\mathcal{F}_{\text{P4TC}}$ simply outputs OK to the SA who does not learn anything about any of the user’s transactions beyond that.

(P11) Protection Against False Accusation: Given a user ID and a bit string the task Guilt Verification checks if the bit string has been recorded as a legitimate proof-of-guilt for this user. A proof-of-guilt can only be registered via successful invocation of Double-Spending Detection. Hence soundness of the proof and thus protection against false accusation is guaranteed by the internal bookkeeping of $\mathcal{F}_{\text{P4TC}}$.

$\mathcal{F}_{\text{P4TC}}$ provides unlinkability of transactions (P9) up to information gained from user attributes, RSU attributes, and the total debt. As discussed in Section 2 we assume the attributes to be sufficiently indistinct to not enable any tracking. This is not ensured within the scope of $\mathcal{F}_{\text{P4TC}}$ —apart from outputs to the user, which enable him to check attributes. Their real-world impact on the privacy level crucially depends on the concrete deployment of the system. Please note that this constitutes a line of research on its own.

B Information Leakage and Discussion on Privacy Implications

As briefly discussed in Section 2 and Appendix A.3, 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 π_{P4TC} to be indistinguishable from the ideal functionality $\mathcal{F}_{\text{P4TC}}$, it is ensured that an adversary attacking π_{P4TC} in the real world can only learn as much about a user as an adversary in the ideal model. Table 5 summarizes what an adversary learns about the users in each task. We omitted the serial number s and the fraud detection ID ϕ in the table as these are independently and uniformly drawn randomness and thus cannot be exploited

Table 5. Information an adversary learns about honest users.

Protocol	Leakage					
	pk_U	\mathbf{a}_U	$\mathbf{a}_R/\mathbf{a}_T$	$\mathbf{a}_R^{\text{prev}}$	p	b^{bill}
User Registration	•					
Wallet Issuing	•	•	•			
Debt Accumulation		•	•	•	•	
Debt Clearance	•	(•)		•		•
Prove Participation	•					

(see (P9) in Appendix A.3). In all tasks except Debt Accumulation the user’s public key pk_U is leaked. The variables \mathbf{a}_U , \mathbf{a}_R , $\mathbf{a}_R^{\text{prev}}$ and \mathbf{a}_T refer to attributes of the participating parties. The variable p denotes the price of a Debt Accumulation transaction, and b^{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_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 5).

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 \mathbf{a}_U , \mathbf{a}_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^{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 $\mathbf{a}_{\mathcal{U}}$, the attribute $\mathbf{a}_{\mathcal{R}}^{\text{prev}}$ of the target node equal $\mathbf{a}_{\mathcal{R}}$ 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 $\mathbf{a}_{\mathcal{U}}$ and $\mathbf{a}_{\mathcal{R}}$. 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^{bill}) the problem might become fixed-parameter tractable [5]. 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 Collect²¹ for trucks and analyze the statistics from January, 2018 [32], as a more concrete example. This system uses 24 distinct²² 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.²³ 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 B.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 [63] 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 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.

C Protocol Assumptions 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.

²¹ <https://www.toll-collect.de/en/>

²² The actual system uses more attribute values (i.e., 720) for statistical purposes. However, we only counted attribute values that the price depends on.

²³ 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.

C.1 Algebraic Setting and Assumptions

Our protocol instantiations are based on an asymmetric bilinear group setting $\mathbf{gp} := (G_1, G_2, G_T, e, \mathfrak{p}, g_1, g_2)$. We adopt the following definition from [44].

Definition C.1 (Prime-order Bilinear Group Generator).

A *prime-order bilinear group generator* is a PPT algorithm SetupGrp that on input of a security parameter 1^n outputs a tuple of the form

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

where G_1, G_2, G_T are descriptions of cyclic groups of prime order \mathfrak{p} , $\log \mathfrak{p} = \Theta(n)$, 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*: $\forall a \in G_1, b \in G_2, x, y \in \mathbb{Z}_{\mathfrak{p}}: e(a^x, b^y) = e(a, b)^{xy}$.
- *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.

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

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

and the advantage

$$\text{Adv}_{\text{SetupGrp}, i, \mathcal{A}}^{\text{DDH}}(1^n) := \left| \text{Succs}_{\text{SetupGrp}, i, \mathcal{A}}^{\text{DDH}}(1^n) - \frac{1}{2} \right|$$

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

2. We say that the *SXDH assumption* holds with respect to SetupGrp if the above holds for both $i = 1$ and $i = 2$.

The co-CDH assumption is defined as follows:

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

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

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 SetupGrp if the advantage $\text{Adv}_{\text{SetupGrp}, \mathcal{A}}^{\text{co-DLIN}}(1^n)$ defined by

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

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 ($F_{\mathbf{gp}}$ -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 SetupGrp be a bilinear group generator (cp. [Definition C.1](#)) that outputs descriptions of asymmetric bilinear groups $\mathbf{gp} \leftarrow \text{SetupGrp}(1^n)$. The following building blocks all make use of 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 $stmt$ is contained in L without V learning anything beyond that fact. In a non-interactive zero-knowledge (NIZK) proof, only one message, the proof π , 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 (gp, x, w) . We call gp the group setup, x the statement, and w the witness. Given some gp , let L_{gp} be the language containing all statements x such that $(gp, x, w) \in R$. Let $\text{POK} := (\text{SetupGrp}, \text{SetupPoK}, \text{Prove}, \text{Vfy})$ be a tuple of PPT algorithms such that

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

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

1. *Perfect completeness:* For all $gp \leftarrow \text{SetupGrp}(1^n)$, $\text{CRS}_{\text{pok}} \leftarrow \text{SetupPoK}(gp)$, $(gp, x, w) \in R$, 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

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

is 1.

3. *Perfect F_{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

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

is zero.

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

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

is 1.

4. *Composable Zero-knowledge:* There exists a polynomial-time simulator $(\text{SetupSPoK}, \text{SimProof})$ and hint generator 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

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

is negligible in n .

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

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

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

We wish to point out some remarks.

Remark C.6.

1. The considered language L_{gp} may depend on gp .
2. F_{gp} -extractability actually implies soundness independent of F_{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 ExtractW —given a trapdoor td_{epok} —is able to extract $F_{\text{gp}}(\text{wit})$ for an NP-witness wit for $\text{stmt} \in L_{\text{gp}}$ from any valid proof π . If F_{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 [35, 42] 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_{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 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 $\text{P}X.\text{Prove}$ and $\text{P}X.\text{Vfy}$, for $1 \leq X \leq 3$.

C.2.2.1 Range Proofs

For one particular task²⁴ we need range proofs in order to show that some \mathbb{Z}_p -element λ'_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 [19]. Here, the verifier initially chooses parameters q and t such that every possible λ'_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 λ'_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 [2].

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_{gp} -binding commitment scheme for a bijective function F_{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_{gp} the implicit message space.

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

- SetupGrp takes as input a security parameter 1^n and outputs public parameters 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 gp is given as implicit input to all algorithms.
- Gen is a PPT algorithm, which takes gp as input and outputs public parameters CRS_{com} .
- Com is a PPT algorithm, which takes as input parameters CRS_{com} and a message $m \in \mathcal{M}$ and outputs a commitment c to m and some decommitment value d .
- Open is a deterministic polynomial-time algorithm, which takes as input parameters CRS_{com} , commit-

²⁴ More precisely: The task Wallet Issuing

ment c , an implicit message $M \in \mathcal{M}'$, and opening d . It returns either 0 or 1.

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 COM is a (computationally) *hiding*, F_{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

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

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_{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 SimGen , SimCom and 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

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

is negligible in n .

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

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

is zero.

4. *Extractable*: There exist PPT algorithms ExtGen and 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

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

is negligible in n .

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

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

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 α -message-commitment scheme from Abe et al. [4]. This commitment scheme has message space \mathbb{Z}_p^α , commitment space G_2 and opening value space G_1 . It is statistically hiding, additively homomorphic, equivocal, and F'_{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 $\text{CRS } \text{CRS}_{\text{com}}^1$ 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 [42]. 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'_{gp} -Binding for $F'_{\text{gp}}(m) := m$. In our system, we use this commitment scheme as C2 with 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 σ on a message m using its secret signing key sk such that anybody can publicly verify that σ is a valid signature for m using the public verification key pk of the signer but nobody can feasibly forge a signature without knowing sk .

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

- SetupGrp takes as input a security parameter 1^n and outputs public parameters gp . We assume that gp is given as implicit input to all algorithms.
- Gen takes gp as input and outputs a key pair (pk, sk) . The public key and gp define a message space \mathcal{M} .
- Sgn takes as input the secret key sk and a message $m \in \mathcal{M}$, and outputs a signature σ .
- Vfy takes as input the public key pk , a message $m \in \mathcal{M}$, and a purported signature σ , 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 from Abe et al. [2], 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 [2], 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. [2] 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 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:

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

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

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

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

is negligible with

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

and

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

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 [15].

Our Instantiation

We will make use of two different IND-CCA2-secure encryption schemes:

- We implicitly use the encryption scheme by Cash, Kiltz, and Shoup [26] to realize the secure channels underlying our model.
- We use a variant of Camenisch et al. [20] 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. [20]. 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. [3] 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 [3] (or can also be easily seen by inspecting the original proof in [20]). We present the modified scheme in full detail.

Definition C.10 (Type-3 variant of Camenisch et al. [20]). Let $\text{gp} := (G_1, G_2, G_T, e, \mathfrak{p}, g_1, g_2) \leftarrow \text{SetupGrp}(1^n)$ be as in Definition C.1. Let φ be the dimension of the message space G_1^φ . The algorithms Gen, Enc and Dec are depicted in Figs. 17 to 19.

```

Gen(gp, ϕ)
-----
parse (G1, G2, GT, e, p, g1, g2) := gp
α1, …, αϕ, β0, …, β3, γ1, …, γϕ  $\stackrel{R}{\leftarrow}$  Zp3
sk := ({αi}i=1, …, ϕ, {βi}i=0, …, 3, {γi}i=1, …, ϕ)
ξ1, …, ξ3  $\stackrel{R}{\leftarrow}$  Zp}^*
h̃1 := g1ξ1, h̃2 := g1ξ2, h̃3 := g1ξ3
ĥ1 := g2ξ1, ĥ2 := g2ξ2, ĥ3 := g2ξ3
xi,1 := h̃1αi,1 h̃3αi,3, xi,2 := h̃2αi,2 h̃3αi,3, for i = 1, …, ϕ
yi,1 := ĥ1βi,1 ĥ3βi,3, yi,2 := ĥ2βi,2 ĥ3βi,3, for i = 0, …, 3
zi,1 := ĥ1γi,1 ĥ3γi,3, zi,2 := ĥ2γi,2 ĥ3γi,3, for i = 1, …, ϕ
pk := (h̃1, h̃2, h̃3, ĥ1, ĥ2, ĥ3,
      {xi,1, xi,2}i=1, …, ϕ, {yi,1, yi,2}i=0, …, 3, {zi,1, zi,2}i=1, …, ϕ)
return (pk, sk)

```

Fig. 17. The key generation algorithm Gen of the adapted CCA-secure encryption scheme by Camenisch et al. [20] with parameter φ and message space G_1^φ

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

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.)

A symmetric encryption scheme E is *IND-CCA2-secure* if for all PPT adversaries \mathcal{A} it holds that the

```

Enc(pk, m)
-----
parse (h̃1, h̃2, h̃3, ĥ1, ĥ2, ĥ3, {xi,1, xi,2}i=1, …, ϕ,
      {yi,1, yi,2}i=0, …, 3, {zi,1, zi,2}i=1, …, ϕ) := pk
r, s  $\stackrel{R}{\leftarrow}$  Zp}^*
ũ1 := h̃1r   ũ2 := h̃2s   ũ3 := h̃3r+s
û1 := ĥ1r   û2 := ĥ2s   û3 := ĥ3r+s
ci = mi xi,1r xi,2s for i = 1, …, ϕ
v = ∏i=03 e(ũi, yi,1r yi,2s) ∏i=1ϕ e(ci, zi,1r zi,2s) with ũ0 := g1
u := (ũ1, ũ2, ũ3, û1, û2, û3)
c := (c1, …, cϕ)
c := (u, c, v)
return (c)

```

Fig. 18. The encryption algorithm Enc of the adapted CCA-secure encryption scheme by Camenisch et al. [20] with parameter φ and message space G_1^φ

```

Dec(sk, c)
-----
parse ({αi}i=1, …, ϕ, {βi}i=0, …, 3, {γi}i=1, …, ϕ) := sk
parse (u, c, v) := c
parse (ũ1, ũ2, ũ3, û1, û2, û3) := u
parse (c1, …, cϕ) := c
ũ0 := g1
if v ≠ ∏i=03 e(ũi, û1βi,1 û2βi,2 û3βi,3) ∏i=1ϕ e(ci, û1γi,1 û2γi,2 û3γi,3)
  abort
if e(ũi, g2) ≠ e(g1, ûi) for any i ∈ {1, 2, 3}
  abort
mi := ci ũ1-αi,1 ũ2-αi,2 ũ3-αi,3 for i ∈ {1, …, ϕ}
m := (m1, …, mϕ)
return (m)

```

Fig. 19. The decryption algorithm Dec of the adapted CCA-secure encryption scheme by Camenisch et al. [20] with parameter φ and message space G_1^φ

advantage $\text{Adv}_{E,\mathcal{A}}^{\text{IND-CCA-sym}}(1^n)$ defined by

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

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 [14].

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

- SetupGrp takes as input a security parameter 1^n and outputs public parameters gp . We assume that gp is given as implicit input to all algorithms. The input domain \mathcal{X}_{gp} , the key space \mathcal{K}_{gp} , and the codomain \mathcal{Y}_{gp} may all depend on gp .

- Gen takes gp as input and outputs a key $k \in \mathcal{K}_{\text{gp}}$. (Typically, we have $k \stackrel{R}{\leftarrow} \mathcal{K}_{\text{gp}}$.)
- 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

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

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 [33]. This function is defined by $F(k, x) : \mathbb{Z}_{\mathfrak{p}}^2 \rightarrow G_1$, $(k, x) \mapsto g_1^{\frac{1}{x+k}}$, where $k \stackrel{R}{\leftarrow} \mathbb{Z}_{\mathfrak{p}}$ is the random PRF key. It is secure for inputs $\{0, \dots, n_{\text{PRF}}\} \subset \mathbb{Z}_{\mathfrak{p}}$ under the n_{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 SetupGrp be a bilinear group generator (cf. [Definition C.1](#)) for which the SXDH assumption holds and $\text{gp} := (G_1, G_2, G_T, e, \mathfrak{p}, g_1, g_2, g_T) \leftarrow \text{SetupGrp}(1^n)$ denotes the output of SetupGrp . Furthermore, let $X_1, \dots, X_{m_1} \in G_1$, $x_1, \dots, x_{m_2} \in \mathbb{Z}_{\mathfrak{p}}$, $Y_1, \dots, Y_{m_3} \in G_2$, and $y_1, \dots, y_{m_4} \in \mathbb{Z}_{\mathfrak{p}}$ denote variables in the following types of equations:

- *Pairing-Product Equation (PPE)*:

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

for constants $A_i \in G_1$, $B_i \in G_2$, $t_T \in G_T$, $\gamma_{i,j} \in \mathbb{Z}_{\mathfrak{p}}$.

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

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

for constants $A_i, t_1 \in G_1$, $b_i, \gamma_{i,j} \in \mathbb{Z}_{\mathfrak{p}}$.

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

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

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

– *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_{gp} be a language containing statements described by the conjunction of n_1 pairing-product equations over 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_{gp} , as introduced by [42], is perfectly correct, perfectly sound, and satisfies F_{gp} -extractability [35, 42] 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

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

There, F_{gp} 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 $s \in \mathbb{Z}_p$. It is also known to be composable zero-knowledge [35, 42] as long as for all PPEs in L_{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 td_{spok} . Also note that if these discrete logarithms are not known, there is a workaround which consists of adding new helper variables to L_{gp} [42].

D Full Protocol

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

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

The proof that π_{P4TC} is a GUC-realization of $\mathcal{F}_{\text{P4TC}}$ is given in Appendix E. The style of the presentation follows the same structure as the presentation of the ideal model $\mathcal{F}_{\text{P4TC}}$ in Appendix A.2: Although π_{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. 20. 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 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. 21) 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. 46). 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. 42 and 43) because they are so simple that splitting it each into two yields no advantage.

UC-Protocol π_{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 $\{pk_{\mathcal{R}} \mapsto \mathbf{a}_{\mathcal{R}}\}$ 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 Ω^{dsp} , Ω^{bl} and $\Omega_{\mathcal{R}}^{\text{PP}}$ of double-spending information, blacklisting information and prove participation information.
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 $\Omega_{\mathcal{U}}^{\text{PP}}$ of prove participation information.

II. Behavior

- | | | |
|--------------------------------------|--------------------------------------|--|
| – <i>DR Registration</i> (Fig. 22) | – <i>RSU Certification</i> (Fig. 30) | – <i>Prove Participation</i> (Fig. 40) |
| – <i>TSP Registration</i> (Fig. 24) | – <i>Wallet Issuing</i> (Fig. 32) | – <i>Double-Spending Detection</i> (Fig. 42) |
| – <i>RSU Registration</i> (Fig. 26) | – <i>Debt Accumulation</i> (Fig. 35) | – <i>Guilt Verification</i> (Fig. 43) |
| – <i>User Registration</i> (Fig. 28) | – <i>Debt Clearance</i> (Fig. 38) | – <i>User Blacklisting</i> (Fig. 44) |

Fig. 20. The UC-protocol π_{P4TC}

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.

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

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

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

The alterable components consist of the *RSU commitment* $c_{\mathcal{R}}$ (a commitment on λ , b , u_1^{next} and x^{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^{next} for the next transaction, the *PRF counter* x^{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 6.

D.3 System Setup

To setup the system once (see Fig. 21), the public parameter CRS must be generated in a trustworthy way. The

Setup($1^n, \mathcal{B}$)
$\text{gp} := (G_1, G_2, G_T, e, p, g_1, g_2) \leftarrow \text{SetupGrp}(1^n)$
$\text{CRS}_{\text{com}}^1 \leftarrow \text{C1.Gen}(\text{gp})$
$\text{CRS}_{\text{com}}^2 \leftarrow \text{C2.Gen}(\text{gp})$
$\text{CRS}_{\text{pok}} \leftarrow \text{SetupPoK}(\text{gp})$
$\text{CRS} := (\text{gp}, \mathcal{B}, \text{CRS}_{\text{com}}^1, \text{CRS}_{\text{com}}^2, \text{CRS}_{\text{pok}})$
return CRS

Fig. 21. System Setup Algorithm

CRS consists of a description of the underlying algebraic framework gp , 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 DR , \mathcal{T} , \mathcal{R} and U are all presented with a wrapper protocol π_{P4TC} (see Figs. 22, 24, 26 and 28) and a core protocol (see Figs. 23, 25, 27 and 29). 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.

The DR computes a key pair $(\text{pk}_{DR}, \text{sk}_{DR})$ which can be used to remove the unlinkability of user transactions in case of a dispute between the user and the TSP (see Figs. 22 and 23). The DR could be a (non-governmental) organization trusted by both, users to protect their privacy and the TSP to protect operator security.

The TSP must also generate a key pair (see Figs. 24 and 25). Therefore, the TSP generates several signature key pairs $(\text{pk}_{\mathcal{T}}^{\mathcal{T}}, \text{sk}_{\mathcal{T}}^{\mathcal{T}})$, $(\text{pk}_{\mathcal{T}}^{\text{cert}}, \text{sk}_{\mathcal{T}}^{\text{cert}})$, $(\text{pk}_{\mathcal{T}}^{\mathcal{R}}, \text{sk}_{\mathcal{T}}^{\mathcal{R}})$, where $\text{sk}_{\mathcal{T}}^{\mathcal{T}}$ is used in the Wallet Issuing task to sign the TSP commitment $c_{\mathcal{T}}$ and the user attributes \mathbf{a}_U , $\text{sk}_{\mathcal{T}}^{\text{cert}}$ is used to sign RSU public keys in the RSU Certification task and $\text{sk}_{\mathcal{T}}^{\mathcal{R}}$ is used in the Wallet Issuing task to sign the RSU commitment $c_{\mathcal{R}}$ and the serial number s in place

Table 6. 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 $\mathbf{a}_{\mathcal{U}}$
$c_{\mathcal{R}}$	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
$\text{cert}_{\mathcal{R}}$	$(G_1^3 \times G_2) \times G_1^y \times (G_2^2 \times G_1)$	RSU certificate
$\text{cert}_{\mathcal{T}}^{\mathcal{R}}$	$(G_1^3 \times G_2) \times G_1^y \times (G_2^2 \times G_1)$	TSP certificate
$\sigma_{\mathcal{R}}^{\text{cert}}$	$G_2^2 \times G_1$	signature on $pk_{\mathcal{R}}$ and $\mathbf{a}_{\mathcal{R}}$
$\sigma_{\mathcal{T}}^{\text{cert}}$	$G_2^2 \times G_1$	signature on $pk_{\mathcal{T}}^{\mathcal{R}}$ and $\mathbf{a}_{\mathcal{T}}$
c_{hid}	G_2	hidden ID
d_{hid}	G_1	opening of hidden ID
c'_{seed}	G_2	commitment on the user half of the wallet ID
d'_{seed}	G_1	opening of c'_{seed}
c''_{ser}	G_1^2	commitment on the RSU half of the serial number
d''_{ser}	\mathbb{Z}_p^2	opening of c''_{ser}
ω^{dsp}	$G_1 \times \mathbb{Z}_p \times \mathbb{Z}_p$	double-spending information
ω^{bl}	$G_1 \times \mathbb{Z}_p$	blacklisting information
$\omega_{\mathcal{U}}^{\text{pp}}$	$G_1 \times G_2 \times G_1$	user's prove participation information
$\omega_{\mathcal{R}}^{\text{pp}}$	$G_1 \times G_2$	RSU's prove participation information
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_{PRF}	\mathbb{N}	maximum value of the PRF counter

UC-Protocol π_{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 $(pk_{DR}, sk_{DR}) \leftarrow DRRegistration(CRS)$ (see Fig. 23).
- (4) Record (pk_{DR}, sk_{DR}) internally and call $\bar{\mathcal{G}}_{bb}$ with input (register, pk_{DR}).

DR output: (pk_{DR})

Fig. 22. The UC-protocol π_{P4TC} (cont. from Fig. 20)

DRRegistration(CRS)

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

Fig. 23. DR Registration Core Protocol

of an RSU. The TSP also generates a certificate $cert_{\mathcal{T}}^{\mathcal{R}}$ for its own key $pk_{\mathcal{T}}^{\mathcal{R}}$.

Each RSU must generate a key pair as well (see Figs. 26 and 27). For that purpose, each RSU generates a signature key pair $(pk_{\mathcal{R}}, sk_{\mathcal{R}})$ that is used in the Debt Accumulation task to sign the RSU commitment $c_{\mathcal{R}}$.

Each user also has to generate a key pair (see Figs. 28 and 29). The public key $pk_{\mathcal{U}}$ 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 π_{P4TC} (see Fig. 30) and a core protocol $RSUCertification$ (see Fig. 31). 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

$$\left(\begin{array}{c} (cert_{\mathcal{R}}), \\ (0K) \end{array} \right) \leftarrow RSUCertification \left\langle \begin{array}{c} \mathcal{R}(pk_{\mathcal{T}}, pk_{\mathcal{R}}), \\ \mathcal{T}(pk_{\mathcal{T}}, sk_{\mathcal{T}}, pk_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}}) \end{array} \right\rangle.$$

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 $pk_{\mathcal{R}}$ and the associated certificate $cert_{\mathcal{R}}$ has to be refreshed from time to time. For the ease of presentation we assume that the same RSU (identified by its PID $pid_{\mathcal{R}}$) 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 π_{P4TC} (see Fig. 32) and a core protocol $WalletIssuing$ (see Figs. 33 and 34). 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. 46) after the core protocol. The wrapper protocol internally invokes the core protocol

$$\left(\begin{array}{c} (\tau), \\ (s, htd) \end{array} \right) \leftarrow WalletIssuing \left\langle \begin{array}{c} \mathcal{U}(pk_{DR}, pk_{\mathcal{U}}, sk_{\mathcal{U}}), \\ \mathcal{T}(pk_{DR}, sk_{\mathcal{T}}, \mathbf{a}_{\mathcal{U}}, cert_{\mathcal{T}}^{\mathcal{R}}, bl_{\mathcal{T}}) \end{array} \right\rangle.$$

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

The protocol fulfills four objectives:

UC-Protocol π_{P4TC} (cont.) – Task *TSP Registration*

TSP input: (register, $\mathbf{a}_{\mathcal{T}}$)

- (1) If a key pair $(\mathbf{pk}_{\mathcal{T}}, \mathbf{sk}_{\mathcal{T}})$ has already been recorded, output \perp and abort.
- (2) Obtain CRS CRS from \mathcal{F}_{CRS} .
- (3) Run $(\mathbf{pk}_{\mathcal{T}}, \mathbf{sk}_{\mathcal{T}}, \text{cert}_{\mathcal{T}}^{\mathcal{R}}) \leftarrow \text{TSPRegistration}(\text{CRS}, \mathbf{a}_{\mathcal{T}})$ (see Fig. 25).
- (4) Record $(\mathbf{pk}_{\mathcal{T}}, \mathbf{sk}_{\mathcal{T}})$ and $(\text{cert}_{\mathcal{T}}^{\mathcal{R}})$ internally and call $\bar{\mathcal{G}}_{\text{bb}}$ with input (register, $\mathbf{pk}_{\mathcal{T}}$).

TSP output: $(\mathbf{pk}_{\mathcal{T}})$

Fig. 24. The UC-protocol π_{P4TC} (cont. from Fig. 20)

TSPRegistration(CRS, $\mathbf{a}_{\mathcal{T}}$)

$\text{parse } (\mathbf{gp}, \mathcal{B}, \text{CRS}_{\text{com}}^1, \text{CRS}_{\text{com}}^2, \text{CRS}_{\text{pok}}) := \text{CRS}$
 $(\mathbf{pk}_{\mathcal{T}}^{\mathcal{T}}, \mathbf{sk}_{\mathcal{T}}^{\mathcal{T}}) \leftarrow \text{S.Gen}(\mathbf{gp})$
 $(\mathbf{pk}_{\mathcal{T}}^{\text{cert}}, \mathbf{sk}_{\mathcal{T}}^{\text{cert}}) \leftarrow \text{S.Gen}(\mathbf{gp})$
 $(\mathbf{pk}_{\mathcal{T}}^{\mathcal{R}}, \mathbf{sk}_{\mathcal{T}}^{\mathcal{R}}) \leftarrow \text{S.Gen}(\mathbf{gp})$
 $(\mathbf{pk}_{\mathcal{T}}, \mathbf{sk}_{\mathcal{T}}) := ((\mathbf{pk}_{\mathcal{T}}^{\mathcal{T}}, \mathbf{pk}_{\mathcal{T}}^{\text{cert}}, \mathbf{pk}_{\mathcal{T}}^{\mathcal{R}}), (\mathbf{sk}_{\mathcal{T}}^{\mathcal{T}}, \mathbf{sk}_{\mathcal{T}}^{\text{cert}}, \mathbf{sk}_{\mathcal{T}}^{\mathcal{R}}))$
 $\sigma_{\mathcal{T}}^{\text{cert}} \leftarrow \text{S.Sgn}(\mathbf{sk}_{\mathcal{T}}^{\text{cert}}, (\mathbf{pk}_{\mathcal{T}}^{\mathcal{R}}, \mathbf{a}_{\mathcal{T}}))$
 $\text{cert}_{\mathcal{T}}^{\mathcal{R}} := (\mathbf{pk}_{\mathcal{T}}^{\mathcal{R}}, \mathbf{a}_{\mathcal{T}}, \sigma_{\mathcal{T}}^{\text{cert}})$
 return $(\mathbf{pk}_{\mathcal{T}}, \mathbf{sk}_{\mathcal{T}}, \text{cert}_{\mathcal{T}}^{\mathcal{R}})$

Fig. 25. TSP Registration Core Protocol

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 λ' and λ'' , 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 toss and omit the last message. This way, the wallet ID λ 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 λ under the public key \mathbf{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 λ from g_1^λ due to the hardness of the DLOG problem in G_1 . Therefore, the user splits its share λ' 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 λ'_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 $\mathbf{pk}_{\mathcal{U}}$ under the public key \mathbf{pk}_{DR} of the DR and sends the ciphertext e^* to the TSP. The TSP's share $g_1^{\lambda''}$ and the public key $\mathbf{pk}_{\mathcal{U}}$ 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 $h_{td} := (\mathbf{pk}_{\mathcal{U}}, s, \lambda'', e^*)$. In the case that the user commits a fraud, the TSP sends the h_{td} for each wallet of the fraudulent user to the DR and the DR recovers the wallet ID λ for each wallet. For more details see the User Blacklisting task in Fig. 45.

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 λ and his secret user key $\mathbf{sk}_{\mathcal{U}}$ for the TSP commitment $c_{\mathcal{T}}$. For the preliminary RSU commitment $c_{\mathcal{R}}$, he commits to the wallet ID λ , the balance $b := 0$, a fresh double-spending mask u_1^{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 λ'_i is smaller than \mathcal{B} . More precisely, P1 is used to compute a proof π for a statement $stmnt$ from the

UC-Protocol π_{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 $(pk_{\mathcal{R}}, sk_{\mathcal{R}}) \leftarrow \text{RSURegistration}(CRS)$ (see Fig. 27).
- (4) Store $(pk_{\mathcal{R}}, sk_{\mathcal{R}})$ internally and call $\bar{\mathcal{G}}_{bb}$ with input (register, $pk_{\mathcal{R}}$).

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

Fig. 26. The UC-protocol π_{P4TC} (cont. from Fig. 20)

RSURegistration(CRS)

parse $(gp, \mathcal{B}, CRS_{com}^1, CRS_{com}^2, CRS_{pok}) := CRS$

$(pk_{\mathcal{R}}, sk_{\mathcal{R}}) \leftarrow S.Gen(gp)$

return $(pk_{\mathcal{R}}, sk_{\mathcal{R}})$

Fig. 27. RSU Registration Core Protocol

language $L_{gp}^{(1)}$ defined by

$$\left(\begin{array}{l} \text{pk}_{\mathcal{U}} \\ \text{pk}_{DR} \\ e^* \\ c_{\mathcal{T}} \\ c_{\mathcal{R}} \\ c'_{seed} \\ \Lambda'' \\ \lambda'' \end{array} \right)^T \left\{ \begin{array}{l} \exists \lambda, \lambda', \lambda'_0, \dots, \lambda'_{\ell-1}, r_1, r_2 \in \mathbb{Z}_p; \\ \Lambda, \Lambda', \Lambda'_0, \dots, \Lambda'_{\ell-1}, U_1^{next}, d_{\mathcal{T}}, d_{\mathcal{R}}, d'_{seed} \in G_1; \\ SK_{\mathcal{U}} \in G_2; \\ e(pk_{\mathcal{U}}, g_2) = e(g_1, SK_{\mathcal{U}}) \\ C1.Open(CRS_{com}^1, (\Lambda, pk_{\mathcal{U}}), c_{\mathcal{T}}, d_{\mathcal{T}}) = 1 \\ C1.Open(CRS_{com}^1, (\Lambda, 1, U_1^{next}, g_1), c_{\mathcal{R}}, d_{\mathcal{R}}) = 1 \\ C1.Open(CRS_{com}^1, \Lambda', c'_{seed}, d'_{seed}) = 1 \\ e^* = E1.Enc(pk_{DR}, (\Lambda'_0, \dots, \Lambda'_{\ell-1}, \Lambda'', pk_{\mathcal{U}}); r_1, r_2) \\ \lambda = \lambda' + \lambda'' \\ \Lambda = g_1^\lambda, \Lambda' = g_1^{\lambda'} \\ \lambda' = \sum_{i=0}^{\ell-1} \lambda'_i \cdot \mathcal{B}^i \\ \forall i \in \{0, \dots, \ell-1\}: \\ \quad \lambda'_i \in \{0, \dots, \mathcal{B}-1\} \\ \quad \Lambda'_i = g_1^{\lambda'_i} \end{array} \right. \quad (5)$$

Note that the first equation in Eq. (5) actually proves the knowledge of $g_2^{sk_{\mathcal{U}}}$ (rather than $sk_{\mathcal{U}}$ itself).²⁵ However, computing $g_2^{sk_{\mathcal{U}}}$ without knowing $sk_{\mathcal{U}}$ (only given $pk_{\mathcal{U}}$) 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 $pk_{\mathcal{U}}$ and starts the Blum coin toss for the wallet ID by sending c'_{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 λ'' and starts the other Blum coin toss for the serial number by sending c''_{ser} . Moreover, the TSP sends its own certificate $cert_{\mathcal{T}}^{\mathcal{R}}$ and the attributes $\mathbf{a}_{\mathcal{U}}$ 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 π 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 π) 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 π verifies, the TSP creates two signatures: $\sigma_{\mathcal{T}}$ on $c_{\mathcal{T}}$ together with $\mathbf{a}_{\mathcal{U}}$ 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''_{ser} . At this point the user assembles all part to obtain a fully functional wallet

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

At the end of the protocol, the user returns the wallet τ . The TSP returns the serial number s 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

²⁵ Note that proving a statement $\exists sk_{\mathcal{U}} \in \mathbb{Z}_p : pk_{\mathcal{U}} = g_1^{sk_{\mathcal{U}}}$ instead would not help as we can only extract $g_1^{sk_{\mathcal{U}}}$ from the proof.

UC-Protocol π_{P4TC} (cont.) – Task *User Registration*

User input: (register)

- (1) If a key pair $(\text{pk}_{\mathcal{U}}, \text{sk}_{\mathcal{U}})$ has already been stored, output \perp and abort.
- (2) Obtain CRS CRS from \mathcal{F}_{CRS} .
- (3) Run $(\text{pk}_{\mathcal{U}}, \text{sk}_{\mathcal{U}}) \leftarrow \text{UserRegistration}(\text{CRS})$ (see Fig. 29).
- (4) Store $(\text{pk}_{\mathcal{U}}, \text{sk}_{\mathcal{U}})$ internally and call $\bar{\mathcal{G}}_{\text{bb}}$ with input (register, $\text{pk}_{\mathcal{U}}$).

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

Fig. 28. The UC-protocol π_{P4TC} (cont. from Fig. 20)

UserRegistration(CRS)

parse $(\text{gp}, \mathcal{B}, \text{CRS}_{\text{com}}^1, \text{CRS}_{\text{com}}^2, \text{CRS}_{\text{pok}}) := \text{CRS}$

$y \xleftarrow{\mathbb{R}} \mathbb{Z}_{\mathfrak{p}}$

$(\text{pk}_{\mathcal{U}}, \text{sk}_{\mathcal{U}}) := (g_1^y, y)$

return $(\text{pk}_{\mathcal{U}}, \text{sk}_{\mathcal{U}})$

Fig. 29. User Registration Core Protocol

is presented in two parts: a wrapper protocol π_{P4TC} (see Fig. 35) and a core protocol *DebtAccumulation* (see Figs. 36 and 37). 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. 46) after the core protocol has terminated. The wrapper protocol internally invokes the core protocol

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

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^{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. 40 and 41) to prove to the TSP that the user behaved honestly in this transaction.

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

$$\left(\begin{array}{c} \text{pk}_{\mathcal{T}}^{\mathcal{T}} \\ \text{pk}_{\mathcal{R}}^{\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}}, q_{\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{R}}^{\text{cert}^{\text{prev}}}, \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. \quad (6)$$

This proof essentially shows that the wallet τ 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 ϕ^{prev} from the last transaction

UC-Protocol π_{P4TC} (cont.) – Task *RSU Certification*

RSU input: (certify)

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

(1) At the RSU side:

- Load the internally recorded $(pk_{\mathcal{R}}, sk_{\mathcal{R}})$.[⊥]
- Receive $pk_{\mathcal{T}}$ from the bulletin-board $\overline{\mathcal{G}}_{bb}$ for PID $pid_{\mathcal{T}}$.[⊥]

(2) At the TSP side:

- Load the internally recorded $(pk_{\mathcal{T}}, sk_{\mathcal{T}})$.[⊥]
- Receive $pk_{\mathcal{R}}$ from the bulletin-board $\overline{\mathcal{G}}_{bb}$ for PID $pid_{\mathcal{R}}$.[⊥]
- Check that no mapping $pk_{\mathcal{R}} \mapsto \mathbf{a}'_{\mathcal{R}}$ 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. 31)

$$((cert_{\mathcal{R}}), (OK)) \leftarrow \text{RSUCertification} \langle \mathcal{R}(pk_{\mathcal{T}}, pk_{\mathcal{R}}), \mathcal{T}(pk_{\mathcal{T}}, sk_{\mathcal{T}}, pk_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}}) \rangle.$$

(4) At the RSU side:

- Parse $\mathbf{a}_{\mathcal{R}}$ from $cert_{\mathcal{R}}$.
- Record $cert_{\mathcal{R}}$ internally.

(5) At the TSP side:

- Record $pk_{\mathcal{R}} \mapsto \mathbf{a}_{\mathcal{R}}$ internally.

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

TSP output: (OK)

[⊥]If this does not exist, output \perp and abort.

Fig. 30. The UC-protocol π_{P4TC} (cont. from Fig. 20)




Fig. 31. RSU Certification Core Protocol

UC-Protocol π_{P4TC} (cont.) – Task *Wallet Issuing*

User input: (issue)

TSP input: (issue, \mathbf{a}_U , $bl_{\mathcal{T}}$)

(1) At the user side:

- Load the internally recorded (pk_U, sk_U) .[⊥]
- Receive $pk_{\mathcal{T}}$ from the bulletin-board $\overline{\mathcal{G}}_{bb}$ for PID $pid_{\mathcal{T}}$.[⊥]
- Receive pk_{DR} from the bulletin-board $\overline{\mathcal{G}}_{bb}$ for PID pid_{DR} .[⊥]

(2) At the TSP side:

- Load the internally recorded $(pk_{\mathcal{T}}, sk_{\mathcal{T}})$.[⊥]
- Load the internally recorded $cert_{\mathcal{T}}^R$.[⊥]
- Receive pk_{DR} from the bulletin-board $\overline{\mathcal{G}}_{bb}$ for PID pid_{DR} .[⊥]

(3) Both sides: Run the code of *WalletIssuing* between the user and the TSP (see [Figs. 33](#) and [34](#))

$$((\tau), (s, htd)) \leftarrow \text{WalletIssuing} \langle \mathcal{U}(pk_{DR}, pk_U, sk_U), \mathcal{T}(pk_{DR}, sk_{\mathcal{T}}, \mathbf{a}_U, cert_{\mathcal{T}}^R, bl_{\mathcal{T}}) \rangle.$$

(4) At the user side:

- Run the code of *WalletVerification*($pk_{\mathcal{T}}, pk_U, \tau$) (see [Fig. 46](#)).
- If *WalletVerification* returns 0, output \perp and abort.
- Record τ internally.
- Parse s and \mathbf{a}_U from τ .

(5) At the TSP side:

- Insert htd into *HTD*.

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

TSP output: (s)

[⊥]If this does not exist, output \perp and abort.

Fig. 32. The UC-protocol π_{P4TC} (cont. from [Fig. 20](#))




Fig. 33. Wallet Issuing Core Protocol




Fig. 34. Wallet Issuing Core Protocol (cont. from Fig. 33)

has been computed correctly. It also shows that c_{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 ϕ 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 π verifies and the fraud detection ID ϕ is on the RSU blacklist $bl_{\mathcal{R}}$ 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 information: The blacklisting information $\omega^{\text{bl}} := (\phi, p)$ can be used at a later point in the User Blacklisting task (cf. Fig. 44) to calculate the total fee of a fraudulent user. The double-spending 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. 42). The RSU's prove participation information $\omega_{\mathcal{R}}^{\text{PP}} := (s, c_{\text{hid}})$ can be used later in the Prove Participation task (cf. Fig. 40).

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's prove participation 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 information ω^{bl} , ω^{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 π_{P4TC} (see Fig. 38) and a core protocol DebtClearance (see Fig. 39). 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

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

UC-Protocol π_{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 $(pk_{\mathcal{U}}, sk_{\mathcal{U}})$.[⊥]
- Receive $pk_{\mathcal{T}}$ from the bulletin-board $\overline{\mathcal{G}}_{\text{bb}}$ for PID $pid_{\mathcal{T}}$.[⊥]
- Load the internally recorded token τ^{prev} for serial number s^{prev} .[⊥]

(2) At the RSU side:

- Load the internally recorded $(pk_{\mathcal{R}}, sk_{\mathcal{R}})$.[⊥]
- Load the internally recorded $\text{cert}_{\mathcal{R}}$.[⊥]
- Receive $pk_{\mathcal{T}}$ from the bulletin-board $\overline{\mathcal{G}}_{\text{bb}}$ for PID $pid_{\mathcal{T}}$.[⊥]

(3) Both sides: Run the code of DebtAccumulation between the user and the RSU (see Figs. 36 and 37)

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

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. 46).
- If WalletVerification returns 0, output \perp and abort.
- Record τ and $\omega_{\mathcal{U}}^{\text{pp}}$ internally.
- Parse s , $\text{cert}_{\mathcal{R}}$, p and b from τ .
- 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 ϕ from ω^{bl} .

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

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

[⊥]If this does not exist, output \perp and abort.

Fig. 35. The UC-protocol π_{P4TC} (cont. from Fig. 20)




Fig. 36. Debt Accumulation Core Protocol

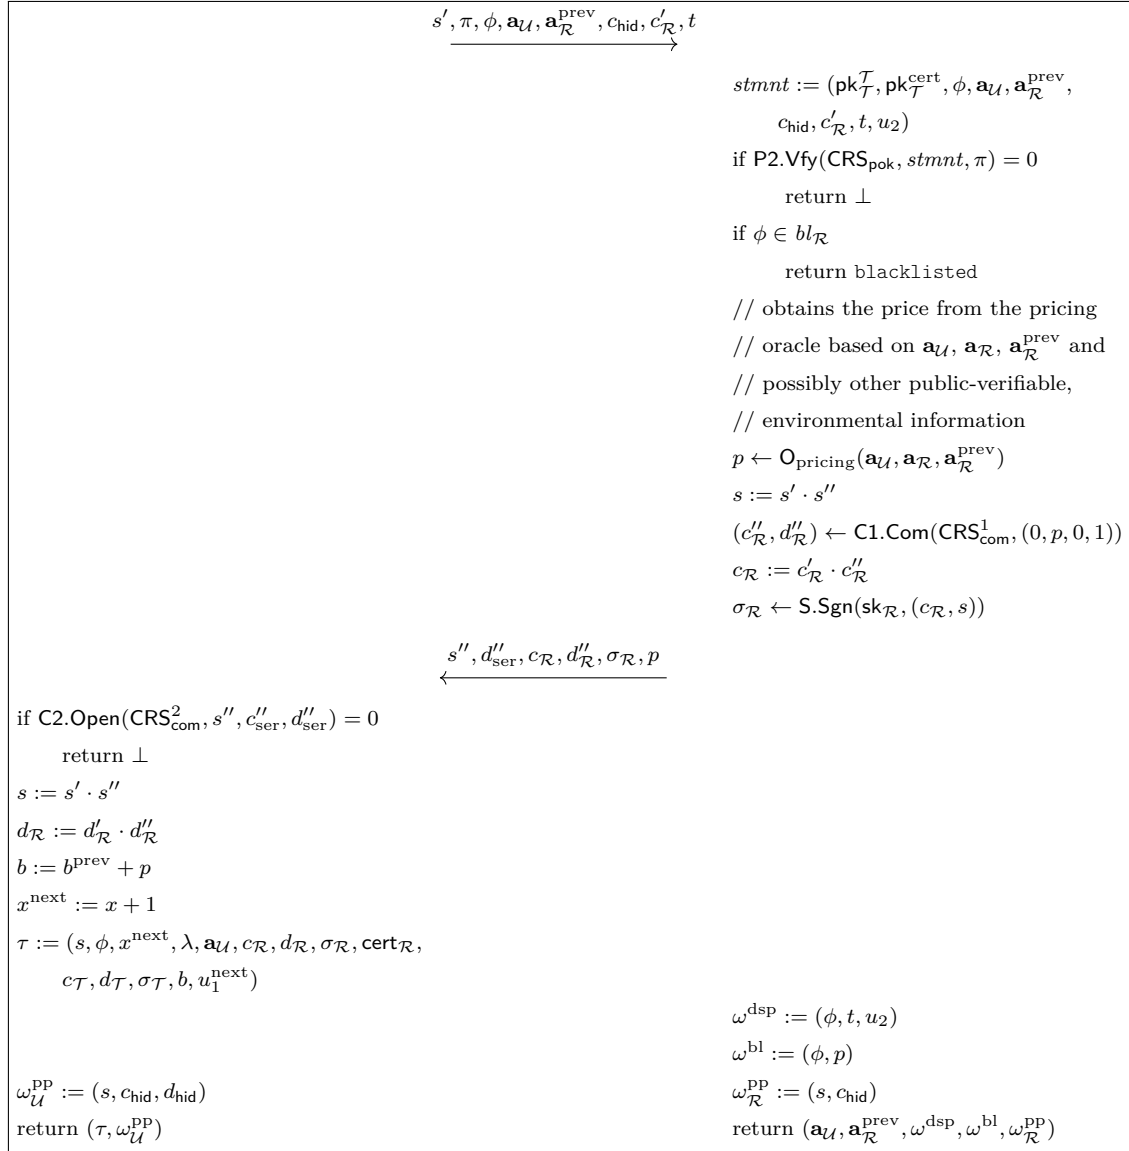


Fig. 37. Debt Accumulation Core Protocol (cont. from Fig. 36)

UC-Protocol π_{P4TC} (cont.) – Task *Debt Clearance*

User input: $(\text{clear_debt}, s^{\text{prev}})$

TSP input: (clear_debt)

(1) At the user side:

- Load the internally recorded $(\text{pk}_{\mathcal{U}}, \text{sk}_{\mathcal{U}})$.[⊥]
- Receive $\text{pk}_{\mathcal{T}}$ from the bulletin-board $\overline{\mathcal{G}}_{\text{bb}}$ for PID $\text{pid}_{\mathcal{T}}$.[⊥]
- Load the internally recorded token τ^{prev} for serial number s^{prev} .[⊥]

(2) At the TSP side:

- Load the internally recorded $(\text{pk}_{\mathcal{T}}, \text{sk}_{\mathcal{T}})$.[⊥]

(3) Both sides: Run the code of DebtClearance between the user and the TSP (see Fig. 39)

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

(4) At the TSP side:

- Record ω^{bl} and ω^{dsp} internally.
- Parse ϕ and $-b^{\text{bill}}$ from ω^{bl} .

User output: (b^{bill})

TSP output: $(\text{pk}_{\mathcal{U}}, \phi, b^{\text{bill}})$

[⊥]If this does not exist, output \perp and abort.

Fig. 38. The UC-protocol π_{P4TC} (cont. from Fig. 20)

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 $t := \text{sk}_{\mathcal{U}}u_2 + u_1 \bmod \mathfrak{p}$ and the fraud detection ID ϕ for this transaction. He then continues by preparing a proof of knowledge. More precisely, P3 is used to compute a proof π for a statement

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

$$\left(\begin{array}{l} \text{pk}_{\mathcal{U}} \\ \text{pk}_{\mathcal{T}} \\ \text{pk}_{\mathcal{T}}^{\text{cert}} \\ \phi \\ \mathbf{a}_{\mathcal{U}} \\ \mathbf{a}_{\mathcal{R}}^{\text{prev}} \\ B^{\text{prev}} \\ t \\ u_2 \end{array} \right)^{\text{T}} \left\{ \begin{array}{l} \exists \lambda, x, u_1, \text{sk}_{\mathcal{U}} \in \mathbb{Z}_{\mathfrak{p}}; \phi^{\text{prev}}, s^{\text{prev}}, X, \Lambda, U_1, \\ d_{\mathcal{R}}^{\text{prev}}, d_{\mathcal{T}} \in G_1; \text{pk}_{\mathcal{R}}^{\text{prev}} \in G_1^3, c_{\mathcal{R}}^{\text{prev}}, c_{\mathcal{T}} \in G_2; \\ \sigma_{\mathcal{R}}^{\text{prev}}, \sigma_{\mathcal{R}}^{\text{cert}^{\text{prev}}}, \sigma_{\mathcal{T}} \in G_2^2 \times G_1; \\ \text{C1.Open}(\text{CRS}_{\text{com}}^1, (\Lambda, \text{pk}_{\mathcal{U}}), c_{\mathcal{T}}, d_{\mathcal{T}}) = 1 \\ \text{C1.Open}(\text{CRS}_{\text{com}}^1, (\Lambda, B^{\text{prev}}, U_1, X), c_{\mathcal{R}}^{\text{prev}}, d_{\mathcal{R}}^{\text{prev}}) = 1 \\ \text{S.Vfy}(\text{pk}_{\mathcal{T}}^{\text{cert}}, \sigma_{\mathcal{T}}, (c_{\mathcal{T}}, \mathbf{a}_{\mathcal{U}})) = 1 \\ \text{S.Vfy}(\text{pk}_{\mathcal{R}}^{\text{prev}}, \sigma_{\mathcal{R}}^{\text{prev}}, (c_{\mathcal{R}}^{\text{prev}}, s^{\text{prev}})) = 1 \\ \text{S.Vfy}(\text{pk}_{\mathcal{T}}^{\text{cert}}, \sigma_{\mathcal{R}}^{\text{cert}^{\text{prev}}}, (\text{pk}_{\mathcal{R}}^{\text{prev}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}})) = 1 \\ \phi^{\text{prev}} = \text{PRF}(\lambda, x - 1), \phi = \text{PRF}(\lambda, x), \\ t = \text{sk}_{\mathcal{U}}u_2 + u_1 \\ \text{pk}_{\mathcal{U}} = g_1^{\text{sk}_{\mathcal{U}}}, U_1 = g_1^{u_1}, X = g_1^x, \Lambda = g_1^\lambda \end{array} \right. \quad (7)$$

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_{hid} .

The user then sends $(\text{pk}_{\mathcal{U}}, \pi, \phi, \mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}}, b^{\text{prev}}, t)$ 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.




Fig. 39. Debt Clearance Core Protocol

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

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.²⁶ 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.

The task is presented in two parts: a wrapper protocol π_{P4TC} (see Fig. 40) and a core protocol $\text{ProveParticipation}$ (see Fig. 41). In the wrapper protocol, the SA interacts with other UC entities and processes the set $S_{\mathcal{R}}^{\text{PP}}$ 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 RSU’s internally recorded set $\Omega_{\mathcal{R}}^{\text{PP}}$ of all prove participation information that contain serial numbers that are also in $S_{\mathcal{R}}^{\text{PP}}$. Afterwards, the wrapper protocol invokes the core protocol

$$\left(\begin{array}{c} (\text{out}_{\mathcal{U}}), \\ (\text{out}_{SA}) \end{array} \right) \leftarrow \text{ProveParticipation} \left\langle \begin{array}{c} \mathcal{U}(\Omega_{\mathcal{U}}^{\text{PP}}), \\ SA(\text{pk}_{\mathcal{U}}, S_{\mathcal{R}}^{\text{PP}}, \Omega_{\mathcal{R}}^{\text{PP}}) \end{array} \right\rangle.$$

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 prove participation information $\Omega_{\mathcal{R}}^{\text{PP}}$ from the RSU as input. The user gets his internally recorded set $\Omega_{\mathcal{U}}^{\text{PP}}$ of all prove participation information as input.

²⁶ 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.

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 his prove participation 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_{hid} is the opening of c_{hid} under $\text{pk}_{\mathcal{U}}$. If both checks succeed, the SA’s output bit out_{SA} is set to OK, if at least one check fails, out_{SA} is set to NOK. At the end of the protocol both parties output their bits $\text{out}_{\mathcal{U}}$ and out_{SA} , respectively.²⁷

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 information ω^{dsp} collected by the RSUs are periodically transmitted to the TSP’s database, which is regularly checked for two double-spending information associated with the same fraud detection ID. If the database contains two such records, the Double-Spending Detection task (see Fig. 42) can be used by the TSP to extract the public key of the user these double-spending 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 fraud detection ID ϕ as input and searches the internal database for two double-spending information $\omega^{\text{dsp}} = (\phi, t, u_2)$ and $\omega^{\text{dsp}'} = (\phi, t', u_2')$ that contain the same fraud detection ID $\phi = \phi'$ but not the same double-spending randomness $u_2 \neq u_2'$. Then the fraudulent user’s secret key can be recovered as $\text{sk}_{\mathcal{U}} := (t - t') \cdot (u_2 - u_2')^{-1} \bmod \mathfrak{p}$. 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. 43).

Every user that is convicted of double-spending is added to the blacklist via the User Blacklisting task (cf. Figs. 44 and 45) and additional measures are taken (these are out of scope).

²⁷ 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.

UC-Protocol π_{P4TC} (cont.) – Task *Prove Participation*

User input: (prove_participation)

SA input: (prove_participation, $\text{pk}_U, S_{\mathcal{R}}^{\text{PP}}$)

(1) At the SA side:

- Load the internally recorded set $\Omega_{\mathcal{R}}^{\text{PP}}$ of all prove participation transaction information with serial numbers in $S_{\mathcal{R}}^{\text{PP}}$.

(2) At the user side:

- Load the internally recorded set Ω_U^{PP} of all prove participation transaction information.

(3) Both sides: Run the code of ProveParticipation between the User and the SA (see Fig. 41)

$$((\text{out}_U), (\text{out}_{SA})) \leftarrow \text{ProveParticipation} \langle \mathcal{U}(\Omega_U^{\text{PP}}), SA(\text{pk}_U, S_{\mathcal{R}}^{\text{PP}}, \Omega_{\mathcal{R}}^{\text{PP}}) \rangle.$$

User output: (out_U)

SA output: (out_{SA})

\perp If this does not exist, output \perp and abort.

Fig. 40. The UC-protocol π_{P4TC} (cont. from Fig. 20)




Fig. 41. Prove Participation Core Protocol

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

TSP input: $(\text{scan_for_fraud}, \phi)$

- (1) Load the internally recorded set Ω^{dsp} of all double-spending information.
- (2) Pick double-spending information $\omega^{\text{dsp}} = (\phi, t, u_2)$ and $\omega^{\text{dsp}'} = (\phi', t', u_2')$ from the database Ω^{dsp} , such that $\phi = \phi'$ and $u_2 \neq u_2'$.[⊥]
- (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: $(\text{pk}_{\mathcal{U}}, \pi)$

[⊥]If this does not exist, output \perp and abort.

Fig. 42. The UC-protocol π_{P4TC} (cont. from Fig. 20)

D.11 Guilt Verification

Whether a user is indeed guilty of double-spending can be verified using the Guilt Verification task depicted in Fig. 43. This algorithm may be run by anyone, in particular by justice. Essentially, the algorithm checks if a given public user key $\text{pk}_{\mathcal{U}}$ and a proof of guilt π match. This is easily accomplished because they match if and only if $g_1^\pi = \text{pk}_{\mathcal{U}}$ holds. This equation holds if and only if π equals the user's secret key $\text{sk}_{\mathcal{U}}$ that was recovered using the Double-Spending Detection task (cf. Fig. 42).

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 π_{P4TC} (see Fig. 44) and a core protocol **BlacklistingAndRecalculation** (see Fig. 45). The

wrapper protocol interacts with other UC entities, pre-processes the input and afterwards internally invokes the core protocol

$$\left(\begin{array}{l} (\text{OK}), \\ (\Phi_{\mathcal{U}}) \end{array} \right) \leftarrow \text{BlacklistingAndRecalculation} \left\langle \begin{array}{l} DR(\text{pk}_{DR}, \text{sk}_{DR}, \text{pk}_{\mathcal{U}}^{DR}), \\ \mathcal{T}(HTD_{\mathcal{U}}) \end{array} \right\rangle.$$

The DR gets as input its public and private key $(\text{pk}_{DR}, \text{sk}_{DR})$ together with the public key $\text{pk}_{\mathcal{U}}^{DR}$ of the user to be blacklisted. The TSP gets a set $HTD_{\mathcal{U}}$ containing all his hidden user trapdoors from the current billing period as input.²⁸ 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 λ for every $htd := (\text{pk}_{\mathcal{U}}^{\mathcal{T}}, s, \lambda'', e^*) \in HTD$. To this end, the DR decrypts e^* to get $(\Lambda'_0, \dots, \Lambda'_{\ell-1}, \Lambda'', \text{pk}_{\mathcal{U}}^{\mathcal{T}})$. Firstly, the DR checks if the decrypted public key $\text{pk}_{\mathcal{U}}^{\mathcal{T}}$ equals the expected public $\text{pk}_{\mathcal{U}}^{DR}$ of the correct user in question.²⁹ This way, the DR cannot be tricked into recovering the wallet ID of another (possibly innocent) user. Since each λ'_i is small ($\lambda'_i < \mathcal{B}$), the DR can compute the discrete logarithms

²⁸ 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.

²⁹ N.b.: The public key $\text{pk}_{\mathcal{U}}^{\mathcal{T}}$ is stored redundantly: in clear as part of $htd := (\text{pk}_{\mathcal{U}}^{\mathcal{T}}, s, \lambda'', e^*)$ and encrypted as part of the cipher text $e^* = \text{Enc}(\Lambda'_0, \dots, \Lambda'_{\ell-1}, \Lambda'', \text{pk}_{\mathcal{U}}^{\mathcal{T}})$. 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 $\text{pk}_{\mathcal{U}}^{\mathcal{T}}$.

UC-Protocol π_{P4TC} (cont.) – Task *Guilt Verification*

Party input: $(\text{verify_guilt}, \text{pk}_{\mathcal{U}}, \pi)$

- (1) Receive $\text{pid}_{\mathcal{U}}$ from the bulletin-board $\overline{\mathcal{G}}_{\text{bb}}$ for key $\text{pk}_{\mathcal{U}}$.[⊥]
- (2) If $g_1^\pi = \text{pk}_{\mathcal{U}}$, then $\text{out} := \text{OK}$, else $\text{out} := \text{NOK}$.

Party output: (out)

[⊥]If this does not exist, output \perp and abort.

Fig. 43. The UC-protocol π_{P4TC} (cont. from Fig. 20)

UC-Protocol π_{P4TC} (cont.) – Task *User Blacklisting*

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

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

- (1) At the DR side:

- Load the internally recorded $(\text{pk}_{DR}, \text{sk}_{DR})$.[⊥]
- Receive $\text{pk}_{\mathcal{U}}^{DR}$ from the bulletin-board $\overline{\mathcal{G}}_{\text{bb}}$ for PID $\text{pid}_{\mathcal{U}}^{DR}$.[⊥]

- (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 `BlacklistingAndRecalculation` between the DR and the TSP (see Fig. 45)

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

- (4) At the TSP side:

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

DR output: (OK)

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

[⊥]If this does not exist, output \perp and abort.

Fig. 44. The UC-protocol π_{P4TC} (cont. from Fig. 20)

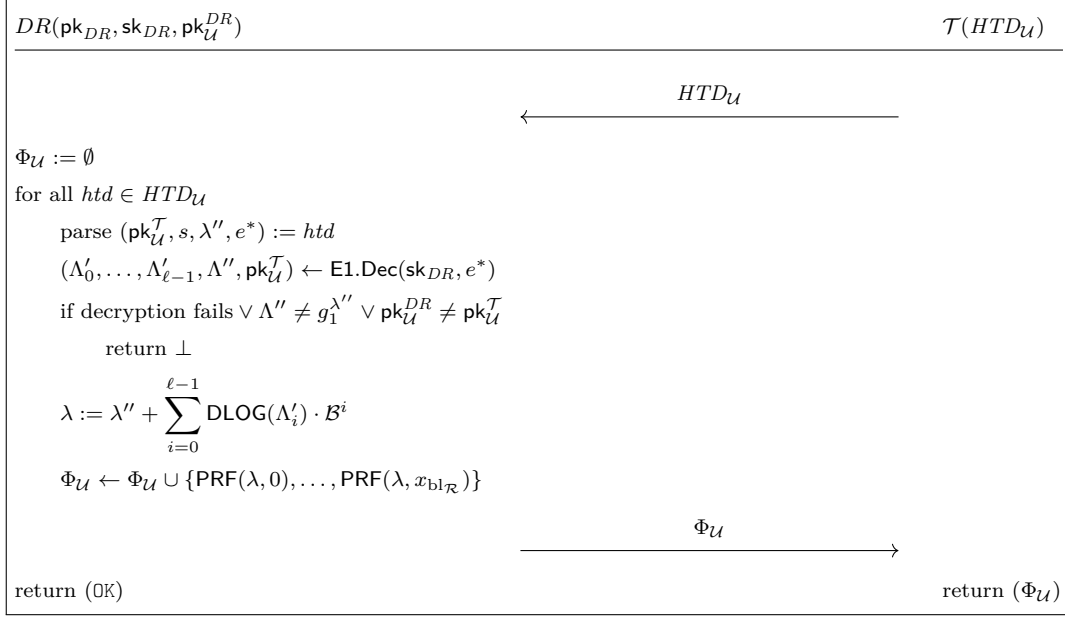


Fig. 45. User Blacklisting Core Protocol

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 λ'' of the wallet ID is consistent to the decrypted $\Lambda'' \stackrel{?}{=} g_1^{\lambda''}$. (Remember that λ'' 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 $\{\text{PRF}(\lambda, 0), \dots, \text{PRF}(\lambda, x_{\text{bl}_{\mathcal{R}}})\}$ of fraud detection IDs for every $htd \in HTD$ to the TSP. The blacklist parameter $x_{\text{bl}_{\mathcal{R}}}$ is chosen in such a way that a user is expected to perform at most $x_{\text{bl}_{\mathcal{R}}}$ 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 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^{bill} the user owes.

After the wrapper protocol has terminated, the fraud detection IDs $\Phi_{\mathcal{U}}$ are added to the RSU blacklist $\text{bl}_{\mathcal{R}}$ and the user's public key $\text{pk}_{\mathcal{U}}$ is added on the TSP blacklist $\text{bl}_{\mathcal{T}}$.

In the Wallet Issuing task the TSP uses the blacklist $\text{bl}_{\mathcal{T}}$ to prevent a user that did not pay its invoice from

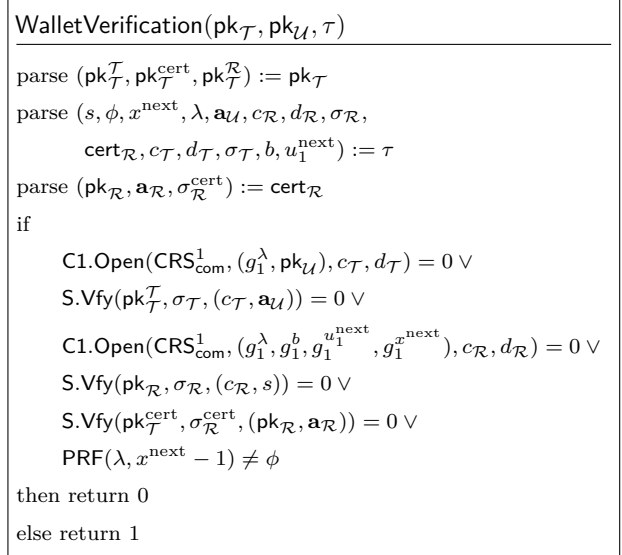


Fig. 46. Algorithm for Wallet Verification

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 $\text{bl}_{\mathcal{R}}$.

D.13 Wallet Verification

The WalletVerification algorithm is depicted in Fig. 46. 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 $c_{\mathcal{T}}$ and $c_{\mathcal{R}}$

are valid and contain the values they are supposed to contain, that $\sigma_{\mathcal{T}}$ is a valid signature under $\text{sk}_{\mathcal{T}}^{\mathcal{T}}$ of $c_{\mathcal{T}}$ and $\mathbf{a}_{\mathcal{U}}$, that $\sigma_{\mathcal{R}}$ is a valid signature under $\text{sk}_{\mathcal{R}}$ of $c_{\mathcal{R}}$ and s , that the certificate $\text{cert}_{\mathcal{R}}$ containing $\text{pk}_{\mathcal{R}}$ is valid and that the fraud detection id ϕ 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 Security Proof

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

Theorem E.1 (Security Statement). *Assume that the SXDH-problem is hard for $\text{gp} := (G_1, G_2, G_T, e, \mathfrak{p}, g_1, g_2)$, the Co-CDH problem is hard for (G_1, G_2) , the n_{PRF} -DDHI problem is hard for G_1 , 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*

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

holds under 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, or
4. all RSUs, TSP and SA as well as a subset of users.

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 E.1](#) in three steps:

- In [Appendix E.2](#) we first show some structural properties of $\mathcal{F}_{\text{P4TC}}^{\bar{\mathcal{G}}_{\text{bb}}}$.
- In [Appendix E.3](#), [Theorem E.10](#) proves [Theorem E.1](#) for the corruption scenario [1](#) and [2](#). We call this case “Operator Security”.
- In [Appendix E.4](#), [Theorem E.25](#) proves [Theorem E.1](#) for the corruption scenario [3](#) and [4](#). We call this case “User Security and Privacy”.

Proof of [Theorem E.1](#). The theorem is immediately implied by [Theorems E.10](#) and [E.25](#). \square

Before giving the proof in full detail and complexity in [Appendices E.2](#) to [E.4](#), [Appendix E.1](#) explains our adversarial model.

E.1 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.

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 is tightly related to deniability which is not part of our desired properties.

Secondly, we only consider adversaries \mathcal{Z} that corrupt one of the following sets³⁰:

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

³⁰ Note that “subset” also includes the empty or full set.

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.

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.

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




Fig. 47. An entry $trdb \in TRDB$ visualized as an element of a directed graph

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.2 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 of the ideal functionality is mostly obvious. In contrast, our ideal functionality \mathcal{F}_{P4TC} is already a complex system on its own with polynomially many parties that can reactively interact forever, i.e., \mathcal{F}_{P4TC} itself has no inherent exit point except that at some point the polynomially bounded runtime of the environment is exhausted. In this appendix no security reduction occurs, because we only consider the ideal functionality \mathcal{F}_{P4TC} . We start with a series of simple lemmas which also help to develop a good conception about how the individual tasks/transactions are connected. Moreover, these lemmas are closely associated to the desired properties of a toll collection scheme (cp. [Appendix A.3](#) and [Section 2](#)).

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

$$trdb = (s^{\text{prev}}, s, \phi, x, \lambda, pid_U, pid_R, p, b).$$

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 $(s, \phi, x, \lambda, pid_U, b)$ and identified by s . The edge to the predecessor is identified by (s^{prev}, s) and labeled with (pid_R, p) . See [Fig. 47](#) 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 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 E.2 (Ideal Transaction Graph (informal)).

The transaction database $TRDB = \{trdb_i\}$ with $trdb = (s^{\text{prev}}, s, \phi, x, \lambda, pid_{\mathcal{U}}, pid_{\mathcal{R}}, p, b)$ is a directed, labeled graph as defined above. This graph is called the *Ideal Transaction Graph*.

Lemma E.3 (Ideal Transaction Forest). *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 E.4 (Tree-wise Uniqueness of the Wallet ID).

The wallet ID λ 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 λ 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 λ 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 E.5. *Within a tree of the Ideal Transaction Graph the PID $pid_{\mathcal{U}}$ of the corresponding user is constant.*

Proof. Same proof as “ \Leftarrow ” in the proof of **Lemma E.4**. \square

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

Lemma E.6. *Within a tree of $TRDB$, every node $trdb = (s^{\text{prev}}, s, \phi, x, \lambda, pid_{\mathcal{U}}, pid_{\mathcal{R}}, p, b)$ has depth x and all nodes of the same depth in the same tree have the same fraud detection ID ϕ . 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 ϕ is chosen. In the scope of Debt Accumulation or Debt Clearance, x is calculated as $x := x^{\text{prev}} + 1$, where by induction x^{prev} is the depth of its predecessor. With respect to ϕ 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 λ and x . \square

Lemma E.7 (Billing Correctness). *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 E.6**. \square

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

Lemma E.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_{Π} 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 E.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. Let the TSP be honest. 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. We prove both claims separately.

1. Let $TRDB_{\mathcal{U}} \subseteq TRDB$ be the subset of all transaction entries $trdb = (\cdot, \cdot, \cdot, \cdot, pid_{\mathcal{U}}, \cdot, p, \cdot)$ 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 λ 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^{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

E.3 Proof of Operator Security

In this appendix we show the following theorem.

Theorem E.10 (Operator Security). *Under the assumptions of Theorem E.1*

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

holds under static corruption of

1. a subset of users, or
2. all users and a subset of RSUs, TSP and SA.

The definition of the UC-simulator $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$ for Theorem E.10 can be found in Figs. 48 to 51. Please note that while the real protocol π_{P4TC} lives in the $(\mathcal{F}_{\text{CRS}}, \bar{\mathcal{G}}_{\text{bb}})$ -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.

While the protocol executes, the simulator $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$ records certain information similar to what the parties or the ideal functionality internally record, namely the set of simulated double-spending detection information $\bar{\Omega}^{\text{dsp}}$, the set of simulated prove participation information $\bar{\Omega}_{\mathcal{U}}^{\text{pp}}$, and the simulated transaction graph \overline{TRDB} . Basically, $\bar{\Omega}^{\text{dsp}}$, $\bar{\Omega}_{\mathcal{U}}^{\text{pp}}$ and \overline{TRDB} correspond to Ω^{dsp} , $\Omega_{\mathcal{U}}^{\text{pp}}$ and $TRDB$ resp., but exist in the head of the simulator and are augmented by additional information. The simulator uses them as “lookup tables” to keep up a consistent simulation in later parts of the protocol. Obviously, this implies information is stored redundantly: In the head of $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$ as $\bar{\Omega}_{\mathcal{U}}^{\text{pp}}$ and \overline{TRDB} and inside the ideal functionality $\mathcal{F}_{\text{P4TC}}$ (in case of $TRDB$) or the environment (in case of $\Omega_{\mathcal{U}}^{\text{pp}}$ for a corrupted user).³¹ A crucial part of the security proof is to show that these sets stay in sync.

Before starting with the security proof we explain the *Simulated Transaction Graph* \overline{TRDB} and the additional information beyond the *Ideal Transaction Graph* (cp. Definition E.2) in more details. The *Ideal Transaction Graph* is a theoretical construct that helps us to link the interactions of the parties across the various tasks our protocol provides. An *Simulated Transaction*

³¹ Note, that we get rid of Ω^{dsp} during the proof, because honest user are only dummy parties and only $\bar{\Omega}^{\text{dsp}}$ remains.

Simulator $\mathcal{S}_{\pi_{P4TC}}^{\text{sys-sec}}$

Setup:

- (1) Run a modified version of the algorithm $\text{CRS} \leftarrow \text{Setup}(1^n)$ with
 - (a) $\text{CRS}_{\text{com}}^2 \leftarrow \text{C2.Gen}$ being replaced by $(\text{CRS}_{\text{com}}^2, \text{td}_{\text{eqcom}}) \leftarrow \text{C2.SimGen}$, and
 - (b) $\text{CRS}_{\text{pok}} \leftarrow \text{SetupPoK}$ being replaced by $(\text{CRS}_{\text{pok}}, \text{td}_{\text{epok}}) \leftarrow \text{SetupEPoK}$.
- (2) Record CRS , td_{eqcom} and td_{epok} .
- (3) Set $\overline{\text{TRDB}} := \emptyset$.
- (4) Set $\overline{\mathcal{Q}}^{\text{dsp}} := \emptyset$.
- (5) Set $\overline{\mathcal{Q}}_{\mathcal{U}}^{\text{pp}} := \emptyset$.

DR Registration: Upon receiving $(\text{registering_dr}, \text{pid}_{DR})$ run $(\text{pk}_{DR}, \text{sk}_{DR}) \leftarrow \text{DRRegistration}(\text{CRS})$, record $\text{pid}_{DR} \mapsto (\text{pk}_{DR}, \text{sk}_{DR})$, and return pk_{DR} to \mathcal{F}_{P4TC} .

TSP Registration: Upon receiving $(\text{registering_tsp}, \text{pid}_{\mathcal{T}}, \mathbf{a}_{\mathcal{T}})$ run $(\text{pk}_{\mathcal{T}}, \text{sk}_{\mathcal{T}}, \text{cert}_{\mathcal{T}}^{\mathcal{R}}) \leftarrow \text{TSPRegistration}(\text{CRS}, \mathbf{a}_{\mathcal{T}})$, record $\text{pid}_{\mathcal{T}} \mapsto (\text{pk}_{\mathcal{T}}, \text{sk}_{\mathcal{T}}, \text{cert}_{\mathcal{T}}^{\mathcal{R}})$, and return $\text{pk}_{\mathcal{T}}$ to \mathcal{F}_{P4TC} .

RSU Registration: Upon receiving $(\text{registering_rsu}, \text{pid}_{\mathcal{R}})$ run $(\text{pk}_{\mathcal{R}}, \text{sk}_{\mathcal{R}}) \leftarrow \text{RSURegistration}(\text{CRS})$, and record $\text{pid}_{\mathcal{R}} \mapsto (\text{pk}_{\mathcal{R}}, \text{sk}_{\mathcal{R}})$, and return $\text{pk}_{\mathcal{R}}$ to \mathcal{F}_{P4TC} .

User Registration: If user is corrupted, then nothing to do, as this is a local algorithm.

If user is honest: Upon receiving $(\text{registering_user}, \text{pid}_{\mathcal{U}})$ run $(\text{pk}_{\mathcal{U}}, \text{sk}_{\mathcal{U}}) \leftarrow \text{UserRegistration}(\text{CRS})$, record $\text{pid}_{\mathcal{U}} \mapsto (\text{pk}_{\mathcal{U}}, \text{sk}_{\mathcal{U}})$, and return $\text{pk}_{\mathcal{U}}$ to \mathcal{F}_{P4TC} .

RSU Certification: Upon receiving $(\text{certifying_rsu}, \text{pid}_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}}) \dots$

- (1) Load the recorded $\text{pid}_{\mathcal{T}} \mapsto (\text{pk}_{\mathcal{T}}, \text{sk}_{\mathcal{T}}, \text{cert}_{\mathcal{T}}^{\mathcal{R}})$, and $\text{pid}_{\mathcal{R}} \mapsto (\text{pk}_{\mathcal{R}}, \text{sk}_{\mathcal{R}})$; if any of these do not exist, let \mathcal{F}_{P4TC} abort.
- (2) Generate $\text{cert}_{\mathcal{R}} := (\text{pk}_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}}, \sigma_{\mathcal{R}}^{\text{cert}})$ with $\sigma_{\mathcal{R}}^{\text{cert}} \leftarrow \text{S.Sgn}(\text{sk}_{\mathcal{T}}^{\text{cert}}, (\text{pk}_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}}))$ faithfully.
- (3) Update record $\text{pid}_{\mathcal{R}} \mapsto (\text{pk}_{\mathcal{R}}, \text{sk}_{\mathcal{R}}, \text{cert}_{\mathcal{R}})$.

Fig. 48. The simulator for Operator Security

Simulator $\mathcal{S}_{\pi_{P4TC}}^{\text{sys-sec}}$

Wallet Issuing:

- (1) Load the recorded $pid_{\mathcal{T}} \mapsto (\text{pk}_{\mathcal{T}}, \text{sk}_{\mathcal{T}}, \text{cert}_{\mathcal{T}}^{\mathcal{R}})$, $pid_{DR} \mapsto (\text{pk}_{DR}, \text{sk}_{DR})$; if any of these do not exist, let \mathcal{F}_{P4TC} abort.
- (2) Upon receiving $(\text{pk}_{\mathcal{U}}, c'_{\text{seed}})$ from $\mathcal{Z}^{\text{sys-sec}}$ in the name of \mathcal{U} with $pid_{\mathcal{U}}^*$...
 - (a) Look up $pid_{\mathcal{U}}$ from $\overline{\mathcal{G}}_{\text{bb}}$ for which $\text{pk}_{\mathcal{U}}$ has been recorded; if no $pid_{\mathcal{U}}$ exists abort.
 - (b) If $pid_{\mathcal{U}} \notin \mathcal{PID}_{\text{corrupt}}$ give up simulation.
 - (c) Call \mathcal{F}_{P4TC} with input (issue)
- (3) Upon receiving leaked $(s, \mathbf{a}_{\mathcal{U}})$ from \mathcal{F}_{P4TC} ...
 - (a) $(c''_{\text{ser}}, \overline{d}_{\text{ser}}) \leftarrow \text{C2.SimCom}(\text{CRS}_{\text{com}}^2)$.
 - (b) $\lambda'' \xleftarrow{\mathcal{R}} \mathbb{Z}_p$.
 - (c) Send $(\text{cert}_{\mathcal{T}}^{\mathcal{R}}, \mathbf{a}_{\mathcal{U}}, c''_{\text{ser}}, \lambda'')$ to $\mathcal{Z}^{\text{sys-sec}}$ as the 1st message from \mathcal{T} to \mathcal{U} .
- (4) Upon receiving $(s', e^*, c_{\mathcal{T}}, c_{\mathcal{R}}, \pi)$ from $\mathcal{Z}^{\text{sys-sec}}$ in the name of \mathcal{U} with $pid_{\mathcal{U}}^*$...
 - (a) $stmnt := (\text{pk}_{\mathcal{U}}, \text{pk}_{DR}, e^*, c_{\mathcal{T}}, c_{\mathcal{R}}, \pi)$.
 - (b) If $\text{P1.Vfy}(\text{CRS}_{\text{pok}}, stmnt, \pi) = 0$ let \mathcal{F}_{P4TC} abort.
 - (c) Extract $Wit = (\Lambda, \Lambda', \Lambda'_0, \dots, \Lambda'_{\ell-1}, \dots, U_1^{\text{next}}, d_{\mathcal{T}}, d_{\mathcal{R}}, d'_{\text{seed}}, \text{SK}_{\mathcal{U}}) \leftarrow \text{P1.ExtractW}(\text{CRS}, \text{td}_{\text{epok}}, stmnt, \pi)$.
 - (d) Assert that $(stmnt, Wit)$ fullfills the projected equations from $L_{\text{gp}}^{(1)}$, else abort (event $E1$)
 - (e) $\lambda := \lambda'' + \sum_{i=0}^{\ell-1} \text{DLOG}(\Lambda'_i) \cdot \mathcal{B}^i$
 - (f) Provide alternative user PID $pid_{\mathcal{U}}$ to \mathcal{F}_{P4TC} .
- (5) Upon being asked by \mathcal{F}_{P4TC} to provide ϕ ...
 - (a) $\phi := \text{PRF}(\lambda, x)$ with $x := 0$
 - (b) Provide ϕ to \mathcal{F}_{P4TC} .
- (6) Upon receiving output $(s, \mathbf{a}_{\mathcal{U}})$ from \mathcal{F}_{P4TC} for \mathcal{U} ...
 - (a) $s'' := s \cdot s'^{-1}$
 - (b) Equivoke $d''_{\text{ser}} \leftarrow \text{C2.Equiv}(\text{CRS}_{\text{com}}^2, \text{td}_{\text{eqcom}}, s'', c''_{\text{ser}}, \overline{d}_{\text{ser}})$.
 - (c) $\sigma_{\mathcal{T}} \leftarrow \text{S.Sgn}(\text{sk}_{\mathcal{T}}^{\mathcal{R}}, (c_{\mathcal{T}}, \mathbf{a}_{\mathcal{U}}))$
 - (d) $\sigma_{\mathcal{R}} \leftarrow \text{S.Sgn}(\text{sk}_{\mathcal{R}}^{\mathcal{R}}, (c_{\mathcal{R}}, s))$
 - (e) Set $s^{\text{prev}} := \perp$, $p := 0$, $b := 0$, $c_{\mathcal{T}}^{\text{in}} := \perp$, $d_{\mathcal{T}}^{\text{in}} := \perp$, $M_{\mathcal{T}}^{\text{in}} := \perp$, $c_{\mathcal{R}}^{\text{in}} := \perp$, $d_{\mathcal{R}}^{\text{in}} := \perp$, $M_{\mathcal{R}}^{\text{in}} := \perp$, $c_{\mathcal{T}}^{\text{out}} := c_{\mathcal{T}}$, $d_{\mathcal{T}}^{\text{out}} := d_{\mathcal{T}}$, $M_{\mathcal{T}}^{\text{out}} := (\Lambda, \text{pk}_{\mathcal{U}})$, $c_{\mathcal{R}}^{\text{out}} := c_{\mathcal{R}}$, $d_{\mathcal{R}}^{\text{out}} := d_{\mathcal{R}}$, and $M_{\mathcal{R}}^{\text{out}} := (\Lambda, g_1^b, U_1, g_1^{x+1})$.
 - (f) Append $(s^{\text{prev}}, s, \phi, x, \lambda, pid_{\mathcal{U}}, pid_{\mathcal{T}}, p, b, c_{\mathcal{T}}^{\text{in}}, d_{\mathcal{T}}^{\text{in}}, M_{\mathcal{T}}^{\text{in}}, c_{\mathcal{R}}^{\text{in}}, d_{\mathcal{R}}^{\text{in}}, M_{\mathcal{R}}^{\text{in}}, c_{\mathcal{T}}^{\text{out}}, d_{\mathcal{T}}^{\text{out}}, M_{\mathcal{T}}^{\text{out}}, c_{\mathcal{R}}^{\text{out}}, d_{\mathcal{R}}^{\text{out}}, M_{\mathcal{R}}^{\text{out}})$ to \overline{TRDB} .
 - (g) Send $(s'', d''_{\text{ser}}, \sigma_{\mathcal{T}}, \sigma_{\mathcal{R}})$ to $\mathcal{Z}^{\text{sys-sec}}$ as the 2nd message from \mathcal{T} to \mathcal{U} .

Fig. 49. The simulator for Operator Security (cont. from Fig. 48)

Simulator $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$ (cont.)

Debt Accumulation:

- (1) Load the recorded $pid_{\mathcal{T}} \mapsto (\text{pk}_{\mathcal{T}}, \text{sk}_{\mathcal{T}}, \text{cert}_{\mathcal{T}}^{\mathcal{R}})$, and $pid_{\mathcal{R}} \mapsto (\text{pk}_{\mathcal{R}}, \text{sk}_{\mathcal{R}}, \text{cert}_{\mathcal{R}})$; if any of these do not exist, let $\mathcal{F}_{\text{P4TC}}$ abort.
- (2) Pick $u_2 \xleftarrow{\mathcal{R}} \mathbb{Z}_{\mathfrak{p}}$.
- (3) $(c''_{\text{ser}}, \bar{d}_{\text{ser}}) \leftarrow \text{C2.SimCom}(\text{CRS}_{\text{com}}^2)$.
- (4) Send $(u_2, c''_{\text{ser}}, \text{cert}_{\mathcal{R}})$ to \mathcal{Z} as the 1st message from \mathcal{R} to \mathcal{U} .
- (5) Upon receiving $(s', \pi, \phi, \mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}}, c_{\text{hid}}, c'_{\mathcal{R}}, t)$ from \mathcal{Z} as the 2nd message from \mathcal{U} to \mathcal{R} ...
 - (a) $stmnt := (\text{pk}_{\mathcal{T}}^{\mathcal{T}}, \text{pk}_{\mathcal{T}}^{\text{cert}}, \phi, \mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}}, c_{\text{hid}}, c'_{\mathcal{R}}, t, u_2)$.
 - (b) If $\text{P2.Vfy}(\text{CRS}_{\text{pok}}, stmnt, \pi) = 0$ let $\mathcal{F}_{\text{P4TC}}$ abort.
 - (c) Extract $Wit = (X, \Lambda, \text{pk}_{\mathcal{U}}, U_1, s^{\text{prev}}, \phi^{\text{prev}}, X, \Lambda, \text{pk}_{\mathcal{U}}, B^{\text{prev}}, U_1, U_1^{\text{next}}, d_{\text{hid}}, d_{\mathcal{R}}^{\text{prev}}, d'_{\mathcal{R}}, d_{\mathcal{T}}, \text{pk}_{\mathcal{R}}^{\text{prev}}, c_{\mathcal{R}}^{\text{prev}}, c_{\mathcal{T}}, \sigma_{\mathcal{R}}^{\text{prev}}, \sigma_{\mathcal{R}}^{\text{cert}^{\text{prev}}}, \sigma_{\mathcal{T}}) \leftarrow \text{P2.ExtractW}(\text{CRS}, \text{td}_{\text{epok}}, stmnt, \pi)$.
 - (d) Assert that $(stmnt, Wit)$ fullfills the projected equations from $L_{\text{gp}}^{(2)}$, else abort (event $E1$)
 - (e) 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 $s^* = s^{\text{prev}}$ being used as key; if no unique entry exists, give up simulation (event $E2$).
 - (f) Give up simulation if any of these conditions meet: $c_{\mathcal{R}}^{\text{out}*} \neq c_{\mathcal{R}}^{\text{prev}}$ (event $E3$), $\Lambda \neq g_1^{\lambda^*}$ (event $E4$), $c_{\mathcal{T}}^{\text{out}*} \neq c_{\mathcal{T}}$ (event $E5$), $\text{pk}_{\mathcal{U}} \neq \text{pk}_{\mathcal{U}}^*$ with $(\Lambda^*, \text{pk}_{\mathcal{U}}^*) := M_{\mathcal{T}}^{\text{out}*}$ (event $E6$), $B^{\text{prev}} \neq g_1^{b^*}$ (event $E7$), or $X \neq g_1^{x^*+1}$ (event $E8$).
 - (g) Retrieve $pid_{\mathcal{U}}$ from $\overline{\mathcal{G}}_{\text{bb}}$ for $\text{pk}_{\mathcal{U}}$
 - (h) Call $\mathcal{F}_{\text{P4TC}}$ with input $(\text{pay_to11}, s^{\text{prev}})$ in the name of \mathcal{U}
- (6) Upon being asked by $\mathcal{F}_{\text{P4TC}}$ to provide an alternative PID, return $pid_{\mathcal{U}}$ to $\mathcal{F}_{\text{P4TC}}$.
- (7) If being asked by $\mathcal{F}_{\text{P4TC}}$ to provide ϕ , return $\phi := \text{PRF}(\lambda, x)$ with $x := x^* + 1$ to $\mathcal{F}_{\text{P4TC}}$.^a
- (8) Upon being ask by $\mathcal{F}_{\text{P4TC}}$ 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.
- (9) Upon receiving output $(s, \mathbf{a}_{\mathcal{R}}, p, b)$ from $\mathcal{F}_{\text{P4TC}}$ for \mathcal{U} ...
 - (a) Set $\overline{\omega}_{\mathcal{U}}^{\text{pp}} := (s, c_{\text{hid}}, d_{\text{hid}}, \text{pk}_{\mathcal{U}})$ and append $\overline{\omega}_{\mathcal{U}}^{\text{pp}}$ to $\overline{\Omega}_{\mathcal{U}}^{\text{pp}}$.
 - (b) 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 $\sigma_{\mathcal{R}} \leftarrow \text{S.Sgn}(\text{sk}_{\mathcal{R}}, (c_{\mathcal{R}}, s))$ honestly as the real protocol would do.
 - (c) Set $s'' := s \cdot s'^{-1}$ and equivocate $d''_{\text{ser}} \leftarrow \text{C2.Equiv}(\text{CRS}_{\text{com}}^2, s'', c''_{\text{ser}}, \bar{d}_{\text{ser}})$.
 - (d) Set $c_{\mathcal{T}}^{\text{in}} := c_{\mathcal{T}}$, $d_{\mathcal{T}}^{\text{in}} := d_{\mathcal{T}}$, $M_{\mathcal{T}}^{\text{in}} := (\Lambda, \text{pk}_{\mathcal{U}})$, $c_{\mathcal{R}}^{\text{in}} := c_{\mathcal{R}}^{\text{prev}}$, $d_{\mathcal{R}}^{\text{in}} := d_{\mathcal{R}}^{\text{prev}}$, $M_{\mathcal{R}}^{\text{in}} := (\Lambda, B^{\text{prev}}, U_1, X)$, $c_{\mathcal{T}}^{\text{out}} := c_{\mathcal{T}}$, $d_{\mathcal{T}}^{\text{out}} := d_{\mathcal{T}} \cdot d''_{\mathcal{T}}$, $M_{\mathcal{T}}^{\text{out}} := (\Lambda, \text{pk}_{\mathcal{U}})$, $c_{\mathcal{R}}^{\text{out}} := c_{\mathcal{R}}$, $d_{\mathcal{R}}^{\text{out}} := d_{\mathcal{R}} \cdot d''_{\mathcal{R}}$, and $M_{\mathcal{R}}^{\text{out}} := (\Lambda, g_1^b, U_1, g_1^{x+1})$.
 - (e) 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} .
 - (f) Append $\overline{\omega}^{\text{dsp}} = (\phi, t, u_2)$ to $\overline{\Omega}^{\text{dsp}}$
 - (g) 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) $\text{sk}_{\mathcal{U}} := (t - t^{\ddagger}) \cdot (u_2 - u_2^{\ddagger})^{-1} \bmod \mathfrak{p}$
 - (ii) Record $pid_{\mathcal{U}} \mapsto (\text{pk}_{\mathcal{U}}, \text{sk}_{\mathcal{U}})$ internally
 - (h) Send $(s'', d''_{\text{ser}}, c_{\mathcal{R}}, d''_{\mathcal{R}}, \sigma_{\mathcal{R}}, p)$ to \mathcal{Z} as the 3rd message from \mathcal{R} to \mathcal{U} .

^a N.b.: $\mathcal{F}_{\text{P4TC}}$ does not always ask for the next serial number. If the corrupted user re-uses an old token, then $\mathcal{F}_{\text{P4TC}}$ 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. 50. The simulator for Operator Security (cont. from Fig. 48)

Simulator $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$ (cont.)

Debt Clearance:

- (1) Load the recorded $pid_{\mathcal{T}} \mapsto (\text{pk}_{\mathcal{T}}, \text{sk}_{\mathcal{T}}, \text{cert}_{\mathcal{T}}^{\mathcal{R}})$ if this does not exist, let $\mathcal{F}_{\text{P4TC}}$ abort.
- (2) Pick $u_2 \xleftarrow{\mathcal{R}} \mathbb{Z}_p$.
- (3) Send u_2 to \mathcal{Z} as the 1st message from \mathcal{T} to \mathcal{U} .
- (4) Upon receiving $(\text{pk}_{\mathcal{U}}, \pi, \phi, \mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}}, b^{\text{prev}}, t)$ from \mathcal{Z} as the 2nd message from \mathcal{U} to \mathcal{T} ...
 - (a) $stmnt := (\text{pk}_{\mathcal{U}}, \text{pk}_{\mathcal{T}}^{\mathcal{T}}, \text{pk}_{\mathcal{T}}^{\text{cert}}, \phi, \mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}}, g_1^{b^{\text{prev}}}, t, u_2)$.
 - (b) If $\text{P3.Vfy}(\text{CRS}_{\text{pok}}, stmnt, \pi) = 0$ let $\mathcal{F}_{\text{P4TC}}$ abort.
 - (c) Extract $Wit = (X, \Lambda, \text{pk}_{\mathcal{U}}, U_1, s^{\text{prev}}, \phi^{\text{prev}}, X, \Lambda, U_1, d_{\mathcal{R}}^{\text{prev}}, d_{\mathcal{T}}, \text{pk}_{\mathcal{R}}^{\text{prev}}, c_{\mathcal{R}}^{\text{prev}}, c_{\mathcal{T}}, \sigma_{\mathcal{R}}^{\text{prev}}, \sigma_{\mathcal{R}}^{\text{cert}^{\text{prev}}}, \sigma_{\mathcal{T}}) \leftarrow \text{P3.ExtractW}(\text{CRS}, \text{td}_{\text{epok}}, stmnt, \pi)$.
 - (d) Assert that $(stmnt, Wit)$ fullfills the projected equations from $L_{\text{gp}}^{(3)}$, else abort (event $E1$)
 - (e) 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 $s^* = s^{\text{prev}}$ being used as key; if no unique entry exists, give up simulation (event $E2$).
 - (f) Give up simulation if any of these conditions meet: $c_{\mathcal{R}}^{\text{out}*} \neq c_{\mathcal{R}}^{\text{prev}}$ (event $E3$), $\Lambda \neq g_1^{\lambda^*}$ (event $E4$), $c_{\mathcal{T}}^{\text{out}*} \neq c_{\mathcal{T}}$ (event $E5$), $\text{pk}_{\mathcal{U}} \neq \text{pk}_{\mathcal{U}}^*$ with $(\Lambda^*, \text{pk}_{\mathcal{U}}^*) := M_{\mathcal{T}}^{\text{out}*}$ (event $E6$), or $B^{\text{prev}} \neq g_1^{b^*}$ (event $E7$).
 - (g) Retrieve $pid_{\mathcal{U}}$ from $\overline{\mathcal{G}}_{\text{bb}}$ for $\text{pk}_{\mathcal{U}}$
 - (h) Call $\mathcal{F}_{\text{P4TC}}$ with input $(\text{clear_debt}, s^{\text{prev}})$ in the name of \mathcal{U}
- (5) Upon being asked by $\mathcal{F}_{\text{P4TC}}$ to provide an alternative PID, return $pid_{\mathcal{U}}$ to $\mathcal{F}_{\text{P4TC}}$.
- (6) If being asked by $\mathcal{F}_{\text{P4TC}}$ to provide ϕ , return $\phi := \text{PRF}(\lambda, x)$ with $x := x^{\text{prev}} + 1$ to $\mathcal{F}_{\text{P4TC}}$.^a
- (7) Upon receiving output (b^{bill}) from $\mathcal{F}_{\text{P4TC}}$ for \mathcal{U} ...
 - (a) Set $c_{\mathcal{T}}^{\text{in}} := c_{\mathcal{T}}, d_{\mathcal{T}}^{\text{in}} := d_{\mathcal{T}}, M_{\mathcal{T}}^{\text{in}} := (\Lambda, \text{pk}_{\mathcal{U}}), c_{\mathcal{R}}^{\text{in}} := c_{\mathcal{R}}^{\text{prev}}, d_{\mathcal{R}}^{\text{in}} := d_{\mathcal{R}}^{\text{prev}}, M_{\mathcal{R}}^{\text{in}} := (\Lambda, g_1^{b^{\text{prev}}}, U_1, X), c_{\mathcal{T}}^{\text{out}} := \perp, d_{\mathcal{T}}^{\text{out}} := \perp, M_{\mathcal{T}}^{\text{out}} := \perp, c_{\mathcal{R}}^{\text{out}} := \perp, d_{\mathcal{R}}^{\text{out}} := \perp, \text{and } M_{\mathcal{R}}^{\text{out}} := \perp$.
 - (b) 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} .
 - (c) Append $\overline{w}^{\text{dsp}} = (\phi, t, u_2)$ to $\overline{\Omega}^{\text{dsp}}$
 - (d) Check if $\overline{w}^{\text{dsp}\ddagger} = (\phi^{\ddagger}, t^{\ddagger}, u_2^{\ddagger}) \in \overline{\Omega}^{\text{dsp}}$ exists with $\overline{\phi} = \overline{\phi}^{\ddagger}$ and $u_2 \neq u_2^{\ddagger}$; in this case
 - (i) $\text{sk}_{\mathcal{U}} := (t - t^{\ddagger}) \cdot (u_2 - u_2^{\ddagger})^{-1} \bmod p$
 - (ii) Record $pid_{\mathcal{U}} \mapsto (\text{pk}_{\mathcal{U}}, \text{sk}_{\mathcal{U}})$ internally
 - (e) Send (OK) to \mathcal{Z} as the 3rd message from \mathcal{T} to \mathcal{U} .

^a N.b.: $\mathcal{F}_{\text{P4TC}}$ does not always ask for the next serial number. If the corrupted user re-uses an old token, then $\mathcal{F}_{\text{P4TC}}$ 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. 51. The simulator for Operator Security (cont. from Fig. 48)

Simulator $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$ (cont.)

Prove Participation:

- (1) Call $\mathcal{F}_{\text{P4TC}}$ with input (prove_participation) in the name of \mathcal{U}
- (2) Obtain leaked set $S_{\mathcal{R}}^{\text{PP}}$ of serial numbers from $\mathcal{F}_{\text{P4TC}}$.
- (3) Upon receiving output ($\text{out}_{\mathcal{U}}$) from $\mathcal{F}_{\text{P4TC}}$ for $\mathcal{U} \dots$
 - (a) Delay the output of $\mathcal{F}_{\text{P4TC}}$ to SA .
 - (b) Send $S_{\mathcal{R}}^{\text{PP}}$ to \mathcal{Z} as the 1st message from SA to \mathcal{U} .
- (4) Upon receiving ($s, c_{\text{hid}}, d_{\text{hid}}$) from \mathcal{Z} as the 2nd message from \mathcal{U} to $SA \dots$
 - (a) If $(s, c_{\text{hid}}, \cdot, \cdot) \notin \overline{\mathcal{D}}_{\mathcal{U}}^{\text{PP}}$, let $\mathcal{F}_{\text{P4TC}}$ abort.
 - (b) Look up $\text{pk}_{\mathcal{U}}$ from $\overline{\mathcal{G}}_{\text{bb}}$ for PID $\text{pid}_{\mathcal{U}}$; if no $\text{pk}_{\mathcal{U}}$ has been recorded abort.
 - (c) If $\text{C1.Open}(\text{CRS}_{\text{com}}^1, \text{pk}_{\mathcal{U}}, c_{\text{hid}}, d_{\text{hid}}) = 0$, let $\mathcal{F}_{\text{P4TC}}$ abort.
 - (d) If $\text{out}_{\mathcal{U}} = \text{NOK}$ give up simulation.
 - (e) Let $\mathcal{F}_{\text{P4TC}}$ return the delayed output to SA .

Double-Spending Detection: Upon being ask to provide a proof for $\text{pid}_{\mathcal{U}}$, look up $\text{pid}_{\mathcal{U}} \mapsto (\text{pk}_{\mathcal{U}}, \text{sk}_{\mathcal{U}})$, and return $\text{sk}_{\mathcal{U}}$.

Guilt Verification: Upon being ask by $\mathcal{F}_{\text{P4TC}}$ to provide out for $(\text{pid}_{\mathcal{U}}, \pi) \dots$

- (1) Receive $\text{pk}_{\mathcal{U}}$ from $\overline{\mathcal{G}}_{\text{bb}}$ for $\text{pid}_{\mathcal{U}}$.
- (2) If $g_1^\pi = \text{pk}_{\mathcal{U}}$, then return $\text{out} := \text{OK}$, else $\text{out} := \text{NOK}$ to $\mathcal{F}_{\text{P4TC}}$.

User Blacklisting: Upon receiving (λ, x) and being ask to provide a serial number, return $\phi := \text{PRF}(\lambda, x)$.

Fig. 52. The simulator for Operator Security (cont. from Fig. 48)




Fig. 53. An entry $\overline{trdb} \in \overline{TRDB}$ visualized as an element of a directed graph

Entry \overline{trdb} has the form

$$\overline{trdb} = (s^{\text{prev}}, s, \phi, x, \lambda, \text{pid}_{\mathcal{U}}, \text{pid}_{\mathcal{R}}, p, b, \quad (8)$$

$$c_{\mathcal{T}}^{\text{in}}, d_{\mathcal{T}}^{\text{in}}, M_{\mathcal{T}}^{\text{in}}, c_{\mathcal{R}}^{\text{in}}, d_{\mathcal{R}}^{\text{in}}, M_{\mathcal{R}}^{\text{in}},$$

$$c_{\mathcal{T}}^{\text{out}}, d_{\mathcal{T}}^{\text{out}}, M_{\mathcal{T}}^{\text{out}}, c_{\mathcal{R}}^{\text{out}}, d_{\mathcal{R}}^{\text{out}}, M_{\mathcal{R}}^{\text{out}})$$

with c, d and M with equal suffixes denoting a commitment, its decommitment information and the opening in the implicit message space (see Fig. 53). At the beginning of a transaction in the scope of Debt Accumulation or Debt Clearance the user loads his token τ^{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 τ again. We call the initial commitments the *in*-commitments of the transaction and the resulting commitments the *out*-commitments.

Definition E.11 (Simulated Transaction Graph (informal)).

The set $\overline{TRDB} = \{\overline{trdb}_i\}$ with \overline{trdb}_i defined as in Eq. (8) is called the *Simulated 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 nor 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 instead of re-using the term "previous commitment". Actually, these kind of equalities is what we have to show.

The overall proof idea is to define a sequence of hybrid experiments H_i together with simulators \mathcal{S}_i and protocols π_i such that the first hybrid H_0 is identical to

the real experiment and the last hybrid H_{16} is identical to the ideal experiment. Each hybrid is of the form

$$H_i := \text{EXEC}_{\pi_i, \bar{G}_{\text{bb}}, \mathcal{S}_i, \mathcal{Z}^{\text{sys-sec}}}(1^n).$$

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 $\mathcal{Z}^{\text{sys-sec}}$ 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 π_i that honest parties perform gradually declines from the real protocol $\pi_0 = \pi_{\text{P4TC}}$ to a dummy protocol π_{16} , 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}_{16} which can be split up into the ideal functionality $\mathcal{F}_{\text{P4TC}}$ and $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$.

We proceed by giving concrete (incremental) definitions of all hybrids H_i . Please note that *input privacy* for the honest TSP, RSU, DR and SA does not pose a difficulty for the definition of the sequence of the simulators. The users learns most information as part of its prescribed output anyway. In other words, the simulator $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{sys-sec}}$ mimicking the role of an honest operator can perfectly simulate most messages towards the (malicious) user after it has received the user's output from the ideal functionality. The essential part is to ensure that no malicious user can make the Simulated Transaction Graph to deviate from the Ideal Transaction Graph and thereby cause a different (wrong) output at some later point in the protocol. To this end, most hybrids introduce additional “sanity checks” to the simulation: if the sanity check holds, both transaction graphs are still in sync and the simulator proceeds; if the sanity check fails, the adversary has caused the transaction graphs to fall apart and the simulator immediately gives up the simulation. Each sanity check is related to the security of one of the building blocks or cryptographic assumptions. Finally, after the last hybrid all sanity checks collectively assert that no efficient adversary can deviate from the Ideal Transaction Graph.

Hybrid H_0

The hybrid H_0 is defined as

$$H_0 := \text{EXEC}_{\pi_0, \bar{G}_{\text{bb}}, \mathcal{S}_0, \mathcal{Z}^{\text{sys-sec}}}(1^n)$$

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 CRS_{pok} is generated by SetupEPoK , and $\text{CRS}_{\text{com}}^2$ is generated by C2.SimGen . Additionally, \mathcal{S}_1 initializes the internal sets \overline{TRDB} , $\overline{\Omega}^{\text{dsp}}$, and $\overline{\Omega}_{\mathcal{U}}^{\text{pp}}$ as empty sets.

Hybrid H_2

Hybrid H_2 replaces the code in the tasks DR/TSP/RSU/User Registration of the protocol π_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. The protocol π_3 is modified such that the simulator \mathcal{S}_3 receives the message $(\text{certifying_rsu}, \text{pid}_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}})$, creates the certificate $\sigma_{\mathcal{R}}^{\text{cert}}$ and records it.

Whenever the honest TSP or honest RSU running π_3 would send $\sigma_{\mathcal{R}}^{\text{cert}}$ as part of its messages in the scope of Wallet Issuing or Debt Accumulation, they omit $\sigma_{\mathcal{R}}^{\text{cert}}$. Instead, the simulator \mathcal{S}_3 injects $\sigma_{\mathcal{R}}^{\text{cert}}$ into the message.

Hybrid H_4

Hybrid H_4 replaces the code in the tasks Wallet Issuing and Debt Accumulation of the protocol π_4 such that the RSU/TSP do not create signatures, but the simulator \mathcal{S}_4 creates the signatures $\sigma_{\mathcal{T}}$, $\sigma_{\mathcal{R}}$ and $\sigma_{\mathcal{R}}$ resp. and injects them into the messages instead.

Moreover, in Debt Accumulation the RSU running π_4 does not send $c_{\mathcal{R}}$ and $d''_{\mathcal{R}}$ in its final message, but reports the price p to \mathcal{S}_4 , \mathcal{S}_4 creates $c_{\mathcal{R}}$ and $d''_{\mathcal{R}}$ honestly and injects them into the message.

Hybrid H_5

H_5 modifies the tasks of Wallet Issuing and Debt Accumulation. The code of π_5 for the TSP/RSU is modified such that it does not send c''_{ser} in the scope of Wallet Issuing or Debt Accumulation. Instead \mathcal{S}_5 runs $(c''_{\text{ser}}, \bar{d}_{\text{ser}}) \leftarrow \text{C2.SimCom}(\text{CRS}_{\text{com}}^2)$ and injects c''_{ser} into the message. Moreover, π_5 for the TSP/RSU is modified such that it uniformly and independently picks $s \xleftarrow{\mathcal{R}} \mathbb{Z}_p$ and passes s to \mathcal{S}_5 as part of the final message. \mathcal{S}_5 calculates $s'' := s \cdot (s')^{-1}$, executes $d''_{\text{ser}} \leftarrow \text{C2.Equiv}(\text{CRS}_{\text{com}}^2, \text{td}_{\text{eqcom}}, s'', c''_{\text{ser}}, \bar{d}_{\text{ser}})$ and injects s'' together with d''_{ser} into the messages from TSP/RSU to the user.

Hybrid H_6

When \mathcal{S}_6 receives a NIZK proof π in the scope of Wallet Issuing, Debt Accumulation and Debt Clearance, it ex-

tracts the witness, restores $\lambda := \lambda'' + \sum_{i=0}^{\ell-1} \text{DLOG}(\Lambda_i') \cdot \mathcal{B}^i$, assembles \overline{trdb} and appends it to \overline{TRDB} . Additionally, \mathcal{S}_6 also assembles $\overline{\omega}_{\mathcal{U}}^{\text{pp}}$, $\overline{\omega}^{\text{dsp}}$ in the scope of Debt Accumulation and appends entries to $\overline{\Omega}_{\mathcal{U}}^{\text{pp}}$, $\overline{\Omega}^{\text{dsp}}$ resp. If $\overline{\Omega}^{\text{dsp}}$ already contains an entry $\overline{\omega}^{\text{dsp}\ddagger}$ with matching fraud detection ID, the secret key $\text{sk}_{\mathcal{U}}$ is immediately reconstructed and the pair $\text{pid}_{\mathcal{U}} \mapsto (\text{pk}_{\mathcal{U}}, \text{sk}_{\mathcal{U}})$ is also recorded.

Moreover, the verification of the proof is moved from π_6 for the honest TSP/RSU to the simulator. If the verification fails, \mathcal{S}_6 aborts as the TSP/RSU running the real protocol would do.

Additionally, \mathcal{S}_6 checks if the pair of the statement and the extracted witness fulfills the languages $L_{\text{gp}}^{(1)}$, $L_{\text{gp}}^{(2)}$, and $L_{\text{gp}}^{(3)}$ resp. If not, \mathcal{S}_6 abort with failure event ($E1$).

Hybrid H₇

This hybrid modifies the code π_7 for \mathcal{T} , SA and DR in the scope of the tasks Prove Participation, Double-Spending Detection, Guilt Verification and User Blacklisting. The honest parties become dummy parties, the code is moved to the simulator and \mathcal{S}_7 resorts to its “lookup tables” \overline{TRDB} , $\overline{\Omega}_{\mathcal{U}}^{\text{pp}}$, and $\overline{\Omega}^{\text{dsp}}$ that have been introduced by the previous hybrid.

More precisely, in Prove Participation the party SA becomes a dummy party and simply forwards the set of serial numbers $S_{\mathcal{R}}^{\text{pp}}$ to \mathcal{S}_7 . The simulator uses its own set $\overline{\Omega}_{\mathcal{U}}^{\text{pp}}$ to validate the response of the environment (in the name of the malicious user) and returns the result to SA .

In the task Double-Spending Detection the honest \mathcal{T} becomes a dummy party, too. It simply asks the simulator \mathcal{S}_7 to provide a proof. To this end, \mathcal{S}_7 checks if $\text{pid}_{\mathcal{U}} \mapsto (\text{pk}_{\mathcal{U}}, \text{sk}_{\mathcal{U}})$ has been recorded and returns $\text{sk}_{\mathcal{U}}$.

The same applies to the task Guilt Verification. The honest party does not locally run the algorithm itself, but simply forwards its input to the simulator (as the dummy party would do) and \mathcal{S}_7 actually checks if $g_1^\pi = \pi$ holds.

The task User Blacklisting is modified accordingly. The dispute resolver DR becomes a dummy party and simply sends its input ($\text{blacklist_user}, \text{pk}_{\mathcal{U}}^{\text{DR}}$) to the simulator \mathcal{S}_7 in order to signal its consent to blacklist the user. The simulator \mathcal{S}_7 utilizes the Simulated Transaction Graph \overline{TRDB} and runs the code as the ideal functionality $\mathcal{F}_{\text{P4TC}}$ would do eventually.

Hybrid H₈

Hybrid H₈ replaces the code in the tasks Wallet Issuing and Debt Accumulation of the protocol π_8 such that the

RSU/TSP do not neither send λ'' nor u_2 . Instead \mathcal{S}_8 draws λ'' and u_2 and injects them into the appropriate messages. Consequently, the code of the TSP is modified such that it does not longer record htd . Likewise, the code of the RSU is modified such that it does not longer record ω^{dsp} .

Hybrid H₉

In the scope of Debt Accumulation or Debt Clearance, the simulator \mathcal{S}_9 looks up the predecessor entry with s^{prev} being used as the unique key. If this fails, \mathcal{S}_9 gives up the simulation with event $E2$.

Hybrid H₁₀

The simulator \mathcal{S}_{10} additionally checks for $c_{\mathcal{R}}^{\text{out}*} \neq c_{\mathcal{R}}^{\text{prev}}$ and gives up the simulation with event $E3$, if the check succeeds.

Hybrid H₁₁

The simulator \mathcal{S}_{11} additionally checks for $\Lambda \neq g_1^{\lambda^*}$ and gives up the simulation with event $E4$, if the check succeeds.

Hybrid H₁₂

The simulator \mathcal{S}_{12} additionally checks for $c_{\mathcal{T}}^{\text{out}*} \neq c_{\mathcal{T}}$ and gives up the simulation with event $E5$, if the check succeeds.

Hybrid H₁₃

The simulator \mathcal{S}_{13} parses $(\Lambda^*, \text{pk}_{\mathcal{U}}^*) := M_{\mathcal{T}}^{\text{out}*}$ and checks for $\text{pk}_{\mathcal{U}} \neq \text{pk}_{\mathcal{U}}^*$. If the check succeeds, it gives up the simulation with event $E6$.

Hybrid H₁₄

The simulator \mathcal{S}_{14} additionally checks for $B^{\text{prev}} \neq g_1^{b^*}$ and gives up the simulation with event $E7$, if the check succeeds.

Hybrid H₁₅

The simulator \mathcal{S}_{15} additionally checks for $X \neq g_1^{x^*+1}$ and gives up the simulation with event $E8$, if the check succeeds.

For the proof of [Theorem E.10](#) we show the indistinguishability of subsequent hybrids by a series of lemmas. The [Lemmas E.12](#) to [E.14](#) are rather trivial and thus [Lemma E.13](#) handles various hybrids at once.

Lemma E.12 (Indistinguishability between H₀ and H₁).

Under the assumptions of [Theorem E.10](#), $H_0 \stackrel{c}{\equiv} H_1$ holds.

Proof. This hop solely changes how the CRS is created during the setup phase. This is indistinguishable for CRS_{pok} , and $\text{CRS}_{\text{com}}^2$ (see the extractability property of [Definition C.5](#) and the equivocality property of [Definition C.7](#), resp., condition (a) each). \square

Lemma E.13 (Indistinguishability $H_1 \rightarrow H_4$ and $H_6 \rightarrow H_8$). *Under the assumptions of [Theorem E.10](#), $H_1 \stackrel{c}{\equiv} H_2$, $H_2 \stackrel{c}{\equiv} H_3$, $H_3 \stackrel{c}{\equiv} H_4$, $H_6 \stackrel{c}{\equiv} H_7$, and $H_7 \stackrel{c}{\equiv} H_8$ holds.*

Proof. The hops are all indistinguishable as they do not change anything in the view of $\mathcal{Z}^{\text{sys-sec}}$. Please note, that $\mathcal{Z}^{\text{sys-sec}}$ only sees the in-/output of honest parties and these hops only syntactically change what parts of the code are executed by the parties or by the simulator. With each hop the parties degrade more to a dummy party while at the same time more functionality is put into the simulator. \square

Lemma E.14 (Indistinguishability between H_4 and H_5). *Under the assumptions of [Theorem E.10](#), $H_4 \stackrel{c}{\equiv} H_5$ holds.*

Proof. This hop is indistinguishable as the equivocated decommitment information is perfectly indistinguishable from a decommitment that has originally been created with the correct message (cp. [Definition C.7, Item 3](#)). \square

So far, none of the hops between two consecutive hybrids changes anything from the environment's perspective: either the hops are only syntactical or the modification is perfectly indistinguishable. Hence, no reduction argument is required. In the contrary, each of the upcoming security proofs roughly follows the same lines of argument. If the environment $\mathcal{Z}^{\text{sys-sec}}$ can efficiently distinguish between two consecutive hybrids, then we can construct an efficient adversary \mathcal{B} against one of the underlying cryptographic building blocks. 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 internally, it only has *black-box access* to it. In other words, although everything happens inside “the head of \mathcal{B} ” it cannot somehow magically extract \mathcal{Z} 's attack strategy.

Lemma E.15 (Indistinguishability between H_5 and H_6). *Under the assumptions of [Theorem E.10](#), $H_5 \stackrel{c}{\equiv} H_6$ holds.*

Proof. First note that the only effective change between H_5 and H_6 are the additional checks that abort the sim-

ulation with event $E1$, if the extracted witnesses are invalid. Again, the other modifications are purely syntactical. To prove indistinguishability between H_5 and H_6 we split this hop into three sub-hybrids. Each sub-hybrid introduces the check for one of the languages $L_{\text{gp}}^{(1)}$, $L_{\text{gp}}^{(2)}$ and $L_{\text{gp}}^{(3)}$, resp. In the following only the sub-hybrid for the language $L_{\text{gp}}^{(1)}$ is considered, the indistinguishability of the remaining two is proved analogously. Further note, that the view of $\mathcal{Z}^{\text{sys-sec}}$ is perfectly indistinguishable, if the simulation does not abort.

Assume there is an environment $\mathcal{Z}^{\text{sys-sec}}$ that triggers the event $E1$ in the first sub-hybrid with non-negligible advantage. This immediately yields an efficient adversary \mathcal{B} against the extraction property of the NIZK scheme. Internally, \mathcal{B} runs $\mathcal{Z}^{\text{sys-sec}}$ in its head, plays the role of the simulator and all honest parties. Externally, \mathcal{B} plays the adversary in [Definition C.5, Item 3b](#). If the event $E1$ occurs internally, \mathcal{B} outputs the corresponding pair (stmt, π) . In the second and third sub-hybrid \mathcal{B} internally extracts the witness for the previous sub-hybrid using the extraction trapdoor td_{epok} which \mathcal{B} obtains as part of its input. \square

Remark E.16. We observe that [Lemma E.15](#) implies with $m = (\Lambda, \text{pk}_{\mathcal{U}})$:

$$\text{C1.Open}(\text{CRS}_{\text{com}}^1, m, c_{\mathcal{T}}, d_{\mathcal{T}}) = 1,$$

with $m = (\Lambda, 1, U_1^{\text{next}}, g_1)$:

$$\text{C1.Open}(\text{CRS}_{\text{com}}^1, m, c_{\mathcal{R}}, d_{\mathcal{R}}) = 1,$$

with $m = (\Lambda, B^{\text{prev}}, U_1, X)$:

$$\text{C1.Open}(\text{CRS}_{\text{com}}^1, m, c_{\mathcal{R}}^{\text{prev}}, d_{\mathcal{R}}^{\text{prev}}) = 1,$$

with $m = (\Lambda, B^{\text{prev}}, U_1^{\text{next}}, X)$:

$$\text{C1.Open}(\text{CRS}_{\text{com}}^1, m, c'_{\mathcal{R}}, d'_{\mathcal{R}}) = 1,$$

furthermore

$$\text{C1.Open}(\text{CRS}_{\text{com}}^1, \Lambda', c'_{\text{seed}}, d'_{\text{seed}}) = 1,$$

$$\text{C1.Open}(\text{CRS}_{\text{com}}^1, \text{pk}_{\mathcal{U}}, c_{\text{hid}}, d_{\text{hid}}) = 1,$$

with $m = (c_{\mathcal{T}}, \mathbf{a}_{\mathcal{U}})$:

$$\text{S.Vfy}(\text{pk}_{\mathcal{T}}^{\mathcal{T}}, \sigma_{\mathcal{T}}, m) = 1,$$

with $m = (c_{\mathcal{R}}^{\text{prev}}, s^{\text{prev}})$:

$$\text{S.Vfy}(\text{pk}_{\mathcal{R}}^{\text{prev}}, \sigma_{\mathcal{R}}^{\text{prev}}, m) = 1,$$

with $m = (\text{pk}_{\mathcal{R}}^{\text{prev}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}})$:

$$\text{S.Vfy}(\text{pk}_{\mathcal{T}}^{\text{cert}}, \sigma_{\mathcal{R}}^{\text{cert}^{\text{prev}}}, m) = 1,$$

and that all variables can efficiently be extracted. Remember, that F_{gp} acts as the identity function on group elements. Moreover, given the extracted chunks of the Wallet ID $\Lambda'_0, \dots, \Lambda'_{\ell-1}$ the unique Wallet ID λ can be reconstructed. The projection F_{gp} becomes injective if the pre-image is restricted to \mathbb{Z}_p and the inverse, i.e. DLOG, can be efficiently computed as $\lambda'_0, \dots, \lambda'_{\ell-1}$ are sufficiently “small”.

Up to this point, we already know that $H_0 \stackrel{c}{\equiv} H_8$ holds. Except for two small changes (from H_4 to H_5 and from H_5 to H_6) all hops are only syntactical. Moreover, the simulator \mathcal{S}_8 of hybrid H_8 is indeed sufficient to simulate an indistinguishable view for $\mathcal{Z}^{\text{sys-sec}}$ in the ideal model. Note, that all subsequent hybrids from H_{10} to H_{15} only add more sanity checks but do not change any messages. Actually, even the modification introduced by H_6 is not required for a indistinguishable simulation, as H_6 only records \overline{TRDB} , but \overline{TRDB} is not used yet. However, only \overline{TRDB} and the upcoming sanity checks enable a reduction to cryptographic assumptions and thus are vital to proof the indistinguishability between H_8 and the ideal model.

To this end, two additional lemmas about the structure of \overline{TRDB} are necessary. These lemmas are in the same spirit as [Lemmas E.3](#) and [E.4](#). Intuitively, the commitments $c_{\mathcal{T}}, c_{\mathcal{R}}$ induce a graph structure onto \overline{TRDB} comparable to the wallet ID λ and serial number s .

Lemma E.17 (Simulated Transaction Forest).

1. Every $\overline{trdb} = (s^{\text{prev}}, s, \dots) \in \overline{TRDB}$ is uniquely identified by s with overwhelming probability.
2. The Simulated Transaction Graph \overline{TRDB} is a forest with edges defined by (s^{prev}, s) .

Proof. We prove both claims separately.

1. A new entry is only inserted in the scope of Wallet Issuing, Debt Accumulation or Debt Clearance. Proof by Induction: The statement is correct for the empty \overline{TRDB} . For each insertion, the simulator \mathcal{S}_6 (and every following simulator) draws s uniformly and independently. The chance to pick a serial number that has already been used is negligible.
2. As the serial number s of the new node is randomly chosen, no existing node can point to the new node as its predecessor and thus no cycle is closed with overwhelming probability.

Lemma E.18 (Indistinguishability between H_8 and H_9).

Under the assumptions of [Theorem E.10](#), $H_8 \stackrel{c}{\equiv} H_9$ holds.

Proof. Assume there is an environment $\mathcal{Z}^{\text{sys-sec}}$ that trigger the event $E2$ with non-negligible advantage. This immediately yields an efficient adversary \mathcal{B} against the EUF-CMA security of S . We only need to deal with the case that s^* does not exist. If it exists, [Lemma E.17](#), [Item 1](#) implies its uniqueness. We need to distinguish two cases. On an abstract level these cases correspond to the following scenarios: Either the previous RSU exists. Then the signature $\sigma_{\mathcal{R}}^{\text{prev}}$ on $(c_{\mathcal{R}}^{\text{prev}}, s^{\text{prev}})$ is a forgery. Or alternatively, the allegedly previous RSU does not exist but has been imagined by the user. Then $(c_{\mathcal{R}}^{\text{prev}}, s^{\text{prev}})$ may have a honest, valid signature (because the user feigned the RSU), but the certificate $\text{cert}_{\mathcal{R}}^{\text{prev}}$ for the fake RSU constitutes a forgery. Please note, that the simulator always records an entry \overline{trdb} when it legitimately issues a signature $\sigma_{\mathcal{R}}$ and vice versa.

1. A record $\text{pid}_{\mathcal{R}}^{\text{prev}} \mapsto (\text{pk}_{\mathcal{R}}^{\text{prev}}, \text{sk}_{\mathcal{R}}^{\text{prev}})$ has been recorded: In other words, $(c_{\mathcal{R}}^{\text{prev}}, s^{\text{prev}})$ has never been legitimately issued by a the the allegedly previous RSU.³² 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 plays the role of the simulator and all honest parties. Externally, \mathcal{B} plays the EUF-CMA security experiment with a challenger \mathcal{C} and a signing oracle $\text{OS}_{\text{pk}, \text{sk}}^S$. \mathcal{B} needs to guess for which $\text{pid}_{\mathcal{R}}^{\text{prev}}$ the event ($E2$) eventually occurs. When the RSU with $\text{pid}_{\mathcal{R}}^{\text{prev}}$ registers itself, and \mathcal{B} playing \mathcal{S}_9 needs to provide $\text{pk}_{\mathcal{R}}^{\text{prev}}$ it embeds the challenge as $\text{pk}_{\mathcal{R}}^{\text{prev}} := \text{pk}_{\mathcal{C}}$. Whenever \mathcal{B} playing the role of \mathcal{S}_9 needs to issue a signature with respect to $\text{pk}_{\mathcal{R}}^{\text{prev}}$, it does so using its external EUF-CMA oracle $\text{OS}_{\text{pk}, \text{sk}}^S$. When the event ($E2$) 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. N.b., $(c_{\mathcal{R}}^{\text{prev}}, s^{\text{prev}})$ has never been signed with respect to $\text{pk}_{\mathcal{R}}^{\text{prev}} = \text{pk}_{\mathcal{C}}$ by assumption.
2. A record $\text{pid}_{\mathcal{R}}^{\text{prev}} \mapsto (\text{pk}_{\mathcal{R}}^{\text{prev}}, \text{sk}_{\mathcal{R}}^{\text{prev}}, \text{cert}_{\mathcal{R}})$ has not been recorded: We construct an efficient adversary \mathcal{B} against the EUF-CMA security of S along the same lines as above. Internally, \mathcal{B} runs $\mathcal{Z}^{\text{sys-sec}}$ in its head and plays the role of the simulator and all honest parties. Externally, \mathcal{B} plays the EUF-CMA security

□

³² N.b.: RSU may also denote the TSP, if the transaction at hand happens to be the first after a Wallet Issuing and thus s^* has been signed by the TSP playing the role an RSU. For brevity, we only consider RSUs here.

experiment with a challenger \mathcal{C} and a signing oracle $\mathcal{O}_{\text{pk,sk}}^S$. 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}_9 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}_9 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 $\mathcal{O}_{\text{pk,sk}}^S$. When the event $(E2)$ occurs, \mathcal{B} extracts $\text{cert}_{\mathcal{R}}^{\text{prev}} = (\text{pk}_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}}, \sigma_{\mathcal{R}}^{\text{cert}})$ from the proof and outputs $(\text{pk}_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}})$ together with $\sigma_{\mathcal{R}}^{\text{cert}}$ as the forgery. N.b.: $(\text{pk}_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}})$ 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.

The forgeries are indeed valid due to [Remark E.16](#). \square

Remark E.19. Without [Lemma E.18](#) it is unclear in [Lemma E.17, Item 2](#) if the denoted predecessor of edge (s^{prev}, s) actually exists. The simulator extracts the serial number s^{prev} of the predecessor from the proof and puts this serial number into the newly added $\overline{\text{trdb}}$. With this in mind [Lemma E.17, Item 2](#) would have to be interpreted such that an edge (s^{prev}, s) is ignored, if the predecessor did not exist. Nonetheless, $\overline{\text{TRDB}}$ is still a forest and [Lemma E.17, Item 2](#) remains correct. Anyway, this oddity is ruled out by [Lemma E.18](#).

Lemma E.20 (Indistinguishability between H_9 and H_{10}). *Under the assumptions of [Theorem E.10](#), $H_9 \stackrel{c}{\equiv} H_{10}$ holds.*

Proof. Assume there is an environment $\mathcal{Z}^{\text{sys-sec}}$ that trigger the event $E3$ with non-negligible advantage. This immediately yields an efficient adversary \mathcal{B} against the EUF-CMA security of S by the same argument as in the proof of [Lemma E.18](#) as $(c_{\mathcal{R}}^{\text{prev}}, s^{\text{prev}})$ are jointly signed by the same signature $\sigma_{\mathcal{R}}$. \square

Lemma E.21 (Indistinguishability between H_{10} and H_{11}). *Under the assumptions of [Theorem E.10](#), $H_{10} \stackrel{c}{\equiv} H_{11}$ holds.*

Proof. Assume there is an environment $\mathcal{Z}^{\text{sys-sec}}$ that trigger the event $E4$ with non-negligible advantage. 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 plays the role of the simulator and all honest parties. Externally, \mathcal{B} plays the role of the adversary as defined by [Definition C.7, Item 2](#). When the event $(E3)$ occurs, \mathcal{B} sets

$$M_{\mathcal{R}}^{\text{prev}} := (\Lambda, B^{\text{prev}}, U_1, X)$$

from the extracted witness and obtains

$$M_{\mathcal{R}}^{\text{out}^*} = (\Lambda^*, B^*, U_1^*, X^*)$$

from $\overline{\text{TRDB}}$. \mathcal{B} outputs $(c_{\mathcal{R}}^{\text{out}^*}, M_{\mathcal{R}}^{\text{prev}}, d_{\mathcal{R}}^{\text{prev}}, M_{\mathcal{R}}^{\text{out}^*}, d_{\mathcal{R}}^{\text{out}^*})$ to the external game. By assumption $\Lambda \neq \Lambda^*$ holds and [Remark E.16](#) asserts that both openings are valid. \square

Lemma E.22 (Tree-wise Uniqueness of the Wallet ID).

The wallet ID λ maps one-to-one and onto a connected component (i.e., tree) of the Simulated Transaction Graph.

Proof. Same argument as in the proof of [Lemma E.4](#). \square

Lemma E.23 (Indistinguishability between H_{11} and H_{12}). *Under the assumptions of [Theorem E.10](#), $H_{11} \stackrel{c}{\equiv} H_{12}$ holds.*

Proof. We introduce a sub-hybrid that splits between two cases why event $E5$ is triggered: 1 $c_{\mathcal{T}}^{\text{out}^*} \neq c_{\mathcal{T}}$ and $c_{\mathcal{T}}$ is not recorded in any $\overline{\text{trdb}} \in \overline{\text{TRDB}}$. 2 $c_{\mathcal{T}}^{\text{out}^*} \neq c_{\mathcal{T}}$ and $c_{\mathcal{T}}$ is recorded in some record $\overline{\text{trdb}}^{\ddagger} \in \overline{\text{TRDB}}$. An environment $\mathcal{Z}^{\text{sys-sec}}$ that can differentiate between H_{11} and the sub-hybrid yields an efficient adversary \mathcal{B} against the EUF-CMA security of S . An environment $\mathcal{Z}^{\text{sys-sec}}$ that can differentiate between the sub-hybrid and H_{12} yields an efficient adversary \mathcal{B} against the binding property of $C1$.

1. 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 plays the role of the simulator and all honest parties. Externally, \mathcal{B} plays the EUF-CMA security experiment with a challenger \mathcal{C} and a signing oracle $\mathcal{O}_{\text{pk,sk}}^S$. 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}_{12} 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}_{12} needs to issue signatures with respect to $\text{pk}_{\mathcal{T}}$, it does so using its external EUF-CMA oracle $\mathcal{O}_{\text{pk,sk}}^S$. When the event $(E5)$ occurs, \mathcal{B} extracts $c_{\mathcal{T}}$ and $\sigma_{\mathcal{T}}$ from the proof and outputs the forgery.
2. 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 plays the role of the simulator and all honest parties. Externally, \mathcal{B} plays the role of the adversary as defined by [Definition C.7, Item 2](#). As $(E5)$ has not been raised earlier, $c_{\mathcal{T}}^{\text{out}^{(i)}} = c_{\mathcal{T}}^{\text{out}^*} \neq c_{\mathcal{T}}$ holds for all $c_{\mathcal{T}}^{\text{out}^{(i)}}$ in the same tree. Consequently, $\overline{\text{trdb}}^{\ddagger}$ with $c_{\mathcal{T}}^{\text{out}^{\ddagger}} = c_{\mathcal{T}}$ is part of a different tree in $\overline{\text{TRDB}}$ and thus $\Lambda^{\ddagger} \neq \Lambda^* = \Lambda$ follows by

Lemma E.22. \mathcal{B} sets

$$M_{\mathcal{T}} := (\Lambda, \text{pk}_{\mathcal{U}})$$

from the extracted witness and obtains

$$M_{\mathcal{T}}^{\text{out}\ddagger} = (\Lambda^{\ddagger}, \text{pk}_{\mathcal{U}}^{\ddagger})$$

from \overline{TRDB} . \mathcal{B} outputs $(c_{\mathcal{T}}, M_{\mathcal{T}}, d_{\mathcal{T}}, M_{\mathcal{T}}^{\text{out}\ddagger}, d_{\mathcal{T}}^{\text{out}\ddagger})$ to the external game.

Remark E.16 asserts that the forgery in 1 and both openings in 2 are indeed valid. \square

Lemma E.24 (Indistinguishability $H_{12} \rightarrow H_{15}$).

Under the assumptions of [Theorem E.10](#), $H_{12} \stackrel{c}{\equiv} H_{13} \stackrel{c}{\equiv} H_{14} \stackrel{c}{\equiv} H_{15}$ holds.

Proof. If an environment $\mathcal{Z}^{\text{sys-sec}}$ can distinguish between any of the hops from H_{12} to H_{15} this yields an efficient adversary against the binding property of C1. As usual, \mathcal{B} runs $\mathcal{Z}^{\text{sys-sec}}$ in its head and internally plays the role of the simulator and all honest parties. Externally, \mathcal{B} plays the role of the adversary as defined by [Definition C.7, Item 2](#). If event (E7) or (E8) occurs, \mathcal{B} sets

$$M_{\mathcal{R}}^{\text{prev}} = (\Lambda, B^{\text{prev}}, U_1, g_1 X)$$

from the extracted witness and obtains

$$M_{\mathcal{R}}^{\text{out}*} := (\Lambda^*, B^*, U_1^*, X^*)$$

from \overline{TRDB} . \mathcal{B} outputs $(c_{\mathcal{R}}, M_{\mathcal{R}}^{\text{prev}}, d_{\mathcal{R}}^{\text{prev}}, M_{\mathcal{R}}^{\text{out}*}, d_{\mathcal{R}}^{\text{out}*})$ to the external game. If the event (E6) is triggered, \mathcal{B} proceeds analogous but for the fixed part of wallet $c_{\mathcal{T}}$. \square

Taking all the aforementioned statements together, [Theorem E.10](#) from the beginning of this appendix follows. For the sake of formal completeness we recall it again.

Theorem E.10 (Operator Security). Under the assumptions of [Theorem E.1](#)

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

holds under static corruption of

1. a subset of users, or
2. all users and a subset of RSUs, TSP and SA.

Proof. A direct consequence of [Lemmas E.12 to E.15, E.18, E.20, E.21, E.23 and E.24](#). \square

E.4 Proof of User Security and Privacy

In this appendix we show the following theorem.

Theorem E.25 (User Security and Privacy). Under the assumptions of [Theorem E.1](#)

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

holds under static corruption of

1. a subset of RSUs, TSP and SA, or
2. all RSUs, TSP and SA as well as a subset of users.

The definition of the UC-simulator $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{user-sec}}$ for [Theorem E.25](#) can be found in [Figs. 54 to 58](#). Please note that while the real protocol π_{P4TC} lives in the $(\mathcal{F}_{\text{CRS}}, \bar{\mathcal{G}}_{\text{bb}})$ -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 π_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

$$H_i := \text{EXEC}_{\pi_i, \bar{\mathcal{G}}_{\text{bb}}, \mathcal{S}_i, \mathcal{Z}^{\text{user-sec}}} (1^n).$$

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 π_i that honest parties perform gradually declines from the real protocol $\pi_0 = \pi_{\text{P4TC}}$ to a dummy protocol π_{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 $\mathcal{F}_{\text{P4TC}}$ 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

$$H_0 := \text{EXEC}_{\pi_0, \bar{\mathcal{G}}_{\text{bb}}, \mathcal{S}_0, \mathcal{Z}^{\text{user-sec}}} (1^n)$$

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 CRS_{pok} is generated by SetupSPoK , $\text{CRS}_{\text{com}}^1$ is generated by C1.SimGen

Simulator $\mathcal{S}_{\pi_{P4TC}}^{\text{user-sec}}$

Setup:

- (1) Run a modified version of the algorithm $\text{CRS} \leftarrow \text{Setup}(1^n)$ with
 - (a) $\text{CRS}_{\text{com}}^1 \leftarrow \text{C1.Gen}$ being replaced by $(\text{CRS}_{\text{com}}^1, \text{td}_{\text{eqcom}}) \leftarrow \text{C1.SimGen}$,
 - (b) $\text{CRS}_{\text{com}}^2 \leftarrow \text{C2.Gen}$ being replaced by $(\text{CRS}_{\text{com}}^2, \text{td}_{\text{extcom}}) \leftarrow \text{C2.ExtGen}$, and
 - (c) $\text{CRS}_{\text{pok}} \leftarrow \text{SetupPoK}$ being replaced by $(\text{CRS}_{\text{pok}}, \text{td}_{\text{spok}}) \leftarrow \text{SetupSPoK}$.
- (2) Record CRS , td_{eqcom} , $\text{td}_{\text{extcom}}$, and td_{spok} .
- (3) Set $\overline{\Omega}^{\text{dsp}} := \emptyset$.
- (4) Set $\overline{\Omega}_{\mathcal{U}}^{\text{pp}} := \emptyset$.
- (5) Set $\overline{HTD} := \emptyset$.

DR Registration: Upon receiving $(\text{registering_dr}, \text{pid}_{DR})$ run $(\text{pk}_{DR}, \text{sk}_{DR}) \leftarrow \text{DRRegistration}(\text{CRS})$, return pk_{DR} to \mathcal{F}_{P4TC} and record $\text{pid}_{DR} \mapsto (\text{pk}_{DR}, \text{sk}_{DR})$.

TSP Registration: Distinguish two cases:

TSP is corrupted: (nothing to do as this is a local algorithm for a corrupted TSP)

TSP honest: Upon receiving $(\text{registering_tsp}, \text{pid}_{\mathcal{T}}, \mathbf{a}_{\mathcal{T}})$ run $(\text{pk}_{\mathcal{T}}, \text{sk}_{\mathcal{T}}, \text{cert}_{\mathcal{T}}^{\mathcal{R}}) \leftarrow \text{TSPRegistration}(\text{CRS}, \mathbf{a}_{\mathcal{T}})$, return $\text{pk}_{\mathcal{T}}$ to \mathcal{F}_{P4TC} and record $\text{pid}_{\mathcal{T}} \mapsto (\text{pk}_{\mathcal{T}}, \text{sk}_{\mathcal{T}}, \text{cert}_{\mathcal{T}}^{\mathcal{R}})$.

RSU Registration: Distinguish two cases:

RSU is corrupted: (nothing to do as this is a local algorithm for a corrupted RSU)

RSU honest: Upon receiving $(\text{registering_rsu}, \text{pid}_{\mathcal{R}})$ run $(\text{pk}_{\mathcal{R}}, \text{sk}_{\mathcal{R}}) \leftarrow \text{RSURegistration}(\text{CRS})$, return $\text{pk}_{\mathcal{R}}$ to \mathcal{F}_{P4TC} and record $\text{pid}_{\mathcal{R}} \mapsto (\text{pk}_{\mathcal{R}}, \text{sk}_{\mathcal{R}})$.

User Registration: Upon receiving $(\text{registering_user}, \text{pid}_{\mathcal{U}})$ run $(\text{pk}_{\mathcal{U}}, \text{sk}_{\mathcal{U}}) \leftarrow \text{UserRegistration}(\text{CRS})$, return $\text{pk}_{\mathcal{U}}$ to \mathcal{F}_{P4TC} and record $\text{pid}_{\mathcal{U}} \mapsto (\text{pk}_{\mathcal{U}}, \text{sk}_{\mathcal{U}})$.

RSU Certification: Distinguish four cases:

TSP and RSU honest: Upon receiving $(\text{certifying_rsu}, \text{pid}_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}}) \dots$

- (1) Load the recorded $\text{pid}_{\mathcal{T}} \mapsto (\text{pk}_{\mathcal{T}}, \text{sk}_{\mathcal{T}}, \text{cert}_{\mathcal{T}}^{\mathcal{R}})$, and $\text{pid}_{\mathcal{R}} \mapsto (\text{pk}_{\mathcal{R}}, \text{sk}_{\mathcal{R}})$; if any of these do not exist, let \mathcal{F}_{P4TC} abort.
- (2) Generate $\text{cert}_{\mathcal{R}} := (\text{pk}_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}}, \sigma_{\mathcal{R}}^{\text{cert}})$ with $\sigma_{\mathcal{R}}^{\text{cert}} \leftarrow \text{S.Sgn}(\text{sk}_{\mathcal{T}}^{\text{cert}}, (\text{pk}_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}}))$ faithfully.
- (3) Update record $\text{pid}_{\mathcal{R}} \mapsto (\text{pk}_{\mathcal{R}}, \text{sk}_{\mathcal{R}}, \text{cert}_{\mathcal{R}})$.

TSP honest, RSU corrupted: Upon receiving $(\text{certifying_rsu}, \text{pid}_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}}) \dots$

- (1) Load the recorded $\text{pid}_{\mathcal{T}} \mapsto (\text{pk}_{\mathcal{T}}, \text{sk}_{\mathcal{T}}, \text{cert}_{\mathcal{T}}^{\mathcal{R}})$, and obtain $\text{pk}_{\mathcal{R}}$ from $\overline{\mathcal{G}}_{\text{bb}}$ for $\text{pid}_{\mathcal{R}}$; if any of these do not exist, let \mathcal{F}_{P4TC} abort.
- (2) Generate $\text{cert}_{\mathcal{R}} := (\text{pk}_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}}, \sigma_{\mathcal{R}}^{\text{cert}})$ with $\sigma_{\mathcal{R}}^{\text{cert}} \leftarrow \text{S.Sgn}(\text{sk}_{\mathcal{T}}^{\text{cert}}, (\text{pk}_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}}))$ faithfully.
- (3) Record $\text{pid}_{\mathcal{R}} \mapsto (\text{pk}_{\mathcal{R}}, \perp, \text{cert}_{\mathcal{R}})$.
- (4) Output cert to $\mathcal{Z}^{\text{user-sec}}$.

TSP corrupted, RSU honest: Upon receiving $(\text{cert}_{\mathcal{R}})$ from $\mathcal{Z}^{\text{user-sec}}$ in the name of \mathcal{T} with $\text{pid}_{\mathcal{T}} \dots$

- (1) Load the recorded $\text{pid}_{\mathcal{R}} \mapsto (\text{pk}_{\mathcal{R}}, \text{sk}_{\mathcal{R}})$, and obtain $\text{pk}_{\mathcal{T}}$ from $\overline{\mathcal{G}}_{\text{bb}}$ for $\text{pid}_{\mathcal{T}}$; if any of these do not exist, let \mathcal{F}_{P4TC} abort.
- (2) Parse $\mathbf{a}_{\mathcal{R}}$ and $\sigma_{\mathcal{R}}^{\text{cert}}$ from $\text{cert}_{\mathcal{R}}$.
- (3) If $\text{S.Vfy}(\text{pk}_{\mathcal{T}}^{\text{cert}}, \sigma_{\mathcal{R}}^{\text{cert}}, (\text{pk}_{\mathcal{R}}, \mathbf{a}_{\mathcal{R}})) = 0$, let \mathcal{F}_{P4TC} abort.
- (4) Call \mathcal{F}_{P4TC} with input $(\text{certify}, \mathbf{a}_{\mathcal{R}})$ in the name of \mathcal{T} with $\text{pid}_{\mathcal{T}}$.

TSP and RSU corrupted: (nothing to do as $\mathcal{Z}^{\text{user-sec}}$ plays both parties)

Fig. 54. The simulator for User Security and Privacy

Simulator $\mathcal{S}_{\mathcal{F}_{P4TC}}^{\text{user-sec}}$ (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 1st message from \mathcal{U} to \mathcal{T}
- (4) Upon receiving $(\text{cert}_{\mathcal{T}}^{\mathcal{R}}, \mathbf{a}_{\mathcal{U}}, c''_{ser}, \lambda'')$ from $\mathcal{Z}^{\text{user-sec}}$ in the name of \mathcal{T} with $pid_{\mathcal{T}} \dots$ ^a
 - (a) Parse $(pk_{\mathcal{T}}^{\mathcal{R}}, \mathbf{a}_{\mathcal{T}}, \sigma_{\mathcal{T}}^{\text{cert}}) := \text{cert}_{\mathcal{T}}^{\mathcal{R}}$.
 - (b) If $\text{S.Vfy}(pk_{\mathcal{T}}^{\text{cert}}, \sigma_{\mathcal{T}}^{\text{cert}}, (pk_{\mathcal{T}}^{\mathcal{R}}, \mathbf{a}_{\mathcal{T}})) = 0$ abort.
 - (c) $\Lambda'' := g^{\lambda''}$
 - (d) $s'' \leftarrow \text{C2.Extract}(\text{CRS}_{com}^2, \text{td}_{\text{extcom}}, c''_{ser})$.
 - (e) Call \mathcal{F}_{P4TC} with input $(\text{issue}, \mathbf{a}_{\mathcal{U}}, \emptyset)$ in the name of \mathcal{T} with $pid_{\mathcal{T}}$.^b
- (5) Upon receiving output (s) from \mathcal{F}_{P4TC} for $\mathcal{T} \dots$
 - (a) $s' := s \cdot s''^{-1}$.
 - (b) $r_1, r_2 \xleftarrow{\mathbb{R}} \mathbb{Z}_p$.
 - (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 2nd 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 $\overline{htd} := (pk_{\mathcal{U}}, s, \lambda'', e^*)$ and insert \overline{htd} into \overline{HTD} .
 - (c) Create real token τ 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.

^a 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.

^b Use empty set as blacklist. If the TSP intended to blacklist the user, the TSP would not have sent the previous message.

Fig. 55. The simulator for User Security and Privacy (cont. from Fig. 54)

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 $(\text{pay_toll}, \emptyset)^{\mathbf{a}}$ 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 (ϕ) 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 (ϕ, 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) $\pi \leftarrow \text{P2.SimProof}(\text{CRS}_{\text{pok}}^1, \text{td}_{\text{spok}}, \text{stmnt})$.
 - (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 $\mathcal{Z}^{\text{user-sec}}$... **b**
 - (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.

a Use empty set as blacklist.

b 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. 56. The simulator for User Security and Privacy (cont. from Fig. 54)

Simulator $\mathcal{S}_{\mathcal{F}_{P4TC}}^{\text{user-sec}}$ (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 $\mathbf{a}_{\mathcal{R}}^{\text{prev}}$.
 - (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 $(\phi, t', u'_2) \in \overline{\mathcal{O}}^{\text{dsp}}$ has been recorded previously with (ϕ) being used as key. If no, pick $t \xleftarrow{\mathbb{R}} \mathbb{Z}_p$. 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 (ϕ, t, u_2) into $\overline{\mathcal{O}}^{\text{dsp}}$.
 - (b) $stmt := (pk_{\mathcal{U}}, pk_{\mathcal{T}}, pk_{\mathcal{T}}^{\text{cert}}, \phi, \mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}}, g_1^{b^{\text{bill}}}, t, u_2)$.
 - (c) $\pi \leftarrow \text{P3.SimProof}(\text{CRS}_{\text{pok}}, \text{td}_{\text{spok}}, stmt)$.
 - (d) Output $(pk_{\mathcal{U}}, \pi, \phi, \mathbf{a}_{\mathcal{U}}, \mathbf{a}_{\mathcal{R}}^{\text{prev}}, b^{\text{bill}}, t)$ to $\mathcal{Z}^{\text{user-sec}}$.
- (4) Upon receiving (OK) from $\mathcal{Z}^{\text{user-sec}}$,^a let \mathcal{F}_{P4TC} return the delayed output to the User.

^a If no message is received, let \mathcal{F}_{P4TC} abort.

Fig. 57. The simulator for User Security and Privacy (cont. from Fig. 54)

Simulator $\mathcal{S}_{\mathcal{F}_{P4TC}}^{\text{user-sec}}$ (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 $(\text{prove_participation}, pid_{\mathcal{U}}, S_{\mathcal{R}}^{\text{PP}})$.
 - (b) Obtain the SA 's output (out) from \mathcal{F}_{P4TC} .
 - (c) If out = NOK, abort.
 - (d) Pick $\bar{\omega}_{\mathcal{U}}^{\text{PP}} = (s, c_{\text{hid}}, \bar{d}_{\text{hid}})$ from $\bar{\Omega}_{\mathcal{U}}^{\text{PP}}$ 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 $g_{\mathcal{I}}^{\pi} = pk_{\mathcal{U}}$, 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) $HTD_{\mathcal{U}}^{\text{genuine}} = \{(pk_{\mathcal{U}}, s, \lambda'', e^*) \mid (\cdot, \cdot, \lambda'', e^*) \in HTD_{\mathcal{U}} \wedge (pk_{\mathcal{U}}, s, \lambda'', e^*) \in \overline{HTD}\}$
 - (b) $HTD_{\mathcal{U}}^{\text{fake}} = \{(pk_{\mathcal{U}}, s, \lambda'', e^*) \mid (\cdot, s, \lambda'', e^*) \in HTD_{\mathcal{U}} \wedge (\cdot, \cdot, \lambda'', e^*) \notin \overline{HTD} \wedge (\dots, pk_{\mathcal{U}}) \leftarrow \text{E1.Dec}(sk_{DR}, e^*)\}$
 - (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 $(\text{blacklist_user}, pk_{\mathcal{U}}^{\mathcal{T}})$.
- (3) Upon being ask by \mathcal{F}_{P4TC} to provide S_{root} ...
 - (a) $S_{\text{root}} := \{s \mid (\cdot, s, \cdot, \cdot) \in HTD_{\mathcal{U}}^{\text{genuine}}\}$
 - (b) Provide S_{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 Φ^{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. 58. The simulator for User Security and Privacy (cont. from Fig. 54)

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 H₂

Hybrid H₂ replaces the code in the tasks DR/TSP/RSU/User Registration of the protocol π_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 H₃

In hybrid H₃ the task RSU Certification is modified. For an honest TSP or an honest RSU the code of π_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 H₄

H₄ modifies the tasks of Wallet Issuing and Debt Accumulation. The code of π_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.

Hybrid H₅

This hybrid modifies π_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 H₆

H₆ modifies π_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 $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{user-sec}}$.

Hybrid H₇

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 H₄.

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”³³ $\overline{HTD}_{\mathcal{U}}^{\text{genuine}}$ and $\overline{HTD}_{\mathcal{U}}^{\text{fake}}$ (cp. Fig. 58, Steps 2a to 2c). If for a hidden user trapdoor $htd = (\text{pk}_{\mathcal{U}}, s, \lambda'', e^*)$ a corresponding entry $(\cdot, \cdot, \lambda'', e^*)$ is recorded in \overline{HTD} , then we call it a *genuine* hidden user trapdoor and the public key $\text{pk}_{\mathcal{U}}$ 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 $\text{pk}_{\mathcal{U}}$ is set to the decrypted value from e^* .³⁴ \mathcal{S}_7 additionally checks if the hidden user trapdoors of both sets $\overline{HTD}_{\mathcal{U}}^{\text{genuine}}$ and $\overline{HTD}_{\mathcal{U}}^{\text{fake}}$ belong to the same user public key $\text{pk}_{\mathcal{U}}^{\mathcal{T}}$ or else aborts. This equals the behavior of the final simulator.

\mathcal{S}_7 runs the code of an honest DR on $\overline{HTD}_{\mathcal{U}}^{\text{genuine}} \cup \overline{HTD}_{\mathcal{U}}^{\text{fake}}$ to recover $\Phi_{\mathcal{U}}$, i.e., it decrypts every hidden user trapdoor and evaluates the PRF itself. For the DR the code of π_7 is changed such that it simply signals its consent by forwarding its input $\text{pk}_{\mathcal{U}}^{\text{DR}}$ to \mathcal{S}_7 .

Hybrid H₇ 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 $\mathcal{S}_{\pi_{\text{P4TC}}}^{\text{user-sec}}$ as hybrid H₉ 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 λ^{fake} and create a syntactically valid hidden user trapdoor for λ^{fake} and any public user key $\text{pk}_{\mathcal{U}}$. 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.³⁵ The final 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.

³³ 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 $(\text{pk}_{\mathcal{U}}, s, \cdot, \cdot)$ are sanitized.

³⁴ We assume, that Dec returns \perp if e^* cannot be decrypted.

³⁵ 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.

Hybrid H₈

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.³⁶

$\mathcal{F}_{\phi\text{-rand}}$ provides the following functionality: Internally, $\mathcal{F}_{\phi\text{-rand}}$ manages a partial map f_{Φ} , mapping pairs of wallet IDs λ 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 ϕ for (λ, 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 λ , 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 λ . 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 λ and returns the set $\{f_{\Phi}(\lambda, 0), \dots, f_{\Phi}(\lambda, x_{\text{bl}_{\mathcal{R}}})\}$.

An honest user running π_8 does not evaluate the PRF to obtain a fraud detection ID ϕ , 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 H₉

In the scope of Wallet Issuing π_9 is modified such that honest users do not send e^* . Instead, \mathcal{S}_6 injects an encrypted zero-vector as e^* .

Hybrid H₁₀

Hybrid H₁₀ 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 ϕ and sets $f_{\Phi}(\lambda, x) := \phi$.

Hybrid H₁₁

In H₁₁ 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 $t \xleftarrow{\mathbb{R}} \mathbb{Z}_p$, else \mathcal{S}_{11} sets

$t := t' + \text{sk}_{\mathcal{U}}(u_2 - u'_2)$. Finally, \mathcal{S}_{11} inserts (ϕ, 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 H₁₂

The hybrid H₁₂ 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 H₀ to H₁₂ yields

$$\begin{aligned} \text{H}_{12} &= \text{EXEC}_{\pi_{12}, \bar{g}_{\text{bb}}, \mathcal{S}_{12}, \mathcal{Z}^{\text{user-sec}}}(1^n) \\ &= \text{EXEC}_{\mathcal{F}_{\text{P4TC}}, \bar{g}_{\text{bb}}, \mathcal{S}_{\pi\text{P4TC}}^{\text{user-sec}}, \mathcal{Z}^{\text{user-sec}}}(1^n). \end{aligned}$$

With these hybrids we can now give the proof of [Theorem E.25](#). We do not spell out all steps of the proofs in full detail, but rather sketch the necessary reductions.

Proof of [Theorem E.25](#).

From H₀ to H₁: This hop solely changes how the CRS is created during the setup phase. This is indistinguishable for CRS_{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 H₁ to H₂: 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 H₂ to H₃: 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 H₃ to H₄: 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 $\mathcal{S}_{\pi\text{P4TC}}^{\text{user-sec}}$.

From H₄ to H₅: 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

³⁶ I.e., communication is confidential, reliable and trustworthy. One might think of this entity as a preliminary version of the eventual ideal functionality.

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'_{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^{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^{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 $\{\text{PRF}(\lambda, 0), \dots, \text{PRF}(\lambda, x_{\text{bl}_{\mathcal{R}}})\}$, if required. Again, this is possible due to the correctness of E1, i.e., e^* uniquely determines λ and thus maps to a unique s .

From H_8 to H_9 : In this hop every encryption e^* of a wallet ID λ 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^{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 λ 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 λ 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.12](#). Internally, \mathcal{B} runs $\mathcal{Z}^{\text{user-sec}}$ and plays the protocol and simulator for $\mathcal{Z}^{\text{user-sec}}$. Externally, \mathcal{B} in-

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 λ , \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 λ . 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 $\text{sk}_{\mathcal{U}}$.

From H_{11} to H_{12} : In this hop the simulator \mathcal{S}_{12} sends simulated commitments c_{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 $\text{pk}_{\mathcal{U}}$, 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