

Privacy for Public Transportation

Thomas S. Heydt-Benjamin, Hee-Jin Chae, Benessa Defend, and Kevin Fu

University of Massachusetts, Amherst, MA 01003, USA
{tshb, chae, defend, kevinfu}@cs.umass.edu

Abstract. We propose an application of recent advances in e-cash, anonymous credentials, and proxy re-encryption to the problem of privacy in public transit systems with electronic ticketing. We discuss some of the interesting features of transit ticketing as a problem domain, and provide an architecture sufficient for the needs of a typical metropolitan transit system. Our system maintains the security required by the transit authority and the user while significantly increasing passenger privacy. Our hybrid approach to ticketing allows use of passive RFID transponders as well as higher powered computing devices such as smartphones or PDAs. We demonstrate security and privacy features offered by our hybrid system that are unavailable in a homogeneous passive transponder architecture, and which are advantageous for users of passive as well as active devices.

1 Introduction

Public transportation ticketing systems must be able to handle large volumes of passenger transactions while providing the minimum possible impedance to travel. Therefore, it is hardly surprising that some of the world's busiest public transportation systems are at the forefront of electronic payment technology. Unfortunately, current systems have been designed such that passengers sacrifice privacy in order to take advantage of the convenience of electronic payment. Moreover, because of the inherent broadcast nature of RF, as systems migrate from contact based technologies like mag-stripe to contactless technologies there is increased risk to privacy and security [21, 27, 6, 15, 33].

The traditional passive RFID transponder is a severely resource constrained computing device. Manufacturing cost is usually a primary design criterion, resulting in transponders with little memory and processing power. Even in more expensive passive transponders, current technology limits the amount of memory and the complexity of the microprocessor that can fit into common form-factors. Furthermore, since passive transponders are powered by electrical induction from the reader's antenna, an RFID tag must power up, receive, process, and transmit within the brief time that a user holds the tag within the reader's electric

This research was partially supported by NSF CNS-052072 and a Ford Foundation Diversity Fellowship.

field. Consequently, many of the security protocols that we would use for communication between other kinds of computers are inappropriate for the RFID platform [38]. However, despite their resource constraints, cards with cryptographic co-processors are capable of executing carefully crafted protocols [18, 37, 26, 29, 20].

As the abilities of contactless smart cards have increased, new cryptographic primitives suitable for these resource constrained devices have been developed. Not only do recent contributions to the field of e-cash and anonymous credentials require much less memory, but the communications required for the zero-knowledge proofs are also greatly reduced [10, 11, 9]. The key management problem for a transit system involving hundreds of readers and hundreds of thousands of tickets has traditionally been difficult. We apply recent advances in re-encryption and re-signatures to place the burden of key management on the more powerful computers in the system, requiring the tickets to store only the public portion of a single highly secure key pair whose private portion can be protected in offline storage [3, 4].

1.1 Background

In 2004, passengers took approximately 9 billion trips through public transportation systems in the United States [16]. Existing systems maintain a database of all transactions, associating them with the identities of passengers whenever possible, such as when a credit card is used in conjunction with the transit card [35, 41, 31]. If communication between a ticket and a transit authority is not properly secured, arbitrary third party adversaries might then have inappropriate access to user data. Many of the currently deployed systems are proprietary [24], and thus closed to scientific scrutiny. Recent historical examples, such as the black-box cryptanalysis of TI's major RFID security mechanism [7], reinforce that eschewing peer review often leads to insecure systems. Even if the RF communication in a transit system is secure, the user's data may still be at risk. The Washington D.C. Metro operated for years without a clearly defined privacy policy [36, 40, 39]. Until recently, users of this system had no legal protection preventing the sale or sharing of their data with third parties. Privacy preserving protocols are needed to protect this large volume of sensitive data.

The utility of privacy to the individual consumer is clear, however the very consumer data that we wish to protect has long been considered valuable to the transit authority. We feel that at a certain point organizations such as transit authorities may wish to scale back on the amount of consumer data they collect. They may come to view such information as a greater liability than an asset since they stand to lose both money and reputation if the data leaks to adversarial parties. Additionally, growing public unease about ubiquitous surveillance may lead to legislation, commercial pressure, or societal pressure forcing companies to adopt stronger privacy technologies. Ultimately a new equilibrium may be achieved in which systems may be designed to permit gathering of useful business data while reassuring the consumer by providing scientific guarantees that such data will be appropriately anonymized.

Many large transit systems are still in the process of choosing and implementing new ticket technologies. The San Francisco Bay Area Rapid Transit (BART) system, for example, had over 91 million passengers in 2004 [32] and is currently in the process of considering how best to implement future RFID ticketing. BART has expressed willingness to listen to suggestions from the scientific community. We hope that our community will respond with protocols that give transit authorities the proper tools.

1.2 Our Contributions

Our research makes three primary contributions that address the challenges of privacy and security in public transportation:

1. We motivate the study of transit system payment as a problem domain with interesting properties and many open problems for research.
2. We propose a framework for reasoning about transit system payment security and privacy.
3. We present novel designs for systems offering RF electronic payment which we discuss in the context of our proposed framework.

Our design provides a payment system suitable for the needs of a typical transit system, in which the transit agency retains the ability to implement a variable rate fare structure. The movements of a user of our protocols cannot be tracked through the transit system by the transit authority nor by a third-party adversary. Our novel authentication protocol built around the re-encryption primitive [3] provides verification of reader authorization and also provides a secure channel in a manner well suited to the resource constraints of the various systems. Reader authorization in our design is efficient and secure, and does not require propagation of revocation information.

2 Related Work

Other researchers have proposed the use of actively powered devices for payment system or RFID anonymity [13, 17, 22]. By contrast we propose a hybrid system which takes advantage of the abilities of more powerful devices such as smartphones, while remaining compatible with more commonly deployed passively powered RFID transponders. Additionally, whereas much prior work exists relating to electronic payment, our focus is on a specific real-world problem domain, with consideration for such issues as the trade-offs between anonymity and certain mandatory and optional transit requirements.

The Advanced Fare Payment Systems Company [25] provides an overview of different types of cards that could be used for transit systems ticketing. They do not give details about card security and mention gathering user data as an advantage of implementing RFID, which contrasts with our goal of protecting user data.

There is much existing work on RFID security, including resistance to tracking and hotlisting attacks [2, 15, 23, 30]. However, our paper is unique in considering them in the specific context of public transit. RFID privacy techniques such as Blocker Tags [19] and Faraday cages, which prevent communication with a transponder, are insufficient since they do not protect privacy when a ticket must be legitimately read.

Systems exist which address the security needs of RF transit ticketing, but do not significantly consider privacy of user data. Many publications consider unique card serial numbers as a requirement for fraud detection [5, 34]. We propose the use of advanced anonymous credentials and e-cash systems, which can detect fraud while maintaining the anonymity of the honest user.

3 The Problem Domain of Transit System Payment

Transit systems have historically been at the forefront of experimenting with new payment technologies [35, 24]. Yet increased security often comes at the expense of privacy. For instance, a transit card that records a passenger's travel history may reduce fraud at the expense of privacy. Below we discuss several challenges to providing freedom from ubiquitous surveillance while also maintaining or increasing security.

A cryptographic transit ticket is a resource constrained computing device. Such tickets are currently implemented on passive RFID transponders with severe limits on power, memory, and CPU, and in more advanced systems on higher powered embedded computing devices (HPDs), such as cell phones or PDAs. These resource constraints raise many compelling questions, as the systems requirements frequently force trade-offs with security or privacy features. However, it is also the case that considering a cryptosystem in the context of a very specific problem, rather than examining it in its general and abstract theory, may permit abbreviations of feature sets which lower the resource requirements of strong cryptosystems. For example, we may assume that the value stored on an e-cash based transit ticket will decrease monotonically. Not offering support for adding tokens to the ticket's wallet may allow savings of memory, transaction time, and CPU time. Transit tickets have limited communications bandwidth, but we will see below that this is an asset to security as well as a constraint.

A remarkable element of the problem domain lies in consideration of hybrid systems, which include both HPDs and passive transponders. HPDs can offer security and privacy benefits not only to the HPD user, but also to the passive transponder user. We examine one such case in section 6. In order for an HPD to enhance the security and privacy of a passive transponder user, however, the HPD and the passive transponder must be difficult to distinguish from one another. If a system permits heterogeneous HPDs based on different technologies from different hardware manufacturers, it may be challenging to ascertain this difficulty. We believe that there are many interesting problems related to this issue such as the problem of building an HPD which behaves as much like a passive transponder as possible, the problem of building a passive transponder with

less predictable power and communications patterns, the problem of building readers with highly accurate antenna power analysis for attacking transponder indistinguishability, and other similar problems.

Cloning detection for temporally bounded tickets (such as weekly or monthly commuting passes) is another fundamental problem related to transit system payment. In general, smartcard manufacturers rely on tamper resistance for cloning prevention [28]. We consider this to be insufficient, as tamper resistance has been shown to be weak in many cases [1]. Cloning detection for anonymous credentials systems exists [14] (and we assume such detection in our design), however these detection schemes are most effective in systems that require credential holders to be online simultaneously. Such mechanisms are insufficient for the needs of transit systems, yet the nature of transit systems may allow other bounds (like the aforementioned communications constraints) on adversarial behavior which will serve to provide more appropriate cloning detection.

For simplicity, we have assumed a strongly connected transit system in which all readers have a reliable network connecting them to central transit authority computers. There are many things to be considered if support is to be offered for weakly-connected networks. For example, we believe that ticket revocation information and other such data could be propagated using disruption tolerant networking techniques, such as packet ferrying [8, 42, 43]. Bus readers with embedded wireless networked computers and even tickets themselves may ferry data as they move through the system.

In this paper we assume that the user of the HPD is able to back up their virtual tickets through some mechanism external to our protocols. Another interesting facet of the domain of transit system payment is the question of how the user of a traditional transponder can back up their ticket without compromising their privacy. Ideally, the transit authority could retain a secure copy of the ticket at the time of purchase, but it is critical that user authorization be required in order to decrypt this backup copy. One possibility that we have considered is that when a smart credit card is used to purchase a ticket, the credit card could provide a mechanism for encryption of the ticket data which could then safely be stored by the transit authority. An ideal such mechanism could optionally allow for anonymity revocation by an authorized entity such as a judicial system when such is desirable or required by law.

The forward secrecy available to a user, should an adversarial transit authority obtain physical possession of the user's ticket, will be highly dependent on the underlying e-cash and anonymous credentials systems. The choice of a specific cryptographic primitive will determine what information the authority can learn given full knowledge of the ticket's data and all past transactions.

These are just a few of the problems worthy of study in the domain of public transit payment. Some of these problems are not unique to transit payment, but consideration of a specific real-world application may lead to development of more general techniques.

4 Definitions and Notation

We apply the traditional meanings of security, anonymity and privacy. We consider a protocol to be secure if it is as difficult to violate the semantics of the protocol as it is to break the underlying cryptographic primitives, and we consider anonymity to mean indistinguishability within a group of transit users. Therefore the degree of anonymity provided to some user u with the transit system in a particular state, is the size of the set of users that are indistinguishable from u up to the strength of the underlying cryptographic primitives. We consider the degree of privacy offered by a transit system to be the degree of difficulty with which an adversary can link a user's identity (such as name or credit card number) with their actions within the transit system over time. For example, the purchase of a ticket with a credit card will be an identifying transaction since the user's name is presented to the system, but the overall system provides privacy to the extent that it is difficult to link this purchase with other events such as entrances and exits. Thus the system may know when and where a user purchases a ticket, but will not know to where the user travels, nor whether they transfer or otherwise re-enter the system.

4.1 Adversaries

Transit system user = U : U possesses the ticket TX and may read or modify any of the ticket data. We assume that U will do any thing she can that will maximize her expected economic utility. U is willing to break the rules of any protocol if it is to her advantage, and if it helps her U may have a non-standard transponder with any reasonable design parameters. We assume that U will help other users steal service from the transit authority as long as such action does not require significant resources from her.

The Transit Authority = TA : The TA is assumed to be controlled by entities hostile to anonymity who wish to identify and track all users of the transit system. If it can be in any way advantageous, TA will carry out extra (unauthorized) transactions using concealed readers both inside and outside of the transit system. We categorize any entity which knows TA 's private key as being equivalent to TA .

We will not consider the various denial-of-service attacks that TA may perform upon U . It may be interesting in the future to consider mechanisms by which U can be protected from spurious charges or cancellation of valid tickets, but existing systems do not offer any such protection and this is not the focus of this paper.

The malicious 3rd party M : M is assumed to be able to read and or modify all data that is broadcast via RF. In order to provide a worst case analysis, we will assume that M has an RFID reader which can read/write to any tag over any arbitrary distance, and which can act as a perfect man-in-the-middle between a transponder and another reader. M is interested in doing anything that would maximize M 's expected economic utility, and will also attempt to degrade the anonymity of users of the system.

4.2 Semantics Required from Credential and E-cash Systems

Our protocols are designed to work with any anonymous credentials system that obeys the semantics outlined below. These semantics are similar to those described in [10] which provides a system compatible with our needs. We have found [11] to be quite suitable to the resource constraints of a passive RFID transponder, and we have begun working on a proof-of-concept implementation of our system based on this credentials system.

FormNym(TX, TA): A session between the ticket TX and the credentials granting organization or their designee. TX and TA negotiate a pseudonym $N_{(TX,TA)}$ part of which is stored by each party, and additional cryptographic validating tags sufficient for TA to later verify $N_{(TX,TA)}$ without TA having any knowledge of TX 's private information (such as a master private key).

GrantCred(N, λ, TA): A session between TX and TA (or designee) in which TX identifies itself to TA by the previously negotiated pseudonym $N_{(TX,TA)}$, and specifies the range of parameters $\lambda \subseteq A$ for which the credential shall be valid. TA creates a credential $C_{(TX,\lambda,TA)}$ which shall be valid only for the specified $\lambda \subseteq A$, and grants it to TX . The credential thus formed can be demonstrated without revealing TX or N , but given TX or N the credential can be revoked.

VerifyCred(N, τ, F): A session between the TX and some verifier (the fare-gate F) in which TX proves possession of credential $C_{(TX,\lambda,TA)}$, and TX furthermore demonstrates that the credential is valid for the parameter $\tau \in \lambda$. This can be accomplished without F knowing any private information belonging to TX or TA , and F cannot determine anything about λ other than $\tau \in \lambda$.

RevokeCred(N, TA): Given knowledge of N the TA may revoke all credentials granted to pseudonym N , and given TX the TA may revoke all credentials granted to TX regardless of pseudonym. We assume that N is revealed when cloning of multi-show credentials is detected (as in [14]), and when double-spending of single-show credentials is detected.

The e-cash semantics that we require are similar to those provided by [9, 12]:

CreateTokens(TA, TX, ν): A session between the Transit Authority TA and the ticket TX in which the TA creates a number ν of valid tokens. These tokens are transmitted and stored on TX . The wallet thus generated must be small enough to fit on a contactless smart card.

SpendTokens(TA, TX, ν): A session between the TA and the TX in which the TX spends ν tokens. The TX transmits the tokens to the TA and then deletes those tokens. This transaction must be unlinkable with *CreateTokens* unless double-spending has occurred.

5 A Semantic Framework for Transit System Payment

In this section we examine the basic properties of transit system payment and offer an analysis of the degree of identification inherent in each transaction. For

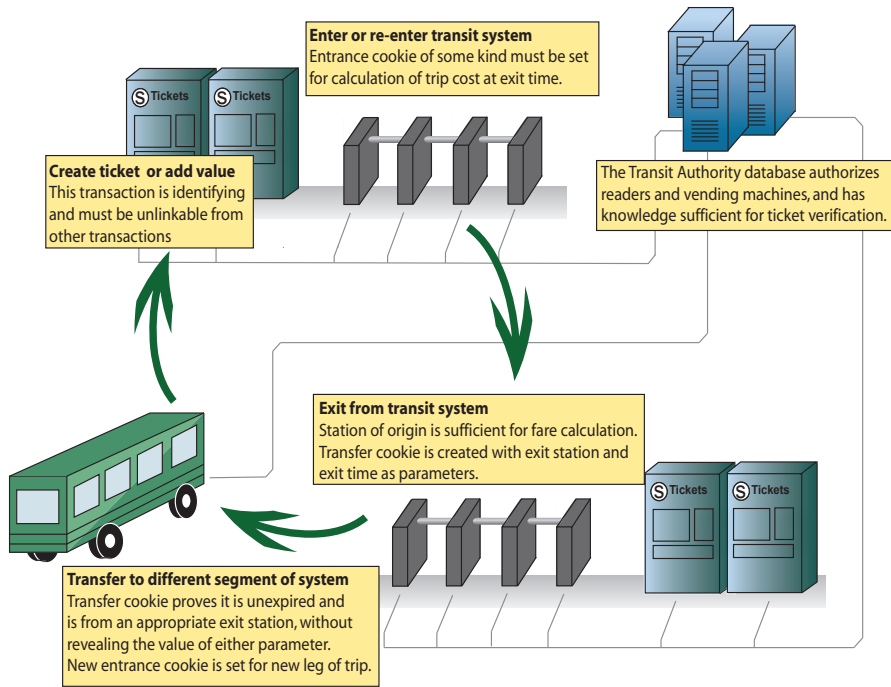


Fig. 1. The major pieces of a transit payment system

simplicity we will consider only the case of a strongly connected transit system, in which every faregate has a network connection to a central TA database. Our system is tolerant of intermittently connected faregates, however for brevity we will save further consideration of these features for future work. We will discuss the semantics of a transit system similar to that depicted in Figure 1 which is a variable-rate system with transfer between two components (bus and subway). Fixed-rate payment and additional transfer components can be trivially composed from these semantics.

$createTicket(TV, TX, U) \rightarrow TX$: A transaction between a ticket vendor (TV , designated by the TA through appropriate cryptographic keys), a ticket TX , and some form of payment external to the transit system. Such forms of payment may be anonymous, such as paper currency, or they may be identifying, such as credit cards. For simplicity we will assume the worst case: that all ticket creation transactions fully identify the user (indicated by U). A privacy-preserving transit system must therefore ensure that future transactions of TX cannot be linked to the ticket creation transaction.

$enter(F, TX) \rightarrow C_E(F)$: A transaction between a faregate F , and a ticket TX . In order to enter the system the ticket must demonstrate its validity to F . The final fare cannot be calculated until exit-time, therefore an entry cookie

C_E of some sort must be generated so that TX may later demonstrate at which F it entered the system. TA and F require no information other than that sufficient to accomplish payment and proof of validity. In a privacy-preserving system no other information should be leaked.

$exit(F, TX, C_E) \rightarrow C_T(F, e)$: When the user leaves a section of the transit system, a transaction is required between the exit faregate F , and the ticket TX . The ticket must prove the validity of its entrance cookie C_E and pay the fare for the trip, and in return it is given a transfer C_T . This cookie is parameterized with the point of exit F (for determining to which segments the user may transfer) and a time epoch e (to permit transfer expiration). TA and F require only payment or proof of appropriate credential and C_E which should reveal only where the user entered. In order to prevent fraud, C_E should prevent double-spending.

$transfer(F, C_T) \rightarrow C_E(F)$: C_T is proven to F . The proof mechanism (which should be zero-knowledge) ensures that C_T 's parameters fall within the range permissible for a valid transfer: i.e. the user is transferring from a permissible section of the transit system, and the transfer is not expired. The C_E generated by this transaction is done so by the same means as in the *enter* transaction. C_T should be double-spending proof. In an ideal privacy-preserving system, the only information that should be revealed is that the transfer has not yet been spent, that it is not yet expired, and that it comes from some exit faregate within a range of acceptable such faregates.

$addValue(TV, TX, U) \rightarrow TX'$: All value is taken off (is spent) of TX , the user provides additional value from an external payment source, and then the transaction proceeds as in initial ticket creation. This provides a new ticket TX' , which is unlinkable with TX and with U .

$cancel(TA, TX, U)$: In this transaction, given full disclosure of TX the TA can cancel the entire ticket, spending all tokens and optionally paying the remaining balance of the card to some entity U . As previously mentioned, the user of the HPD is assumed to have the ability to back up any tickets stored on their HPD. Additionally, the TA may offer a mechanism for secure backup of traditional transponders. A user may reclaim the remaining balance of a lost card by performing a *cancel* on the backup of the ticket. *cancel* is also the transaction that the TA uses to destroy tickets which are identified by fraud detection.

6 A Design for Anonymous Transit System Payment

We consider two kinds of tickets: the passive RFID transponder and the embedded system such as the cell phone or PDA. We refer to the latter as a High Powered computing Device (HPD) in order to distinguish it from the passively powered transponder. We choose these two kinds of tickets because of their wide-deployment and non-trivial security and privacy properties. A HPD with an RF transmitter can follow the same protocols as the passive transponder. Thus, transit systems with existing RFID deployments can implement our hybrid design without requiring separate faregate hardware for each technology.

In addition to traditional HPD features, such as increased security through PIN or biometric user authentication [13], we assume an HPD can be backed up and restored by a user to some external storage (much like common PDA synchronization). We also assume that HPDs exist for which users can observe, debug, and modify the programming. Such devices are important because they provide the basis of assurance of detection of certain kinds of adversarial action on the part of the transit system. An open HPD platform allows interested users to monitor the transit system to observe that it follows its stated protocols and does not, for example, charge too little or too much for a particular transaction. HPD properties specific to transit systems include the ability to kill a transfer immediately after receipt if the user knows that they will not be transferring, and the ability to report a spurious balance in a transit system with protocols requiring the ticket’s remaining balance to be disclosed. Note that this latter property increases not only the privacy of the HPD user, but that of the passive ticket as well.

For simplicity we assume that HPDs and passive transponders are indistinguishable to TA with respect to communications. We defer examination of this assumption to future work as there are many arguments both for it and against it under different circumstances.

In the remainder of this section, we will discuss the details of our design for transit system payment using e-cash and anonymous credentials and consider the security and privacy implications of each transaction.

6.1 Authentication and Session Key Creation

We propose Re-Encryption based Authentication (REA); a novel method for authentication of an authorized reader to a ticket. This method is secure and well matched to the computation and storage resources of the various computers involved in a transit system. The burden of key management is shouldered by the main transit authority computers, which is appropriate since they have the least resource constraints. In our system the TA must daily generate delegation keys which are only good for that day, and then distribute them to each authorized reader. Revocation of a reader is accomplished by simply failing to issue that reader a delegation key for a new day.

Possession of a non-expired delegation key permits F to re-encrypt messages from TX according to the protocol depicted in Figure 2. TX can accomplish authentication and negotiation of a session with only a single public key operation. In this system, TX only needs to store a single public key. This is appropriate to the transponder’s storage constraints, and is superior to a system that would require TX to store revocation information.

In the challenge-response protocol, the authorized reader F , sends the current time t to TX . This step is necessary since some passive RFID transponders are not able to support real time clocks. This time is then concatenated together with a random number of length l_n (which is a security parameter) to form the session key S , which is then transmitted to F encrypted with TA ’s public key (it is reasonable to assume that K_{TA}^+ is built into TX at the factory). F




Fig. 2. Authentication of reader to ticket using re-encryption (RE) allows F to translate ciphertext encrypted with K_{TA}^+ to ciphertext which can be decrypted with K_F^- . Thus the private key of TA remains offline. This re-encryption can only happen if F possesses an appropriate non-expired delegation key. Proof of possession of this delegation key is the mechanism by which F demonstrates that it is authorized. This protocol provides a secure channel while matching the resource constraints of the different devices.

demonstrates that it is authorized by using its delegation key to transform C into a form which it can then decrypt with its own private key. The fact that F is then able to reply to TX with a well-formed message encrypted with session-key S demonstrates that F is authorized (possesses a non-expired delegation key).

Once TX is satisfied that it is talking to an authorized reader it updates its logical clock to value t . If it ever receives a communication with a timestamp less than t , the communication will be assumed to be adversarial, and the protocol will be aborted. TX also uses t to refuse to divulge any information about cookies it holds which have expired.

Since t increases monotonically (which can be monitored by HPDs, and discrepancies will also be eventually caught by passive transponders) and r is chosen by TX , neither F nor TX can cheat at this protocol in such a way as to make a re-play attack possible. S can only be decrypted by a reader with an unexpired delegation key (up to the strength of the underlying public-key and re-signature cryptosystems). This suffices for the security (up to underlying primitives) of the challenge-response.

6.2 $createTicket(TV, TX, U) \rightarrow TX$

Once the session key S is negotiated as discussed above, a stored-value ticket can be created by calling $CreateTokens(TV, TX, \nu)$ resulting in a new wallet which is then stored on TX . The protocol for creation of a temporally bounded ticket is similar, except that in place of $CreateTokens$, $FormNym$ and $GrantCred$ must be executed with respect to some time interval λ which the user has chosen

and purchased. We assume the existence of some function which maps from t to a particular credential expiration epoch $\tau \in \lambda$.

As with all of our other protocols, transactions such as *CreateTokens*, *FormNym*, and *GrantCred* are protected by the session key S , thus preventing the various attacks of middleman adversary M . We omit the means by which payment is proven to TA since solutions to this problem are so well understood.

Privacy of ticket creation: We assume the worst case scenario in which a user purchases or adds value with a credit card, making this a fully-identifying transaction. In this case TV (and by extension TA) gains knowledge of a tuple (U, t, TV, ν) , where ν is initial balance of TX . We will consider the privacy of the rest of the system in terms of how difficult it is to link future transactions with this initial transaction.

6.3 $enter(F, TX) \rightarrow C_E(F)$

Upon entrance to the transit system, TX must either prove that it possesses an unexpired credential using $VerifyCred(N, \tau, F)$, or if it is a stored-value card it must accept an entrance cookie C_E . In our system this cookie is a one-show credential parameterized by the identity of the station at which we enter the system L . The cookie is formed with a call to *FormNym* and then *GrantCred*.

Privacy of *enter*: In the case of the temporally bounded ticket, the only information that is revealed by TX is the possession of a valid credential which is not expired for the present day. TA could attempt to learn TX 's λ by running this transaction multiple times with increasing values of t , but since TX 's logical clock increases monotonically this would have the effect of destroying the ticket, thus preventing this attack from being used for fingerprinting and tracking TX .

In the case of the stored-value ticket, no meaningful information is revealed during this transaction. It is worthy to note, however, that TX could be tricked into carrying an identifying cookie if an adversarial TA could provide a unique L with each transaction. In order to detect such adversarial behavior, HPDs should carry a table mapping from station id L to station name. The HPD software and user can easily detect invalid values of L . Since TA cannot distinguish between an HPD and a passive transponder, this is another example of how HPDs in a hybrid system offer benefits to users of lower cost passive tickets.

6.4 $exit(F, TX, C_E) \rightarrow C_T(F, e)$

In the case where TX is a temporally bounded ticket, exit from the system is exactly the same as entrance.

In the case where TX is a stored-value ticket, the entrance cookie C_E must now be revealed. TX deletes C_E from its memory as soon as its validity has been proven, and therefore avoids being tracked in the future based on any property of C_E . C_E gives TA knowledge of the location at which TX entered the system, so TA can now calculate the cost of the trip. This cost is transmitted to TX and is paid according to the *SpendTokens* primitive. As a convenience to the user,

TX transmits its remaining balance to the faregate so that it may be displayed to the user. After all of this occurs, a transfer cookie C_T is negotiated between TX and TA . The creation of C_T proceeds the same way as with C_E , except that C_T is parameterized by both an expiration time e as well as the exit station identity.

Privacy of *exit*: For simplicity of argument, we shall strengthen our adversary by assuming that TA can calculate with perfect knowledge the amount of time it would take for TX to move from its point of purchase to the exit faregate F .

If the system is to provide the convenience of displaying remaining balance at the faregate, the user of the passive transponder must necessarily lose a certain degree of anonymity. We will consider the size of TX 's anonymity set. Since the station of entry is disclosed by our protocol, let set α be the set of all tickets which if they had travelled here from their station of entry would be arriving now (at time t). Let set β be the set of all tickets which have the same balance as TX and that were purchased within a time interval such that they could just now be arriving at the current faregate F . Let set γ be the set of all tickets which have the same balance as TX but that were not purchased within a time interval such that they could be just now arriving at the current faregate F . With respect to the information possessed by TA , TX 's anonymity set is then $(\alpha \cap \beta) \cup \gamma$. Assuming that the fare structure is set up such that there is a reasonable distribution of possible balances, this anonymity set should be of acceptable size.

It is here that users of HPDs may offer greater privacy to users of traditional transponders. HPDs are capable of displaying remaining balance on their own screens, therefore the HPD user does not need to see the balance on the screen of the faregate. Consequently, HPDs could be programmed to either transmit a random "balance", or could even intentionally choose a "balance" with probability inversely proportional to that of the expected real system balance distribution. With many such HPDs in a system, and with no way of knowing which tickets are reporting false balances, the task of correlating exit balances with identified ticket purchases becomes quite challenging. Let set δ be the set of tickets falsely reporting the same balance as TX , then the new anonymity set of the traditional transponder is $(\alpha \cap \beta) \cup \gamma \cup \delta$. At the same time, this is advantageous to the user of the HPD who now enjoys greater anonymity since they can no longer be distinguished by their balance.

6.5 $transfer(F, C_T) \rightarrow C_F(F)$

For temporally bounded tickets, this transaction is the same as *enter*.

Stored-value tickets begin by proving that their transfer is unexpired, and is from an exit station in the set of stations which may transfer here. If these things are true, then a new type of transfer C_F is minted through the same means as C_E , except that C_F is distinguishable from an entry cookie. When the user exits at the final destination, TA can now compute a balance discounted according to the transit system's transfer rules. TX always deletes cookies from its memory as soon as they have been verified by TA .

Privacy of *transfer*: Some information about both the station at which C_T was minted, and the time epoch τ of genesis are revealed during the transfer verification. The size of the anonymity set will be the number of tickets issued during τ which are valid at the verifying station. Since TX will only agree to verify C_T once, an adversary cannot test different values of τ and F to reduce the size of TX 's anonymity set. At this point in the system, it would be quite difficult indeed to trace TX back to its original purchase. Although the anonymity at each intermediate step is less than total, in a transit system with a reasonable passenger volume tracking a particular user quickly becomes infeasible.

6.6 *addValue(TV, TX, U) → TX'*

In the case of a temporally bounded TX , the remaining time on the card (λ) can be determined by verifying the card's credential for increasing t until the expiration date is found. At this point, of course, the card has been destroyed, and must be re-initialized with a brand new Nym and Credential for a new, longer time period. This is accomplished as in *createTicket*.

The procedure for the stored-value TX is quite similar: The remaining tokens are spent via *SpendToken*, and a new wallet is created for this value plus whatever new value the user has purchased.

Privacy of *addValue*: In the worst case, the user will choose to refill their ticket using a credit card. In this case, of course, this is a fully identifying transaction. We will consider the case where the user pays for additional value through some anonymous means.

For the temporally bounded TX , the size of the anonymity set in this transaction will be the number of tickets in circulation with the same remaining λ . For most transit systems it seems likely that the more distinguishing λ values would be the longest such values (year-long passes and such). Fortunately it is unlikely that the owner of a ticket would desire to add time to a ticket which already has a great deal of time remaining.

There is an attack on the anonymity of the temporally bounded TX in which an adversarial TA reduces the size of TX 's anonymity set by spuriously executing an *addValue* transaction in order to fingerprint TX 's λ , and then creates a new Nym and Credential for TX . In our system, TX has a number of protections against this attack. First of all, such spurious transactions can clearly be detected and reported by HPDs. Secondly, the passive transponder can keep track of the most recent λ and refuse to accept a new λ that is not at least one day greater than the old. This latter defense would mean that an adversarial TA would have to give a free extra day for every fingerprint, and the fingerprint would become increasingly meaningless as the value of λ diverged further and further from the identifying purchase.

The only information that the TA gains from the stored-value TX is the remaining balance. It should be noted that this is a circumstance where the HPD cannot give a false balance, as the balance is checked by actually spending the remaining tokens. The anonymity set of TX during this transaction will be the number of TX in circulation with the same balance. Fortunately it is likely

that users will add value to their tickets only when the ticket runs low. In this case the ticket will be likely to have a common balance.

6.7 *cancel(TA, TX, U)*

Given all of the information on TX , TA may cancel the ticket, and any clones it may have. In the case of a stored-value ticket, TA executes *spendToken* on all remaining tokens in TX 's wallet, optionally reimbursing the user, if it is the user who is canceling the ticket (rather than the TA choosing to cancel due to detected fraud).

In the case of the temporally bounded ticket, TA optionally determines the remaining value by the same mechanism as in the *addValue* transaction, and then executes *RevokeCred* which will cancel TX and any clones thereof.

7 Alternative Approaches

If a transit authority chooses passive transponders which lack sufficient resources for the primary design outlined above, there are many alternatives to be considered which are much cheaper to implement, but do not provide the same strength of anonymity as our primary design.

Entry (C_E) and transfer (C_T) cookies can be realized with no processing required from the transponder above the cost of authenticating the reader. The TA can compute and transmit $C_E := M_{K_{TA}^-}(S)$, where S is the session-key containing a timestamp and a nonce. The TA can store the cookie along with the identity of the entry station. Note that the cookie can be signed by the reader using a re-signature key. For protection against tracking, the ticket will only disclose a cookie once, and thus can be tracked only by the TA and only for the duration of a single trip. This cookie design provides privacy, in that there is still no way to link a ticket to its fully identifying purchase, however it clearly does not meet the anonymity goals outlined in our payment semantics.

Another alternative design provides temporally bounded tickets through the same mechanisms as the stored value tickets. In this design, day passes are created with a large quantity of valueless e-cash tokens parameterized with an expiration date. The tokens are essentially used as single-show credentials. This system falls short of the true semantics of temporally bounded tickets in that the ticket may not be used an unbounded number of times during its period of validity. However, it may be cheaper to implement a system based on one primary underlying set of cryptographic protocols.

8 Future Work

Before a transit payment system is ready for deployment, its various components should individually and in aggregate be stated in formal notation with a clear security model and proofs of security within that model. We hope that future

work will consider specific cryptosystems within the semantic framework that we have proposed here, and will provide appropriate proofs up to the assumptions supported by the chosen cryptographic schemes.

A problem clearly exists with maintaining an anonymous credential on a virtual card which the user has full ability to read and modify. In the naïve system there is little to prevent a dishonest user from selling many copies of a valid credential, and this would be quite difficult to detect given the anonymous nature of the credentials. Some work has been done on fraud detection in simultaneous use of the cloned credentials, which is quite suitable for other problem domains such as online game licensing [14]. In our problem domain, however, many kinds of fraud will go undetected with high probability. We are investigating novel mechanisms for cloning detection in anonymous credential systems, which we hope will offer a solution to this currently open problem.

9 Conclusion

We have (1) demonstrated that transit systems are an important problem domain for the study of security and privacy, (2) presented a framework for formal consideration of transit system ticketing, and (3) provided designs sufficient for implementation of a secure privacy preserving transit system. Our approach uses e-cash, anonymous credentials, and proxy re-encryption to increase passenger privacy without compromising the secure payment requirements of the transit authority. Yet many theoretical and practical challenges remain for further study in how to balance the privacy concerns of passengers with the security needs of transit authorities.

10 Acknowledgments

We are grateful to Anna Lysyanskaya for discussions about cloning detection for anonymous credentials systems, and to Ben Adida and Susan Hohenberger for discussion about the feasibility of re-encryption based authentication. For proofreading and presentation advice we thank Jeremy Barth, Ed Costello, Marc Liberatore, and Boris Margolin. Finally, we thank our anonymous reviewers for their suggestions and encouragement.

References

1. Anderson, R., Kuhn, M.: Tamper Resistance - A Cautionary Note. In: The Second USENIX Workshop on Electronic Commerce Proceedings. (1996) 1–11
2. Ateniese, G., Camenisch, J., de Medeiros, B.: Untraceable RFID Tags via Insubvertible Encryption. In: Conference on Computer and Communications Security – CCS’05, Alexandria, Virginia, USA, ACM, ACM Press (2005)
3. Ateniese, G., Fu, K., Green, M., Hohenberger, S.: Improved Proxy Re-Encryption Schemes with Applications to Secure Distributed Storage. In: Proceedings of the 12th Annual Network and Distributed System Security Symposium (NDSS). (2005)

4. Ateniese, G., Hohenberger, S.: Proxy Re-Signatures: New Definitions, Algorithms, and Applications. In: Proceedings of the 12th ACM conference on Computer and communications security (CCS '05), Alexandria, VA, USA, ACM, ACM Press (2005) 310–319
5. Attoh-Okine, N., Shen, L.: Security Issues of Emerging Smart Cards Fare Collection Application in Mass Transit. In: Vehicle Navigation and Information Systems Conference. (1995) 523–526
6. Avoine, G., Oechslin, P.: RFID Traceability: A Multilayer Problem. In Patrick, A., Yung, M., eds.: Financial Cryptography – FC'05. Volume 3570 of Lecture Notes in Computer Science., Roseau, The Commonwealth Of Dominica, IFCA, Springer-Verlag (2005) 125–140
7. Bono, S., Green, M., Stubblefield, A., Juels, A., Rubin, A., Szydlo, M.: Security Analysis of a Cryptographically-Enabled RFID Device. In: USENIX Security Symposium, Baltimore, Maryland, USA, USENIX (2005) 1–16
8. Burgess, J., Gallagher, B., Jensen, D., Levine, B.: Maxprop: Routing for vehicle-based disruption-tolerant networks. In: Proc. IEEE INFOCOM. (2006)
9. Camenisch, J., Hohenberger, S., Lysyanskaya, A.: Compact E-Cash. In: EUROCRYPT, Aarhus, Denmark, IACR (2005) 302–321
10. Camenisch, J., Lysyanskaya, A.: An Efficient System for Non-transferable Anonymous Credentials with Optional Anonymity Revocation. In: EUROCRYPT, Innsbruck(Tyrol), Austria, IACR (2001)
11. Camenisch, J., Lysyanskaya, A.: Signature Schemes and Anonymous Credentials from Bilinear Maps. In: CRYPTO, Santa Barbara, CA, USA (2004)
12. Chaum, D., Fiat, A., Naor, M.: Untraceable Electronic Cash. In: CRYPTO '88: Proceedings on Advances in Cryptology, New York, NY, USA, Springer-Verlag New York, Inc. (1990) 319–327
13. Chaum, D.: Security without Identification: Transaction Systems to Make Big Brother Obsolete. CACM **28**(10) (1985)
14. Damgård, I., Dupont, K., Pedersen, M.Ø.: Unclonable Group Identification. Cryptology ePrint Archive, Report 2005/170 (2005) <http://eprint.iacr.org/>.
15. Dimitriou, T.: A Lightweight RFID Protocol to Protect Against Traceability and Cloning Attacks. In: Conference on Security and Privacy for Emerging Areas in Communication Networks – SecureComm, Athens, Greece, IEEE (2005)
16. Federal Transit Administration: Federal Transit Administration National Transit Database. WWW (2006) <http://www.ntdprogram.com>.
17. Guerineau, P.: Active RFID Technology Applied to Security Improvement and Statistical Control in Public Transit. In: Automatic Fare Collection. New Horizons in Public Transport with Smart Cards, Brussels, Belgium, International Union of Public Transport (2002)
18. Handschuh, H., Paillier, P.: Smart Card Crypto-Coprocessors for Public Key Cryptography. In Quisquater, J.J., Schneier, B., eds.: Smart Card Research and Applications, SPLNCS. Volume 1820. (2000) 386–394
19. Juels, A., Rivest, R., Szydlo, M.: The Blocker Tag: Selective Blocking of RFID Tags for Consumer Privacy. In Atluri, V., ed.: 8th ACM Conference on Computer and Communications Security. (2003) 103–111
20. Juels, A.: Minimalist Cryptography for Low-Cost RFID Tags. In Blundo, C., Cimato, S., eds.: International Conference on Security in Communication Networks – SCN 2004. Volume 3352 of Lecture Notes in Computer Science., Amalfi, Italia, Springer-Verlag (2004) 149–164

21. Juels, A., Molnar, D., Wagner, D.: Security and Privacy Issues in E-passports. In: Conference on Security and Privacy for Emerging Areas in Communication Networks – SecureComm, Athens, Greece, IEEE (2005)
22. Juels, A., Syverson, P., Bailey, D.: High-Power Proxies for Enhancing RFID Privacy and Utility. In: Proceedings of Privacy Enhancing Technologies workshop (PET 2005). (2005)
23. Kang, J., Nyang, D.: RFID Authentication Protocol with Strong Resistance Against Traceability and Denial of Service Attacks. In Molva, R., Tsudik, G., Westhoff, D., eds.: European Workshop on Security and Privacy in Ad hoc and Sensor Networks – ESAS’05. Volume 3813 of Lecture Notes in Computer Science., Visegrad, Hungary, Springer-Verlag (2005) 164–175
24. Maxey, C., Benjamin, P.: Seamless Fare Collection: Using Smart Cards for Multiple-Mode Transit Trips. WWW (2006) www.apta.com/research/info/briefings/documents/maxey.pdf.
25. McDaniel, T.L., Haendler, F.: Advanced RF Cards for Fare Collection. In: Commercial Applications and Dual-Use Technology Conference Proceedings, National Telesystems Conference (1993) 31–35
26. Mohammed, E., Emarah, A., El-Shennawy, K.: Elliptic Curve Cryptosystems on Smart Cards. In: SEC ’02: Proceedings of the IFIP TC11 17th International Conference on Information Security, Deventer, The Netherlands, The Netherlands, Kluwer, B.V. (2002) 311–322
27. Molnar, D., Wagner, D.: Privacy and Security in Library RFID: Issues, Practices, and Architectures. In Pfitzmann, B., Liu, P., eds.: Conference on Computer and Communications Security – ACM CCS, Washington DC, USA, ACM, ACM Press (2004) 210–219
28. Neve, M., Peeters, E., Samyde, D., Quisquater, J.J.: Memories: A Survey of Their Secure Uses in Smart Cards. In: IEEE Security in Storage Workshop. (2003) 62–72
29. Poupard, G., Stern, J.: On the Fly Signatures Based on Factoring. In: CCS ’99: Proceedings of the 6th ACM conference on Computer and communications security, New York, NY, USA, ACM Press (1999) 37–45
30. Ranasinghe, D., Engels, D., Cole, P.: Low-Cost RFID Systems: Confronting Security and Privacy. In: Auto-ID Labs Research Workshop, Zurich, Switzerland (2004)
31. Roschke, G.: Notes from an Information Law Student. WWW (2006) Last viewed February 24, 2006, <http://luminousvoid.net/archives/16/wmata-responds>.
32. San Francisco Bay Area Rapid Transit District: Bay Area Rapid Transit (BART) Fiscal Year 2004 Annual Report. WWW (2006)
33. Sarma, S., Weis, S., Engels, D.: Radio-Frequency Identification: Security Risks and Challenges. *Cryptobytes*, RSA Laboratories **6**(1) (2003) 2–9
34. Sim, L., Seow, E., Prakasam, S.: Implementing an Enhanced Integrated Fare System for Singapore. *Public Transport International* **53**(6) (2004) 34–37
35. The Smart Card Alliance: Hong Kong Octopus Card. WWW (2006) http://www.smarcardalliance.org/pdf/about_alliance/user_profiles/Hong_Kong_Octopus_Card.pdf.
36. The Smart Card Alliance: Smart Card Talk Standards. The Smart Card Alliance Newsletter (2006) Jan. issue.
37. Trichina, E., Bucci, M., Seta, D.D., Luzzi, R.: Supplemental Cryptographic Hardware for Smart Cards. *IEEE Micro* **21**(6) (2001) 26–35
38. Vajda, I., Buttyán, L.: Lightweight Authentication Protocols for Low-Cost RFID Tags. In: Second Workshop on Security in Ubiquitous Computing – Ubicomp 2003, Seattle, WA, USA (2003)

39. Washington Metropolitan Area Transit Authority: WMATA Privacy Policy. WWW (2006) http://www.wmata.com/about/parp_docs/pi.9.2.0.pdf.
40. Washington Metropolitan Area Transit Authority: WMATA Privacy Policy Proposal. WWW (2006) <http://www.wmata.com/about/parp2.cfm>.
41. Winters, N.: Personal Privacy and Popular Ubiquitous Technology. In: Ubiconf, London, United Kingdom (2004)
42. Zhao, W., Ammar, M., Zegura, E.: A Message Ferrying Approach for Data Delivery in Sparse Mobile Ad Hoc Networks. In: MobiHoc '04: Proceedings of the 5th ACM international symposium on Mobile ad hoc networking and computing, New York, NY, USA, ACM Press (2004) 187–198
43. Zhao, W., Ammar, M.H.: Message Ferrying: Proactive Routing in Highly-Partitioned Wireless Ad Hoc Networks. In: FTDCS '03: Proceedings of the The Ninth IEEE Workshop on Future Trends of Distributed Computing Systems (FTDCS'03), Washington, DC, USA, IEEE Computer Society (2003) 308