Countering Identity Theft through Digital Uniqueness, Location Cross-Checking, and Funneling*

P.C. van Oorschot[1] S. Stubblebine[2]

[1] School of Computer Science, Carleton University, Ottawa, Canada
[2] Stubblebine Research Labs, Madison, NJ, USA

Abstract. One of today's fastest growing crimes is identity theft – the unauthorized use and exploitation of another individual's identity-corroborating information. It is exacerbated by the availability of personal information on the Internet. Published research proposing technical solutions is sparse. In this paper, we identify some underlying problems facilitating identity theft. To address the problem of identity theft and the use of stolen or forged credentials, we propose an authentication architecture and system combining a physical location cross-check, a method for assuring uniqueness of location claims, and a centralized verification process. We suggest that this system merits consideration for practical use, and hope it serves to stimulate within the security research community, further discussion of technical solutions to the problem of identity theft.

1 Introduction and Motivation

Identity theft is the unauthorized use and exploitation of another individual's identity-corroborating information (e.g. name, home address, phone number, social security number, bank account numbers, etc.). Such information allows criminal activities such as fraudulently obtaining new identity credentials, credit cards or loans; opening new bank accounts in the stolen name; and taking over existing accounts. It is one of today's fastest growing crimes. In one Canadian incident reported in April 2004 [13], a single identity theft involving real estate lead to a $540,000 loss. In 2002, reportedly 3.2 million Americans suffered an identity theft which resulted in new bank accounts or loans [1]. The severity of the problem has resulted in a recent U.S. law – the "Identity Theft Penalty Enhancement Act" – boosting criminal penalties for phishing (see below) and other identity fraud ([29]; see also [26]).

* Version: 14 Jan. 2005. To appear in pre-proceedings of *Financial Cryptography'05*, Feb.28–Mar.3, 2005. Author addresses: paulv@scs.carleton.ca, stuart@stubblebine.com

Despite growing media attention and numerous web sites (government-sponsored[3] and other) discussing the problem, its seriousnous continues to be under-estimated by most people other than those who have been victimized. The (U.S.) Identity Theft and Assumption Deterrence Act of 1998 specifically made the actual theft of another's identifying information a federal crime, and created the Identity Theft Clearinghouse Database – a central database of identity theft complaints used by law enforcement. In contrast to this generally reactive tool (which helps victims after-the-fact), the most common preventative measure to date seems to involve efforts to educate the public to more carefully guard personal information. In the research literature to date, there appear to be few effective technical solutions or practical proposals (see below and in §2), none of which to our knowledge have been adopted successfully to the point of decreasing identity thefts in practice.

"Activity profiling" by credit card companies – a form of anomaly detection in customer usage of a credit card – partially addresses the problem of stolen or fraudulent credit cards, but not that of identity theft itself. While consumers have limited liability on use of fraudulent credit cards in their name, protection by credit card companies is limited to the realm of credit cards (see next paragraph). Regarding protection afforded by banks, in the U.S. (but reportedly not Canada), when one major bank puts an alert on a name, a common clearinghouse (limited to banks) allows all major banks to share that warning [17].

Unfortunately, identity theft appears to be a system-level problem that no one really "owns", and thus it is unclear whose responsibility it is to solve. Sadly, individual citizens are poorly positioned to solve this problem on their own, despite being the victims suffering the most in terms of disrupted lives, frustration and lost time to undo the damage – especially when stolen identity information is used to mint new forms of identity-corroborating information (or e.g. new credit cards) unbeknownst to the legitimate name-owner. According to one 2003 report [1], victims averaged 60 hours "to resolve the problem" of an identity theft, e.g. getting government and commercial organizations to stop recognizing stolen identification information, and to re-issue new identity information.

Among those perhaps in the best position to address identity theft are the national consumer credit reporting agencies – e.g. in the U.S., Equifax, Experian, and Trans Union. Among other things, the credit bureaus can when necessary post alerts on credit files of individuals whom they suspect are subjects of identity theft [17]. However, it is unclear how strongly

[3] For example, see http://www.consumer.gov/idtheft/

the business models of credit bureaus motivate them to aggressively address the problem, and surprisingly some have reportedly opposed certain measures which aid in identity theft prevention (e.g. see [1]). Moreover, at least one such organization[4] was itself exploited by criminals in an incident raising fears of large-scale identity theft.

Phishing[5] is a relatively new Internet-based attack used to carry out identity theft. "Phishing kits" now available on the Internet allow even amateurs to create bogus websites and use spamming software to defraud users [32]. A typical phishing attack involves email sent to a list of target victims, encouraging users to visit a major online banking site. By chance a fraction of targeted users actually hold an account at the legitimate site. However the advertised link is to a spoofed site, which prompts users to enter a userid and password. Many legitimate users do so immediately, thereby falling victim. This is a variation of an attack long-known to computer scientists, whereby malicious software planted on a user's machine puts up a fraudulent login interface to obtain the user's userid and login password to an account or application.

Key logging attacks now rival phishing attacks as a serious concern related to online identity and sensitive personal information theft [19]. A recent example involved a trojan program *Bankhook.A* which spread without human interaction beyond web browsing, involved a (non-graphic) file named *img1big.gif*, and exploited a vulnerability in a very widely used web browser. Upon detecting attempted connections to any of about 50 major online banks,[6] it recorded sensitive information (e.g. account userid and password) prior to SSL encryption, and mailed that data to a remote computer [28, 22].

Our contributions. We identify underlying problems facilitating identity theft, and propose a general authentication architecture and system we believe will significantly reduce identity theft in practice. The system combines a physical location cross-check, a method for assuring uniqueness of location claims, and a centralized verification process. We outline how the system prevents a number of potential attacks. We propose an extension addressing the problem of acquiring fraudulent new identity credentials from stolen credentials. A major objective is to stimulate further research and discussion of technical solutions to the "whole"

[4] Equifax Canada recently confirmed that in February 2004, 1400 consumer credit reports were "accessed by criminals posing as legitimate credit grantors" [16, 17].

[5] See http://www.ftc.gov/bcp/conline/pubs/alerts/phishingalrt.htm

[6] Text string searches were made for https connection attempts to URLs containing any of 50 target substrings. See Handler's log (June 29, 2004) at http://isc.sans.org/presentations/banking_malware.pdf.

problem of identity theft (rather than subsets thereof – e.g. phishing and key-logging).

Organization. The sequel is organized as follows. §2 discusses further background and related work. §3 presents an overview of our proposed authentication system and architecture for addressing identity theft, a security analysis considering some potential attacks, and a discussion of preventing privacy loss due to location-tracking. §6 gives concluding remarks.

2 Fundamentals and Related Work

We first discuss credentials, then identify what we see as the fundamental issues facilitating identity theft, thereafter mention a relationship to issues arising in PKI systems, and finally review related work.

Credentials. We define *identity credentials* (*credentials*) rather loosely as "things" generally accepted by verifiers to corroborate another individual's identity. By this definition, a credential may be digital (such as userid-password, or public-key certificate and matching private key) or physical (e.g. physical driver's license, plastic credit card, hardware token including secret key). The looseness arises from situations like the following: the secret key within a hardware token is extracted, and as the key itself is then digital, essentially the important component of the physical token in now available in digital form – which we also call *credential information.* A further looseness is that unfortunately some pieces of information, such as (U.S.) Social Security Number, are used by some parties as identity-corroborating data, even if provided verbally (rather than physical inspection of a paper or plastic card) – even though they are not generally treated as secret.

Fundamental underlying problems. There are numerous reasons why personal identities and credential information are so easily stolen, and why this is so difficult to resolve. We believe the fundamental problems facilitating identity theft include the following.

F1: *ease of duplication:* the ease of duplicating personal data and credentials;

F2: *difficulty of detecting duplication:* the difficulty of detecting when a copy of a credential or credential information is made or exists (cf. [18]);[7] and

[7] Thus one cannot tell when an identity theft occurs. Often copies of identity information are made, used elsewhere, and detected later only after considerable damage has occurred.

F3: *independence of new credentials:* if existing credential information is used by an impersonator to obtain new credentials, the latter are in one sense "owned" by the impersonator, and usually no information flows back to the original credential owner immediately.

In particular due to F3, we see identity theft as a *systemic* problem, which cannot be solved by any single credential-granting organization in isolation. Regarding F2, a *copy* of a cryptographic key is digital data; a copy of a physical credential is another physical object which a verifier might accept as the original.

Identity theft is also facilitated by the availability of personal information (and even full credentials, e.g. stored at servers) on the Internet; and the ease with which many merchants grant credit to new customers without proper verification of identification. While we focus on the theft of credential *information*, the theft of actual physical credentials (e.g. authentic credit cards) is also a concern – but one more easily detected.

Relationship to PKI systems. We note there are similarities between detecting the theft and usage of password-based credentials and that of signature private keys; indeed, passwords and signature private keys are both secrets, and ideally in both cases, some form of theft checkpoint would exist at the time of verification. More generally, issues similar to those arising in identity theft arise in certificate validation within public key infrastructure (PKI) systems – most specifically, the revocation of private keys. There is much debate in practice and in academic research about revocation mechanisms, and which are best or even adequate. This has lead to several *online status checking* proposals (e.g. OCSP [27] and SCVP [25]), to counter latency concerns in offline models. This suggests looking to recent PKI research for ideas useful in addressing identity theft (and vice versa). As a related result, we cite the *CAP principle* [8, 10]: a large-scale distributed system can essentially have at most two of the following three properties: high service availability; strong data consistency; and tolerance of network partitions.

Related work. The U.S. Federal Communications Commission (FCC) requires[8] that by Dec. 31 2005, wireless carriers report precise location information (e.g. within 100 meters) of wireless emergency 911 callers, allowing automatic display of address information on 911 call center phones, as presently occurs for wireline phones. Companies must either use GPS in 95% of their cell phones by Dec. 31 2005, or deploy other location-tracking technology (e.g. triangulation or location determination based

[8] See http://www.fcc.gov/911/enhanced/ (see also [9]).

on distance and direction from base stations); thereafter emergency call centers must deploy related technology to physically locate callers. As of Feb. 2004, 18% of U.S. call centers have this technology [30].

While many technologies and systems exist for determining the physical location of objects, these generally are not designed to operate in a malicious environment – e.g. see the survey by Hightower and Borriello [14]. Sastry et al. [31] propose a solution to the *in-region verification problem* of a verifier checking that a claimant is within the claimed specified region. This differs from the more difficult *secure location determination problem* involving a verifier determining the physical location of a claimant. Gabber and Wool [9] discuss four schemes, all based on available infrastructure, for detecting the movement of user equipment; they include discussion of attacks on these systems, and note that successful cloning, if carried out, would defeat all four. All of the above references address a problem other than identity theft per se, where complicating matters include the minting of new credentials (see F3 above) and uniqueness of a claimant with the claimed identity; the binding of location information to a claimed identity is also critical.

Physical location has long been proposed as a fourth basis on which to build authentication mechanisms, beyond the standard "something you know, something you have, something you are". In 1996, Denning and MacDoran [6] outlined a commercial location-based authentication system using the Global Positioning System (GPS), notwithstanding standard GPS signals being subject to spoofing (e.g. see [9, 33]. Their system did not seek to address the identity theft problem – for example regarding F2, note that in general, location information alone does not guarantee uniqueness (e.g. a cloned object may claim a different physical location than the original object); F3 is also not addressed.

One real-world system-level technique to ameliorate identity-theft is the *credit-check freeze* solution [1],[9] now available in many U.S. states. An individual can place a "fraud alert" on their credit reports, thereby blocking access to it by others for a fixed period of time, or until the individual contacts the credit bureaus and provides previously agreed information (e.g. a PIN). Another option is selective access, whereby a frozen report can be accessed only by a specifically named inquirer. These methods apparently prevent identity thieves from getting (new) credit in a victim's name, or opening new accounts thereunder, but again do not solve the problem of identity theft (e.g. recall F3 above).

[9] See also http://www.ftc.gov/bcp/conline/pubs/general/idtheftfact.htm

Corner and Noble [3] propose a mechanism involving a cryptographic token which communicates over a short-range wireless link, providing access control (e.g. authentication or decryption capabilities) to a local computing device without user interaction. While not proposed as a solution to identity theft per se, this type of solution offers an innovative alternative to easily replicated digital authentication credentials – simultaneously increasing security and decreasing user interaction (e.g. vs. standard password login).

Chou et al. [2] proposed a client-side software plug-in and various heuristics for detecting online phishing scams. Lu and Ali [24] discuss using network smart cards to encrypt sensitive data for remote nodes prior to its availability to local key-logging software.

3 Authentication based on Uniqueness, Location and Funneling

A high-level overview of our proposed authentication system is given in §3.1 . A partial security analysis is given in §3.2. Privacy refinements are discussed in §5.

3.1 High-level Overview of New System

Our goal is a system which prevents, or significantly reduces, occurrences of identity theft in practice. Our design is as follows. Every system user has a hardware-based *personal device*,[10] e.g. cell phone or wireless personal digital assistant (PDA), kept on or near their person, and which can be used to securely detect their location[11] and securely map the person to a location, ideally on a continuous basis. We call this a *heartbeat locator*, perhaps initially simply based on existing infrastructure such as emergency wireless 911 technology (see §2).

Note that in many cases, if someone has your identification credentials, or a reasonable copy thereof, for all intents and purposes they *are* you from the viewpoint of a verifier. We therefore must address both credential theft and cloning. To address cloning, one general solution is

[10] Here "personal" implies that the device be able to identify (or can be associated with) a unique individual.

[11] By *securely detecting location* we mean: the detected location cannot easily be spoofed. In particular, if person P_A is factually at location L_A, then it must be very difficult (ideally infeasible in practice) for an attacker to arrange that a signal is sent indicating that P_A is at a different location $L_B \neq L_A$.

to perform a check (providing reasonably high confidence) that the personal device does in fact remain unique; we call this an *entity uniqueness* mechanism. To aid in this, we require that all identity verifications be *funneled* through a centralized point, allowing a check to be made that no “irregularities” have occurred (based on ongoing device monitoring) for the personal device in question. For discussion of irregularities and more about theft and cloning, see §3.2.

In the process of a transaction being executed/processed, when an identity[12] is simply asserted (or ideally, confirmed by a first means), a secondary confirmation occurs based on the location of the transaction (e.g. merchant's point of sale location) matching the location the central service last recorded for the personal device corresponding to the asserted identity. This can thus be employed as a second-factor authentication system,[13] with the features of (1) combining location determination with continuous location tracking; and (2) funeling all transactions through a single point. This effectively turns an offline or distributed verification system into an online one (cf. §2).

Extension addressing minting of new credentials. We now present a proposal to address issue F3 above (note that *some* such proposal is necessary to fully address identity theft). An extension of the above system is to require that a name-owner give explicit approval before certain actions specifically based on existing identity information – such as the minting of new credential information *not tied to the personal device* – are taken. In practice, a solution might be most effectively put in place by the national credit bureaus as a new service offering, to complement that of freezing access to credit records (see §2). Incoming queries regarding a consumer credit file could be required, by policy, to specify if the inquiry was being used to mint credentials which might reasonably be used as identity credentials by other responsible parties. The major credit bureaus might provide (in a coordinated manner) a central alert-centre to check if such credential minting was currently “allowed” by the legitimate name-owner (e.g. as indicated by a *minting bit* in the existing credit file). Reputable (participating) organizations which created any form of personal creden-

12 An identity per se is not required – e.g. pseudonyms could be used, to enhance privacy (see §5).

13 Again, this is a systemic (multi-application) authentication system addressing identity theft, rather than a second-factor point solution limited to a particular application, such as credit card authorization.

tial would agree[14] to create new credentials only if the response from the centralized service indicated the minting bit was on. In this way, a cautious individual, even without prior identity theft problems, could have minting of new credentials disabled the majority of the time, as a preemptive measure.

3.2 Security Analysis and Discussion

In this section we provide a partial security analysis of the new proposal, and discuss necessary checks regarding the personal device. While we offer no rigorous security arguments here,[15] we discuss a number of attack scenarios and how the system addresses these. We do not "prove" that the proposed system is "secure" in a general practical setting, and believe this would be quite difficult, as "proofs" of security are at best relative to a particular model and assumptions, with increased confidence in the relevance and suitability of these generally gained only over time. However we encourage further analysis to allow the proposal to be iteratively improved.

We begin by referring back to the three fundamental problems of §2. The system proposed in §3.1 addresses these as follows. The ease of credential duplication (F1) is reduced by the use of a hardware device; the capability to detect credential duplication (F2) is provided by the funneling mechanism and ongoing device monitoring (heartbeat mechanism); and the minting of new (fraudulent) credentials based on stolen authentic credentials (F3) is partially[16] addressed by the "minting bit" extension.

Device irregularities, theft and cloning. Fraud mitigation strategies depend on users reporting stolen personal devices in a timely matter.[17] However, some heuristics may also be effective to detect both theft and cloning. Examples of heuristic predictors of cloning include the same personal device appearing multiple times (two heartbeats asserting the same identity, whether at the same or distinct locations), or in two different locations within an unreasonably short period of time (taking into

[14] We recognize that this would require a significant change in behaviour by many organizations, over a long period of time (which legislation might shorten). However, we expect that nothing less will solve the difficult problem of identity theft.

[15] A more complete security analysis will be given in the full paper.

[16] Our proposal does not prevent an attacker from himself forging new credentials; but can prevent the use of stolen credentials to obtain new credentials from an authentic credential-generating organization.

[17] Loaning a personal device to a non-malicious user (e.g. a relative or friend) does not necessarily cause an increase in fraud since those users generally are trusted not to commit fraud using the device.

account usual modes of travel). A heuristic indicator of device theft is a user unable to correctly authenticate even though the location is verifiable (e.g. within range). These are all examples of *irregularities*. In this case, authentication attempts using the device within a short time thereafter may be suspect.

Personal devices flagged as having experienced sufficient irregularities should be disallowed from participating in transactions, or subject to additional checks. As suspicion arises regarding a device (cloning, theft or other misuse), extensions to the basic techniques are possible. For example, the personal device holder might be requested to provide an additional authentication factor to confirm a transaction. In essence, known techniques used for credit card activity profiling, which by system design are currently used only to mitigate credit card fraud, could be adapted to mitigate identity theft in the new system.

Note that a theft deterrent in this system is the risk of physical discovery – device possession allows location-tracking of the thief. Related to this, the deactivation (if featured) and re-activiation of the device's location-tracking feature should also require some means of user authentication, so that a thief cannot disable this feature easily, and if already disabled, the device is unusable for authentication.

Device uniqueness. While ideally the personal device would be difficult to physically duplicate, our proposal only partially relies on this, as duplicate heartbeats will lead to a failed verification check. To enforce device uniqueness, ideally both (1) each device is tracked continuously since registration; and (2) it can be verified that the user originally registering a device remains associated with the tracked device. We may consider the latter issue under the category of theft, and the former under cloning. In practice, monitoring could at best be roughly continuous, e.g. within discrete windows of time, say from sub-second to a minute; we expect this would not pose a significant problem. However there are practical contraints in even roughly monitoring devices – for example, wireless devices are sometimes out of range (e.g. in tunnels, or on airplanes) or turned off. Thus the system must address the situation in which for at least some devices, location-tracking is temporarily disabled. It may be an acceptable risk to allow a device to be "off-air" for a short period of time (e.g. seconds or minutes), provided that it reappears in a reasonably plausible geographic location. Devices "off-air" for a longer period could be required to be re-activated by a user-to-system authentication means (i.e. not user-to-device). Personal devices which have gone "off-air" recently might be given a higher irregularity score, or not be allowed to participate in higher-value transactions (absent additional assurance) for some period of time.

Threats and Potential Attacks. The class of threats we are intending to protect against is essentially the practical world, or more precisely, any plausible real-world attack of "reasonable" cost (relative to the financial gain of the identity theft to the attacker). We consider here a number of potential attacks, and discuss how the system fares against them.

1. *Theft.* If the personal device is stolen or lost, the loss should be reported leading to all further verification checks failing; effectively this is credential revocation. Since often a theft is not immediately noticed or reported, the device should require some explicit user authentication mechanism (such as a user-entered PIN or biometric) as part of any transaction; the device should be shut down upon a small number of incorrect entries (possibly allowing a longer "unblocking PIN" for re-activation).[18]
2. *Cloning.* There can be no absolute certainty that the personal device has not been cloned or mimicked. If a clone exists, either it has a continuous heartbeat (case A), or no heartbeat (case B). In case A, assuming the original device also still has a heartbeat, the system will be receiving two heartbeats with the same device identifier, and flag an irregularity. In case B, if and when the cloned device is used for a transaction, its location will be inconsistent with previous heartbeats (from the legitimate device), and thus the cloned device will be unable to successfully participate in transactions.
3. *Theft, clone, return.* Another potential attack is for a thief to steal a device, clone it (in a tracking de-activated state), then "simultaneously" activate the clone and deactivate the original, and finally return the stolen device. The idea is then to carry out a transaction before the original device owner reactivates or reports the theft. Such an attack, if possible, would nonetheless make identity thefts significantly more difficult than today (and thus our goal would be achieved). A variation has the attacker inject unauthorized software in the original device, to completely control it (including the capability to remotely power it on and off), before returning it. Then at the instance of carrying out a transaction, the attacker remotely powers down the original before powering up the clone, to prevent detection of two heartbeats.

[18] Although a motivated and well-armed attacker can generally defeat user-to-device authentication mechanisms (cf. [9]), we aim to significantly reduce, rather than totally eliminate, occurrences of identity theft. We believe a 100% solution will be not only too expensive or user-unfriendly, but also non-existent.

However a geographic irregularity would arise (as the clone's location would differ from that of the last heartbeat of the real device).

4. *Same-location attack.* An attacker, without possessing a target victim's personal device, might attempt to carry out a transaction at the same physical location (e.g. retail store) as the target victim and that victim's personal device. This attack should be prevented by a requirement that a user take some physical action to commit a transaction (e.g. press a designated key, enter a PIN, or respond to an SMS message). A further refinement is an attacker attempting to carry out a transaction at the same place and the same instant as a legitimate user (and also possessing any other credentials necessary to impersonate the user in the transaction). Here the attacker would be at some physical risk of discovery, and one of the two transactions would go through. While this attack requires further consideration, it appears to be less feasible.

4 Example Applications

In this section we give two example applications for applying our proposed technology. The first example concerns credit card authorization in conjunction with a personal device. The second concerns authenticating government-issued identification credentials. In this section, we use personal device and mobile device synonymously.

4.1 Credit Card Authorization Example

We now give a protocol for integrating our proposed technology within a credit card processing framework. The entities in the protocol are the customer (C), point of sale terminal ($POST$), credit card authorization network ($CCAN$), and location verification service (LVS). The LVS comprises one or more networked entities that track the uniqueness and location of mobile device subscribers on behalf of customers and one or more verifiers. Consequently, we assume personal devices can be tracked using a heartbeat locator (as previously described in §3.1). For reference, we first list the exchanged messages.

Message 1 $C \rightarrow POST : \quad CC$

Message 2 $POST \rightarrow CCAN : ID_{POST}, CC, Transaction$

Message 3 $CCAN \rightarrow POST : Conditional\_Authorization, LVS, Token$

$$\text{Message 4} \quad POST \rightarrow LVS: \; Token \;\; (= \{Time, MRI, Location_{POST}\}_{LVS,CCAN})$$
$$\text{Message 5} \quad LVS \rightarrow POST: \; Status, \{Status, Time, MRI, Location_{POST}\}_{LVS,CCAN}$$

Messages 2-3 and 4-5 represent a query response type of exchange between the parties. The notation $\{X\}_{LVS,CCAN}$ represents confidentiality, integrity, and authenticity of the data X with respect to a mutually shared key between LVS and $CCAN$. In addition to this protection these exchanges employ a secure connection between the pair of entities. This protection includes mutual authentication, confidentiality, and message stream integrity.

Message 1 represents a physical or visual communication whereby the customer presents his credit card to the merchant $POST$ whereupon the merchant acquires credit card information (CC) including the credit card number, expiration date, and security code.

In Message 2, $POST$ initiates an authorization request through the credit card authorization network ($CCAN$). The authorization request includes the merchant identifier (ID_{POST}), CC, and other transactional data ($Transaction$). Upon receipt of this message $CCAN$ uses a table indexed by credit card numbers to determine the particular LVS and mobile reference identifier (MRI – see end of §4.1) for the particular credit card holder. $CCAN$ looks up $Location_{POST}$ (i.e. $POST$'s location per $CCAN$'s records) using another table indexed by $POST$. $CCAN$ uses this information to prepare Message 3 in response to the authorization request. The message contains a conditional authorization for the transaction subject to an obligation on the $POST$'s part (if enabled) to perform a location cross-check query to the LVS identified in the message. Also, within the message is an authorization token (i.e. $\{Time, MRI, Location_{POST}\}_{LVS,CCAN}$) for $POST$ to use when obtaining location verification services from LVS. The authorization token contains the current time ($Time$), the MRI of the device to be queried, and the location of the $POST$ per $CCAN$'s records ($Location_{POST}$).

In Message 4, $POST$ submits a location cross-check query by forwarding the authorization token to LVS. LVS is about to authenticate the authorization token as having come from $CCAN$. Upon receipt of Message 4, LVS checks that $Time$ is recent. Also, LVS uses standard techniques to determine that the request has not been previously responded to. Also the LVS checks $Location_{POST}$ matches the current (or most recent known) location of the mobile device associated with MRI, per LVS knowledge – based on the LVS looking up the unique mobile device using the MRI as an index to the table. (Different MRIs from may refer to the same mobile device.) Finally, the LVS checks that the mobile device is reasonably persistent with no indication of irregularities (e.g. cloning) – as discussed in §3.2, this is a critical aspect of the LVS service.

Message 5 responds to the location cross-check query with a status (*Status*) e.g. verified or unverified, and contains an encrypted component or receipt which can be used to prove the status check to $CCAN$ if the conditional authorization is in dispute. An improvement would be to digitally sign the receipt so that it is not easily repudiated.

Our protocol has a number of attractive properties. First, as a privacy feature, a malicious $POST$ acting alone can not easily determine the location of a mobile device even with a stolen credit card number. A location query is limited to verifying the specific location associated with the $POST$ (i.e. $Location_{POST}$ inserted in the authorization token by $CCAN$). Thus, the $POST$ is unable to test the correctness of a guessed mobile device location, and responses to location queries do not reveal the mobile device's location.

Another attractive property is that LVS is distinct from $CCAN$ and LVS can audit a $CCAN$ which generates abnormal query patterns. (Perhaps $CCAN$ has a dishonest employee taking money to track the whereabouts of people.) Abnormal patterns might include numerous queries for a customer with different locations within a short period of time, or numerous queries to the same location for the same customer at different times. An alternative design might have LVS tightly integrated with the $CCAN$ whereby the $CCAN$ would initiate the location cross-check itself. This would have the advantage of requiring no changes to the $POST$ and reducing the overall delay for transaction processing since location verification could happen in parallel with credit card authorization. The downside is that location privacy assurances to the customer may be diminished.

Mobile Reference Identifier (MRI). The MRI embodies a means for referencing a personal device associated with the customer. A simple approach sets the MRI to the electronic serial number (ESN) (if applicable) of the personal device. However, this is problematic since the MRI could then be used to profile customers across multiple verification points or even across other applications and systems. If implemented carefully, the MRI is not useful to profile customers across multiple verification points or applications. MRI can be assigned uniquely per organization and customer pair. Alternatively, using standard cryptographic techniques, MRIs might be made unique per transaction. A different privacy issue concerns protecting the confidentiality of client locations with respect to even the LVS. The nature of the relationship between the mobile telecommunication company and the mobile subscriber seems to infer that the telecommunication company must be trustworthy (and possibly regulated [15]). For further discussion of privacy, see §5.

4.2 Government Identification Example

We now give an example of authenticating government-issued identification using a location verification service. The general approach is for a person or "customer" to register his mobile device in a manner whereby it may be associated with a government crediential.[19] Laws could require a person's authorization before conveying *any* location information concerning a person's mobile device. A larger issue is the related work [12] on enabling users to control access to their location information as location-based services become popular.

The protocol participants include the person or customer (C) to be authenticated, the point of verification (POV), and the location verification service (LVS). Without loss of generality, we assume that POV also includes or has access to a verification network consisting of one or more (possibly remote) databases. These databases hold government identification information and a "black list" of identities. For the purposes of this example, we assume the government identification is a state driver's license, and the point of verification is a U.S. immigration checkpoint on the U.S. - Canada border. We'll assume that mobile devices are pre-registered with an LVS service and the POV is able to learn an authentic mobile reference identifier (MRI) associated with the user's mobile device. For example, a by-product of registration is initializing the POV with the ability to determine an MRI (perhaps from transaction information) to be used in requesting location cross-checking services from the LVS. Perhaps, a customer benefits by having an expedited border crossing if registered. This pre-registration results in the POV having access to a database of $(Driver\_ID, LVS, MRI)$ tuples.[20] The message exchange sequence follows:

Message 1 $\quad C \rightarrow POV: \quad Driver\_ID$

Message 2 $\quad POV \rightarrow LVS: \quad Token1\,(= [Time, MRI, Location_{POV}]_{POV})$

---

[19] Mandatory (or de facto) mobile device registration for the purposes of tracking could significantly degrade the privacy of individuals. Our goal is not to advocate such tracking but to better understand techniques for addressing privacy issues for credential verification involving location cross-checking.

[20] Registration details will be described in the full paper.

Message 3 $LVS \rightarrow POV$: $[Status, Token1]_{LVS}$

The general approach of this example is similar to the credit card authorization example. In Messages 2 and 3 the square brackets represent a digital signature using the private key of the subscripted entity (e.g. POV in Message 2). This notation implies that the data and the digital signature are conveyed in the message. As with the prior example, we assume that the exchange of Messages 2 and 3 is secured by a connection whose protection include mutual authentication, confidentiality, and message stream integrity.

As illustrated in Message 1, the customer gives the physical driver's license to the POV.

In Message 2, the POV sends a digitally signed message containing the driver's license to an LVS associated with handling LVS inquiries associated with that driver's license. The signed message contains the current time ($Time$), the MRI corresponding to $Driver\_ID$ obtained from a table lookup of $(Driver\_ID, LVS, MRI)$ tuples, and the POV's assertion of its own location ($Location_{POV}$). The signature on the message gives evidence that POV made the request to verify the location information. Later this message can be reconciled with audit trails at the LVS and user authorizations if a privacy dispute evolves.

Upon receipt of Message 2, LVS checks that $Time$ is recent and standard techniques are used to determine that the message has not previously been received. Optionally, LVS refers to a database table of $(POV, POV\_Location)$ to determine that $POV\_Location$ equals the received location, $Location_{POV}$. This counters the threat of a rogue POV (or employee of the POV) from submitting unauthorized location queries (e.g. an employee making queries to determine the location of another individual).

Next, LVS checks whether these locations match the current physical location (per LVS knowledge) of the mobile device indicated by MRI. Finally, LVS checks that the mobile device is not associated with irregularities (see §3.2).

In Message 3, LVS sends a signed response including $Status$ (verified or unverified) and a reference to the original request.

The above illustrates how government identification verification can be made more resilient using our techniques, and how the location-privacy of a customer can be protected. The functionality of the verifier and the location based service is split; the capability of the verifier to choose a device and location to be checked is limited.

5 Privacy Issues

The proposal of §3.1 is a starting point towards a technical system-level approach to addressing identity theft. We acknowledge that it leaves many opportunities for enhancement, and contains some features which some may find unacceptable. Among these is the loss of privacy as a result of continual location-tracking. While there is always a price to pay for increased security, for some users this loss of privacy will clearly be above the acceptable threshold. Thus it is important to explore means to address this privacy issue (cf. [9, 23]).

A user can choose a *trusted third party* (TTP) he is willing to trust to maintain the privacy of his information. In many ways the user is already trusting the communication provider of his personal device (e.g. cell phone, and wireless internet) concerning the privacy of his location information.[21] More generally, while each user could be associated with a particular TTP for location tracking, a relatively large set of TTPs in the overall system could aid scalability and eliminate system-wide single points of failure.

The "Wireless Privacy Protection Act of 2003" [15] requires customer consent related to the collection and use of wireless call location information, and call transaction information. Further it requires that "the carrier has established and maintains reasonable procedures to protect the confidentiality, security, and integrity of the information the carrier collects and maintains in accordance with such customer consents." This or other legislation could mean that straight-forward approaches are practical if organizations can be trusted to adequately protect location data. However, it may be argued that many information-receiving organizations might not be able or trustworthy to guarantee protection of location information and personal transaction data.

As the idea of relying on regulation and the trustworthiness of information holders to protect location and other personal information may cause discomfort to those with strong privacy concerns, we encourage further research on using privacy-preserving techniques to achieve digital uniqueness with a trusted (or minimally trusted) third party. To this end, there exists extensive literature following on from Chaum's early work [4] on digital pseudonyms and mix networks, for protecting privacy including the identities involved in, and the source/destination of communica-

[21] As a side comment, many people enjoy far less privacy than perhaps presumed, due to existing location-tracking technology such as wireless 911 services (see §2). However, this may not bring much comfort.

tions. Privacy-related applications of such techniques include e-elections (e.g. [21]), anonymous email delivery (e.g. [5]), and of particular relevance, location management in mobile communications [7]. (For further recent references, see e.g. [11].) While we do not foresee serious technical roadblocks to integrating largely existing privacy-enhancing technologies to significantly improve the privacy aspects of this proposal, further pursuit of this important topic is beyond the scope of this paper.

Regardless of the means to protect the privacy of data held at a carrier or network service provider it is also important to protect the service interface to restrict query access since this service interface can be used to reveal private location-based information. An important issue becomes specifying individual privacy policies and ensuring that the system enforces these policies.

Individual privacy policies are highly dependent on the application and can be based on a number of factors. Factors include *who* is making the request, the *location* in the request, the *time* of the request, the *frequency* or number of requests, and other *circumstances* for the request. Logically, the individual potentially delegates authorization rights to location-based services based on his privacy policy. In our credit card example (see Section 4.1), implicitly the individual grants location-based verifications to CCAN who further grants a restricted one-time transactional authorization to POST. These transactional authorizations were restricted to location-based queries for the (pre-determined) location of the POST. Furthermore, these authorizations are restricted to queries made at the same time of the credit card transaction.

An alternative direction for controlling access to location-based authentication services is for the individual with the personal device to further retain per-transaction control for granting location-based verifications. This may take place in any number of ways including having the personal device digitally sign the access request after first checking the request conforms to its privacy policy. In the credit card example, the personal device of the credit card holder might issue an authorization based on the current time of the personal device, and the particular identity of the POST. In the government identification example, the personal device might issue an authorization contrained on the POST being a government rather than commerical organization, the current time of the personal device, and the particular location of the personal device.

6 Concluding Remarks

We have proposed an architecture and system for authentication involving a physical location cross-check, and reliance on an entity uniqueness property and funneling within the verification process. While the system is relatively simple – essentially a selective combination of existing technology and techniques – we believe it may be successful at stopping many forms of identity theft. This appears to be among the first technical proposals to address identity theft in a research paper. In our view, part of the problem is that it is not clear which research community is a natural "owner" of the problem. Although in many ways more of a system-engineering than a traditional security problem, we believe that increasingly, technical solutions to identity theft will fall to the security research community. Indeed, phishing for passwords and installation of key-logging software/hardware, which both facilitate identity theft, are problems whose solutions one would naturally seek from the security research community.

It should be clear that we have not yet built the proposed system, even in a test environment, and doing so would not "prove" our proposal was secure in a practical sense. The best, and perhaps only true way to test such a system would be to observe any reduction in identity thefts in a real-world deployment. Nonetheless, we believe this paper lays out sufficient details for security-aware systems-level engineers within appropriate organizations (e.g. major credit card associations, banks, credit rating agencies, or national ID card system designers – cf. [20]) to implement such a system. Any such implementation must be designed keeping scalability in mind, particularly in light of the continuous nature of the tracking.

Effectively, our proposal is a mechanism for enforcing unique ownership of names (i.e. identities), and includes an extension addressing the minting of new (fraudulent) credentials from stolen credentials. We encourage the research community to explore alternate solutions to the latter problem, which is closely linked to that of identity theft.

Acknowledgements. We thank Carl Ellison for bringing references [8] and [10] to our attention. The first author acknowledges the support of the National Sciences and Engineering Research Council of Canada (NSERC) for support under both a Canada Research Chair and an NSERC Discovery Grant. The second author acknowledges that this material is based on work supported by the National Science Foundation under Grant No. 0208983.

References

1. CNN.com, "Anti-identity theft freeze gaining momentum; Credit companies resist measure", Aug.3 2004, http://www.cnn.com/2004/TECH/biztech/08/03/security.freeze.ap.
2. N. Chou, R. Ledesma, Y. Teraguchi, J.C. Mitchell, "Client-side defense against web-based identity-theft", Proc. of *Network and Distributed System Security Symposium* (NDSS'04), Feb. 2004, San Diego.
3. Mark D. Corner, Brian D. Noble, "Zero-Interaction Authentication", Proc. of *MOBICOM'02*, 23–28 Sept. 2002, Atlanta.
4. D. Chaum, "Untraceable electronic mail, return addresses, and digital pseudonyms," *Comm. of the ACM*, 1981, pp.84–88.
5. G. Danezis, R. Dingledine, N. Mathewson, "Mixminion: Design of a Type III Anonymous Remailer Protocol", pp.2–15, *2003 IEEE Symp. Security and Privacy.*
6. D.E. Denning, P.F. MacDoran, "Location-Based Authentication: Grounding Cyberspace for Better Security", *Computer Fraud and Security*, Feb. 1996.
7. H. Federrath, A. Pfitzmann, A. Jerichow, "MIXes in Mobile Communication Systems: Location Management with Privacy", *Workshop on Information Hiding*, Cambridge U.K., 1996.
8. A. Fox, E. Brewer, "Harvest, Yield and Scalable Tolerant Systems", Proc. of *HotOS-VII*, 1999.
9. E. Gabber, A. Wool, "On Location-Restricted Services", *IEEE Network*, November/December 1999.
10. Seth Gilbert, Nancy Lynch, "Brewer's conjecture and the feasability of consistent, available, partition-tolerant web services", *Sigact News* 33(2), June 2002.
11. P. Golle, A. Juels, "Parallel Mixing", pp.220-226, *2004 ACM Conf. Computer and Comm. Security.*
12. Carl A. Gunter, Michael J. May, Stuart G. Stubblebine, "A formal privacy system and its application to location based services", in: *Workshop on Privacy Enhancing Technologies 2004*.
13. Darcy Henton, "Identity-theft case costs taxpayers in Alberta $540,400", 12 April 2004, *The Globe and Mail*, Toronto.
14. J. Hightower, G. Borriello, "Location Systems for Ubiquitous Computing", *IEEE Computer*, Aug. 2001.
15. Wireless Privacy Protection Act of 2003, 108th Cong, H.R. 71 (United States).
16. Mark Hume, "Security breach lets criminals view Canadians' credit reports", 16 March 2004 (page A1/A7), *The Globe and Mail*, Toronto.
17. Mark Hume, "Identity theft cited as threat after Equifax security breach", 17 March 2004 (page A7), *The Globe and Mail*, Toronto.
18. M. Just, P.C. van Oorschot, "Addressing the problem of undetected signature key compromise", Proc. of *Network and Distributed System Security Symp.* (NDSS'99), Feb. 1999, San Diego.
19. G. Keizer, "Internet Scams Cost Consumers $2.4 Billion", TechWeb News, InternetWeek, 16 June 2004.
20. S.T. Kent, L. Millette, eds., *IDs – Not That Easy: Questions About Nationwide Identity Systems*, National Academies Press (U.S.), 2002.

21. A. Kiayias, M. Yung, "The Vector-Ballot e-Voting Approach", pp.72–89, *Financial Cryptography'04*.
22. Robert Lemos, "Pop-up program reads keystrokes, steals passwords", CNET News.com, 29 June 2004, http://news.com.com/2100-7349-5251981.html.
23. P. Lincoln, P. Porras, V. Shmatikov, "Privacy-Preserving Sharing and Correlation of Security Alerts", in: Proc. of *13th USENIX Security Symposium*, August 2004, San Diego.
24. Karen Lu, Asad Ali, "Prevent Online Identity Theft – Using Network Smart Cards for Secure Online Transactions", *2004 Information Security Conference* (ISC'04), Sept. 2004, Palo Alto.
25. A. Malpani, R. Housely, T. Freeman, Simple Certificate Validation Protocol (SCVP), Internet Draft (work in progress), draft-ietf-pkix-scvp-15.txt, July 2004.
26. Declan McCullagh, "Season over for 'phishing'?", CNET News.com, 15 July 2004, http://news.com.com/2100-1028-5270077.html.
27. M. Myers, R. Ankney, A. Malpani, S. Galperin, C. Adams, X.509 Internet Public Key Infrastructure: Online Certificate Status Protocol – OCSP, Internet Request for Comments 2560, June 1999.
28. Panda Software, Bankhook.A (Virus Encyclopedia entry), http://www.pandasoftware.com.
29. Public Law No. 108-275, "Identity Theft Penalty Enhancement Act", United States, July 2004.
30. Jonathan D. Salant, "Call centers lag in cell-phone tracking upgrade, group says", 6 February 2004, (page A8), *The San Diego Union Tribune*.
31. N. Sastry, U. Shankar, D. Wagner, "Security verification of location claims", in: *2003 ACM Workshop on Wireless Security (WiSe 2003)*.
32. James Sherwood, "So you want to be a cybercrook...", CNET News.com (ZDNET UK), Aug.29 2004, http://zdnet.com.com/2100-1105-5317087.html.
33. John A. Volpe National Transportation Systems Centre, Vulnerability Assessment of the Transportation Infrastructure Relying on the Global Positioning System, Final Report, 29 August 2001.