On Identity Theft and a Countermeasure based on Digital Uniqueness and Location Cross-Checking*

P.C. van Oorschot[1] Stuart G. Stubblebine[2]

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

Abstract. We define identity theft as the unauthorized use and exploitation of another individual's identity-corroborating information. Published research proposing technical countermeasures is sparse, in contrast to a number of recent proposals to address the sub-problem of *phishing*. We first identify some underlying problems facilitating identity theft. To address 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 provide two example protocol implementations of this system, for two representative applications. We believe that this system merits consideration for practical use, and will stimulate further technical solutions to the broad problem of identity theft.

1 Introduction and Motivation

We define *identity theft* as 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 widely acknowledged as one of today's fastest growing crimes. A 2003 FTC survey [13] estimated "that nearly 10 million [U.S.] consumers discovered that they were victims of some form of identity theft in the preceding 12 months, costing American businesses an estimated $48 billion in losses, and costing consumers an additional $5 billion in out-of-pocket losses." 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 §3) and other identity fraud [37]. While the problem is not new, identity theft is exacerbated by the availability of personal information on the Internet, and by criminal use of the Internet to extract personal information from individuals.

Despite growing media attention and numerous web sites discussing the problem, its seriousness 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

* Version: 13 December 2005. A preliminary version of parts of this paper appeared as [35], copyright of which is held by IFCA and used here with permission.

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, few effective technical countermeasures have been proposed, and to our knowledge none have been successfully adopted to the point of decreasing identity theft in practice.

Our Contributions. We identify underlying problems facilitating identity theft, and propose a general authentication architecture and system we believe could 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, and give two example protocol instantiations for representative applications: credit card authorization in conjunction with a personal authentication device (e.g., cell phone), and authenticating government-issued identification credentials. We propose an extension addressing the problem of acquiring fraudulent new identity credentials from stolen credentials. Overall, a major objective is to stimulate further research on technical solutions to the "whole" problem of identity theft, as we have defined it (see discussion in §2), rather than subsets thereof – e.g., phishing or key-logging alone, or credit card fraud.

Organization. The sequel is organized as follows. §2 discusses background and fundamental problems underlying identity theft, before related work is briefly discussed in §3. §4 presents an overview and basic proposal for our authentication system and architecture addressing identity theft, and a high-level security analysis considering a number of potential attacks. §5 provides two example applications as mentioned above, including application-specific protocol messages. §6 discusses privacy considerations and refinements to address the obvious concern of privacy loss due to the location-tracking component of our proposal; however proper attention to such important privacy issues is beyond the scope of the present paper. §7 gives concluding remarks.

2 Identity Theft: Background and Underlying Problems

After first discussing credentials, we identify what we see as the fundamental issues facilitating identity theft.

We define *identity credentials* (*credentials*) rather loosely as "evidence" 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 such as the following: if a secret key from a hardware token is extracted, the key itself (i.e., the important component of the token) is then externally available in digital form. So both the physical token, and the critical

data within it, are credentials (or equivalently, *credential information*). A further looseness arises from the fact that some items of information, such as (U.S.) Social Security Number, are often used as identity-corroborating data even though they are not generally treated as secret, and are commonly provided to relying parties only verbally (without display or inspection of a physical form).

To repeat our definition, identity theft is the unauthorized use and exploitation of another individual's identity-corroborating information. This includes both the misuse of existing accounts, and creation of new accounts in a victim's name; the former alone is a simpler sub-problem, which includes password theft to provide access to existing accounts.

There are numerous reasons why personal identities and credential information are so easily stolen, and why this presents a difficult challenge. We believe the fundamental problems underlying and 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. [27]);[1]
F3: *independence of new credentials:* if existing credential information is used by an impersonator to obtain new credentials, the latter are in a sense "owned" by the impersonator, and usually no information flows back to the original credential owner immediately; and
F4: *exterior-system use of credentials:* credentials are often used by relying organizations other than the issuing organizations which created them, and for purposes not intended by the issuing organization. Here, credential abuse in the relying (but non-issuing) system is typically beyond any feedback control of the issuing system.

Regarding F2, note that 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. In particular due to F3, we see identity theft as a problem with the following characteristic, which we state as a proposition.

P1: Identity theft is a *systemic* problem (in the whole-system sense), which cannot be solved by any single credential-issuing organization in isolation.

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.

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.

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

Among other things, the credit bureaus can when necessary post alerts (see §3) on credit files of individuals whom they suspect are subjects of identity theft. However, it is unclear how strongly 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 [3]). Moreover, at least one such organization[2] was itself exploited by criminals in an incident raising fears of large-scale identity theft.

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. Although numbers vary widely depending on the study, and the definition of identity theft may vary widely, a San Diego study [15] covering 2004 reported a mean of 330 hours (ranging from 3 to 5,840 hours) spent by identity theft victims "in the recovery process"; this may include, e.g., getting government and commercial organizations to stop recognizing stolen identification information, and time invested to get new identity information re-issued. Other 2003 studies report averages ranging from 30 to 60 hours of time cost to victims [3, 13].

3 Related Work

Related work is discussed below, as well as in §6.

Physical location determination. The U.S. Federal Communications Commission (FCC) requires [12, 18] that by December 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 this date, or deploy other location-tracking technology (e.g., triangulation or location determination based on distance and direction from base stations); thereafter emergency call centers must deploy related technology to physically locate callers. As of February 2004, 18% of U.S. call centers had this technology [40].

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 [23]. Sastry et al. [41] 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 [18] discuss four schemes, all based on available infrastructure, for detecting the movement of

[2] Equifax Canada recently confirmed that in February 2004, 1400 consumer credit reports were "accessed by criminals posing as legitimate credit grantors" [25, 26].

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 [9] outlined a commercial location-based authentication system using the Global Positioning System (GPS), notwithstanding standard GPS signals being subject to spoofing [18, 43]. 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.

Activity profiling. *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 the broader problem of identity theft. While consumers have limited liability on use of fraudulent credit cards in their name, fraud protection by credit card companies is limited to the realm of credit cards (as one might expect, in a capitalistic world; cf. the business motivation of credit bureaus). Regarding protection afforded by banks, in the U.S., when one major bank puts an alert on a name, a common clearinghouse (limited to banks) allows all major banks to share that warning [26].

Credit-check freezes and fraud alerts. One real-world system-level technique to ameliorate identity-theft is the use of *credit-check freezes* [3], now available in many U.S. states; credit reports become unavailable for viewing by large classes of inquirers. An individual can alternatively place a *fraud alert* on their credit reports, 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), or unless the individual is specifically contacted and agrees to grant access. This may involve a request that no new credit be granted without approval of the individual, although current laws do not generally require credit grantors respect this desire. Another option is selective access, whereby a flagged credit report can be accessed only by specifically named inquirers. These methods are intended to 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).

Phishing and keylogging. *Phishing* is a relatively new Internet-based attack used to carry out identity theft, or more precisely, most often used to obtain sensitive personal information related to existing accounts. "Phishing kits" available on the Internet allow even amateurs to create bogus web sites and use spamming software to defraud users [42]. 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 and thereby fall victim. This is a variation of an old attack whereby malicious software planted on a user machine puts up a fraudulent login interface to obtain the user's userid and login password to an account.

Chou et al. [4] propose a client-side software plug-in and various heuristics for detecting online phishing scams, opening the research area of technical approaches to detect phishing. Ross et al. [39] propose a browser-based technique involving site-specfic password hashing to provide protection against password phishing, although remaining challenges include usability and keyloggers (see below). In contrast to a good number of other recently published anti-phishing proposals (see Dhamija and Tygar [10, §7], including their own "dynamic security skin" proposal), there is a lack of corresponding proposals to the (broader, more difficult) general problem of identity theft, as we have defined it.

Key logging attacks rival phishing as a serious threat to sensitive personal information, particularly in today's environment of ubiquitous malware and software loads on many Internet devices changing daily. For example, the program *Bankhook.A* [36, 31], which spread without human interaction beyond web browsing, involved a (non-graphic) file named *img1big.gif*, and exploited a vulnerability in a widely used web browser. Upon detecting attempted connections to any of about 50 major online banks,[3] it recorded sensitive information (e.g., account userid and password) prior to SSL encryption, and mailed that data to a remote computer.

Relationship to PKI systems. There are similarities between detecting the theft and usage of password-based credentials and that of signature private keys as used in public-key infrastructure (PKI) systems; 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 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 (e.g., see [1]). Among many other proposals (e.g., see Aiello et al. [2]), this has lead to several *online status checking* proposals (e.g., OCSP [34] and SCVP [33]), 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* [16, 19]: 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.

Corner and Noble [5] propose a general 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 in-

[3] Text string searches were made for https connection attempts to URLs containing 50 target substrings.

novative alternative to easily replicated digital authentication credentials – simultaneously increasing security and decreasing user interaction (e.g., vs. standard password login).

4 Authentication based on Uniqueness, Location and Funneling

A high-level overview of our proposed authentication system is given in §4.1. Security is considered in §4.2. More details, e.g., specific protocol messages for example applications, are given in §5. Privacy considerations and refinements are discussed in §6.

4.1 High-level Overview of Proposed System

Our goal is a system which prevents, or significantly reduces, occurrences of identity theft in practice. Our basic design is as follows. Every system user has a hardware-based *personal device*,[4] 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[5] 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 §3). Our heartbeat locator is unrelated to, as far as we know, other uses of this term in the literature (e.g., the *heartbeat messages* of Danezis et al. [8]).

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 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 §4.2.

In the process of a transaction being executed/processed, when an identity[6] 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,[7] with the features of (1) combining location determination with continuous location tracking; and

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

[5] 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 $L_B \neq L_A$.

[6] An identity per se is not required – e.g., pseudonyms could be used, to enhance privacy (see §6).

[7] 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.

(2) funneling 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 §3). 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-center 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 credential would agree[8] 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.

4.2 Security Analysis and Discussion

In this section we provide a high-level security analysis of the new proposal, and discuss necessary checks regarding the personal device. While we offer no rigorous security arguments in the present paper, 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 even for more mathematically-oriented proposals, "proofs" of security are at best relative to a particular model and assumptions, which often differ from the reality of deployed systems. At best, increased confidence in the relevance and suitability of these are gained only over time; but this is not always the case (e.g., see Koblitz and Menezes [30]). Nonetheless, we strongly encourage further analysis to allow the proposal to be iteratively improved.

We begin by referring back to the four fundamental problems of §2. The system proposed in §4.1 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) may

[8] We recognize that this would require a significant change in behavior 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.

be partially[9] addressed by the "minting bit" extension. Exterior-system use of credentials (F4) may be best dealt with by legal means and adherence to best practices (e.g., *not* using credentials issued by an organization with which the relying organization does not have an appropriate relationship).

Device irregularities, theft and cloning. Fraud mitigation strategies depend on users reporting stolen personal devices in a timely matter.[10] 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 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 constraints 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

[9] 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.

[10] 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.

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 biometrics) 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 reactivation).[11]
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. However a geographic irregularity would arise (as the clone's location would differ from that of the last heartbeat of the real device).

[11] Although a motivated and well-armed attacker can generally defeat user-to-device authentication mechanisms (cf. [18]), 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.

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 a text message, e.g., 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.

5 Protocol Examples for Representative 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. Herein, we use the terms personal device and mobile device synonymously, and this device is used, as discussed earlier, for the heartbeat locator functionality.

5.1 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. The protocol entities and message flow are illustrated in Figure 1.

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

Message 2 $POST \rightarrow CCAN : POST, CC, Transaction$

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

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

Message 5 $LVS \rightarrow POST: \ Status, \{Status, Time, MRI, Location_{POST}\}_{LVS,CCAN}$

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

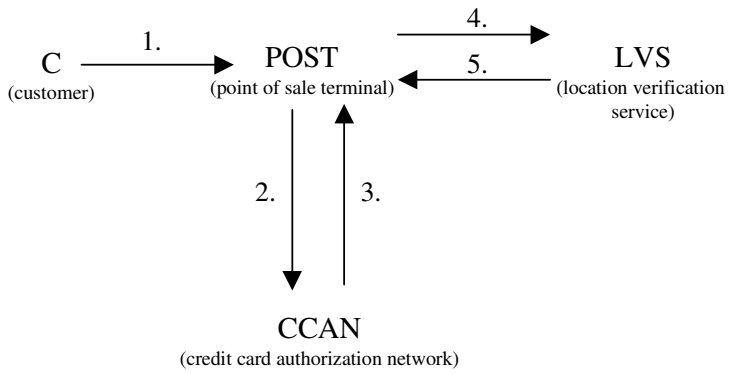

Fig. 1. Message flows (credit card authorization example).

this protection these exchanges employ a secure connection between the pair of entities, providing message stream integrity (e.g., to prevent interleaving attacks) as well as mutual authentication and confidentiality.

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 ($POST$), credit card (CC), and other transactional data ($Transaction$). Upon receipt of this message $CCAN$ uses a table indexed by credit card numbers to lookup the particular LVS and mobile reference identifier (MRI – see end of §5.1) 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. Assuming credit is authorized, the message contains a conditional authorization for the transaction subject to an obligation on $POST$'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 requesting location verification services from LVS. The authorization token contains the current time that $CCAN$ created the authorization ($Time$), an MRI associated with the device to be queried, and the location of $POST$ per $CCAN$'s records ($Location_{POST}$).

In Message 4, $POST$ submits a location cross-check query by forwarding the authorization token to LVS. LVS authenticates the token as having been created by $CCAN$. Upon receipt of Message 4, LVS checks that $Time$ is recent. Also, LVS uses standard techniques

(e.g., caching recent messages) to determine that the request has not been previously responded to. Also the LVS checks that $Location_{POST}$ 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 mobile device using MRI as a record index. Finally, the LVS checks that the mobile device is reasonably persistent with no indication of irregularities (e.g., cloning) – as discussed in §4.2, 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.)

The above 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 wishing to take 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 is a variable used to reference a personal device. In the above credit card authorization example the MRI is used to convey which personal device is to be location cross-checked. Later the LVS queries a database indexed by MRI to access data records associated with the device (e.g., its last known location). Herein the MRI could simply be a non-identifying (unique) number associated with a particular credit card number. MRI is previously generated by LVS and given to $CCAN$ (along with the associated credit card number), e.g., in a registration phase, after C authorizes $CCAN$ to allow $POST$ to query LVS concerning its mobile device (with respect to his credit card transactions). Ideally, LVS authenticates C's authorization in a reliable manner, e.g., by contacting C to confirm an authorization request triggered by $CCAN$. Our protocol is independent of the type of network used to track the personal device.

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 telecommunica-

tion company and the mobile subscriber seems to infer that the telecommunication company must be trustworthy (and possibly regulated [24]). For further discussion of privacy, see §6.

5.2 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 credential.[12] 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 [22] 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.

A goal of the protocol is that the use of LVS services doesn't give POVs an improved ability to profile C over what may have been previously possible. We now discuss a registration protocol that supports this goal.

Registration protocol. For the purposes of this registration protocol, assume that the personal device is a cell phone with text messaging capabilities (e.g., short-messaging service or SMS). Further assume that LVS is a large cellular provider (or affiliated cellular providers). The registration protocol takes place upon issuance of the $Driver\_ID_C$. We assume the customer, C, is in possession of the cell phone. The registration steps are as follows. The user identifies himself and authenticates himself to the drivers license issuing authority. (Typically this step is already done when applying for a drivers license.) The user is asked to input their cell phone number into a trusted registration terminal, e.g., like the terminals used to accept PINs for debit cards. (The user trusts the license authority to have secure terminals.)

The terminal causes an SMS challenge nonce to be issued to the cell phone. The user receives the challenge via the cell phone and enters the nonce into the registration terminal. Upon verifying the correct response, the drivers license issuing authority generates a request to LVS (containing the cell phone number) to enroll the user's device for location cross-checking services. The enrollment authorizes POVs (see below) to query LVS concerning the mobile device.

[12] 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.

Next, LVS generates a unique number, MRI_C, and stores a record entry such as $(MRI_C, Cell\_Number)$ in a table for access by a set of POVs associated with this MRI_C. The table will be used later to retrieve $Cell\_Number$, indexed by MRI_C. Preferably, records of the form $(Driver\_ID_C, MRI_C)$ are stored by LVS or made available in a central database that can be remotely accessed by each distributed POV needing to retrieve MRI_C (using $Driver\_ID_C$ as the index). This prevents the need for replicating this data at each POV.

GOVERNMENT ID PROTOCOL EXAMPLE. The message exchange sequence follows.

Message 1	$C \rightarrow POV:$	$Driver\_ID_C$
Message 2	$POV \rightarrow LVS:$	$Token1\ (= [Time, MRI_C, Location_{POV}]_{POV})$
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 includes 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, POV uses $Driver\_ID_C$ as the index to remotely query the registration table to obtain MRI_C. POV sends a digitally signed message containing the current time ($Time$), the MRI_C, and the POV's assertion of its own location ($Location_{POV}$). The signature on the message gives evidence that POV requested location cross-checking services. If later a privacy dispute evolves, this message can be reconciled with LVS's audit trails and user authorizations.

Upon receipt of Message 2, LVS checks that $Time$ is recent. Again, a cache of recent messages is checked to determine that the message is not a replay. LVS refers to a database table of $(POV, POV\_Location)$ to determine that $POV\_Location$ equals the received location, $Location_{POV}$. This check counters the threat of a rogue POV (or dishonest employee of the POV) from submitting unauthorized location queries (e.g., an employee making queries to determine the location of another individual).

Also, LVS checks whether these locations match the current physical location (per LVS knowledge) of the mobile device. To do this LVS uses MRI_C as an index to retrieve the record $(MRI_C, Cell\_Number)$ and to obtain $Cell\_Number$. LVS uses $Cell\_Number$ to determine the location of the mobile device. The means for determining the location of the mobile device depends on the particular communication network and is beyond the scope of this paper. Finally, LVS checks that the status of the device is not *irregular* (see §4.2).

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.

6 Privacy Considerations and Refinements

The basic proposal of §4.1 is a starting point towards a technical system-level approach to addressing identity theft, but has some characteristics which we expect privacy advocates will find unacceptable, most obviously the potential loss of privacy as a result of continual location-tracking. While there is always a price to pay for increased security, this potential loss of privacy seems to be beyond the acceptable threshold (and evidently, is avoidable). Thus it is important to explore means to address this privacy issue (cf. [9, 18, 32]). We now pursue this briefly.

A user might choose a *trusted third party* (TTP) he is willing to trust to maintain the privacy of his information, including e.g., any un-mapping of digital pseudonyms (see below). 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.[13] More generally, while each user could be associated with one 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" [24] 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 (minimally) trusted third party. To this end, there exists extensive literature beginning with Chaum's seminal work [6] on digital pseudonyms and mix networks, for protecting privacy including the identities involved in, and the source/destination of communications. Privacy-related applications of such techniques include e-elections [29], anonymous email delivery [7], anonymous web browsing [11], cen-

[13] 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 §3). However, this may not bring much comfort.

sorship resistant document storage/retrieval [38, §4], and of particular relevance, location management in mobile communications [14]. For an overview of privacy-preserving technologies to protect identity information on the Internet, see [20, 21, 17]. While we foresee no serious technical roadblocks to integrating or adapting such largely existing privacy-enhancing technologies to significantly enhance the privacy aspects of our 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 could otherwise be exploited to extract private location-based information. An important issue becomes specifying individual privacy policies and ensuring that the system respects these policies.

Individual privacy policies are generally highly dependent on the application and can be based on a number of factors. Factors may 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 5.1), 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 release of location information to the verification service. This may take place in any number of ways including having the personal device digitally sign the access request after confirming that 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 constrained on the POST being a government rather than commercial organization, the current time of the personal device, and the particular location of the personal device.

7 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 would be successful at stopping many forms of identity theft. This appears to be among the first technical proposals to address identity theft in the open literature. Although in many ways more of a system-

engineering than a traditional security problem, we believe that increasingly, technical solutions to Internet-based 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 might naturally seek from the security research community. Note however that our proposal focuses not on preventing theft of credential information, but on detecting fraudulent use of such information at the time of a transaction – much in the spirit of our understanding of how credit card companies do activity profiling today.

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. Independent of this, a more complete and rigorous security analysis of specific instantiations of our proposal is needed. The best, and perhaps only true way, to test such an instantiation 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, governments, or national ID card system designers – cf. [28]) to implement such a system. Any such implementation must be designed keeping scalability in mind, particularly in light of the continuous nature of the location-tracking component of our solution.

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 [16] and [19] to our attention. The first author acknowledges the support of the Natural 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. C. Adams, S. Lloyd, *Understanding PKI*, 2nd edition, Addison Wesley Professional, 2002.
2. W. Aiello, S. Lodha, R. Ostrovsky, "Fast Digital Identity Revocation", pp.137-152, Proc. of Crypto'98, Springer LNCS vol.1462 (1998).
3. 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.
4. 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), Internet Society, Feb. 2004, San Diego.
5. Mark D. Corner, Brian D. Noble, "Zero-Interaction Authentication", Proc. of *MOBICOM'02*, 23–28 Sept. 2002, Atlanta.
6. D. Chaum, "Untraceable electronic mail, return addresses, and digital pseudonyms," *Comm. of the ACM*, 1981, pp.84–88.

7. G. Danezis, R. Dingledine, N. Mathewson, "Mixminion: Design of a Type III Anonymous Remailer Protocol", pp.2–15, *2003 IEEE Symp. Security and Privacy*.
8. G. Danezis, L. Sassaman, "Heartbeat traffic to counter (n-1) attacks: red-green-black mixes", Proc. of 2003 ACM Workshop on Privacy in the Electronic Society (WPES'03).
9. D.E. Denning, P.F. MacDoran, "Location-Based Authentication: Grounding Cyberspace for Better Security", *Computer Fraud and Security*, Feb. 1996.
10. R. Dhamija, J. Tygar, "The Battle Against Phishing: Dynamic Security Skins", Proc. 2005 ACM Symposium on Usable Privacy and Security (SOUPS), pp.77-88, ACM Press, July 2005.
11. R. Dingledine, N. Mathewson, P. Syverson, "Tor: The Second-Generation Onion Router", Proc. 13th USENIX Security Symposium, August 2004.
12. Fourth Memorandum Opinion and Order in the Matter of Revision of the Commission's Rules to Ensure Compatibility with Enhanced 911 Emergency Calling Systems (FCC 00-326, CC Docket No. 94-102, 9/11/00), Federal Communications Commission, Washington, DC, 20554.
13. Deborah Platt Majoras, FTC Chairman, Prepared Statement of the Federal Trade Commission, before the Committee on Commerce, Science and Transportation, U.S. Senate, on Data Breaches and Identity Theft, June 16 2005.
14. H. Federrath, A. Pfitzmann, A. Jerichow, "MIXes in Mobile Communication Systems: Location Management with Privacy", *Workshop on Information Hiding*, Cambridge U.K., 1996.
15. L. Foley, J. Foley, S. Hoffman, T. McGinley, K. Barney, C. Nelson, H. Pontell, A. Tosouni, "Identity Theft: The Aftermath 2004, With comparision to The Aftermath 2003", Identity Theft Resource Center (ITRC), San Diego, September 2005.
16. A. Fox, E. Brewer, "Harvest, Yield and Scalable Tolerant Systems", Proc. of *HotOS-VII*, 1999.
17. Free Haven Project, Anonymity Bibliography (www.freehaven.net/anonbib/) and Related Works: Anonymous Communications Systems (www.freehaven.net/related-comm.html), accessed 4 Aug. 2005.
18. E. Gabber, A. Wool, "On Location-Restricted Services", *IEEE Network*, November/December 1999.
19. Seth Gilbert, Nancy Lynch, "Brewer's conjecture and the feasability of consistent, available, partition-tolerant web services", *Sigact News* 33(2), June 2002.
20. I. Goldberg, D. Wagner, D. Brewer, "Privacy-enhancing technologies for the Internet", Proc. 42nd IEEE Spring COMPCON, Feb. 1997.
21. I. Goldberg, "Privacy-enhancing technologies for the Internet, II: Five years later", Proc. Privacy Enhancing Technologies Workshop (PET 2002), April 2002.
22. 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*.
23. J. Hightower, G. Borriello, "Location Systems for Ubiquitous Computing", *IEEE Computer*, Aug. 2001.
24. Wireless Privacy Protection Act of 2003, 108th Cong, H.R. 71 (United States).
25. Mark Hume, "Security breach lets criminals view Canadians' credit reports", 16 March 2004 (page A1/A7), *The Globe and Mail*, Toronto.
26. Mark Hume, "Identity theft cited as threat after Equifax security breach", 17 March 2004 (page A7), *The Globe and Mail*, Toronto.
27. M. Just, P.C. van Oorschot, "Addressing the problem of undetected signature key compromise", Proc. of *Network and Distributed System Security Symp.* (NDSS'99), Internet Society, Feb. 1999, San Diego.
28. S.T. Kent, L. Millette, eds., *IDs – Not That Easy: Questions About Nationwide Identity Systems*, National Academies Press (U.S.), 2002.
29. A. Kiayias, M. Yung, "The Vector-Ballot e-Voting Approach", pp.72–89, *Financial Cryptography'04*.
30. N. Koblitz, A.J. Menezes, "Another look at 'provable security' ", version of May 4, 2005, IACR ePrint archive, http://eprint.iacr.org/2004/152.pdf
31. Robert Lemos, "Pop-up program reads keystrokes, steals passwords", CNET News.com, 29 June 2004, http://news.com.com/2100-7349-5251981.html.
32. 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.

33. A. Malpani, R. Housely, T. Freeman, Simple Certificate Validation Protocol (SCVP), Internet Draft (work in progress), draft-ietf-pkix-scvp-15.txt, July 2004.
34. 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.
35. P.C. van Oorschot, S. Stubblebine, "Countering Identity Theft through Digital Uniqueness, Location Cross-Checking, and Funneling", *Financial Cryptography and Data Security'05*, Feb.28–Mar.3, 2005, pp.31-43, A.S. Patrick, M. Yung (eds.), proceedings, LNCS 3570, Springer-Verlag 2005.
36. Panda Software, Bankhook.A (Virus Encyclopedia entry), http://www.pandasoftware.com.
37. Public Law No. 108-275, "Identity Theft Penalty Enhancement Act", United States, July 2004.
38. G. Perng, M.K. Reiter, C. Wang, "Censorship Resistance Revisited", 2005 Information Hiding Workshop (IH 2005).
39. B. Ross, C. Jackson, N. Miyake, D. Boneh, J. Mitchell, "Stronger Password Authentication Using Browser Extensions", USENIX Security 2005.
40. Jonathan D. Salant, "Call centers lag in cell-phone tracking upgrade, group says", 6 February 2004, (page A8), *The San Diego Union Tribune*.
41. N. Sastry, U. Shankar, D. Wagner, "Security verification of location claims", in: *2003 ACM Workshop on Wireless Security (WiSe 2003)*.
42. 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.
43. John A. Volpe National Transportation Systems Centre, Vulnerability Assessment of the Transportation Infrastructure Relying on the Global Positioning System, Final Report, 29 August 2001.