

Non-Repudiation Despite the Digital Age

Roger Lines CISA, Health Insurance Commission

WHAT THIS PAPER IS ABOUT

This paper identifies a problem with conventional non-repudiation measures for digitally signed messages. The problem is here named the “false repudiation problem”, and it affects classes of electronic messages *outside* purely commercial relationships. The examples discussed include messages sent to government agencies, and messages containing professional advice which might later turn out to be negligent. Approaches to solving the problem are also discussed.

I have assumed that the reader is already familiar with digital signatures, certification of public keys, and the conventional approaches to non-repudiation based on Certificate Revocation Lists or public directories, covered loosely by the term Public Key Infrastructure [PKI]. If you need to find out more about these, I recommend the ISACA publication *Digital Signatures Security and Controls*, plus a number of articles appearing in the *IS Audit and Control Journal* over the last few years. A particularly relevant article was Meng-Chow Kang, *Dynamic Handwritten Signature Verification System*, appearing in Volume III, 1998.

WHAT IS NON-REPUDIATION

The digital signature, unaided by certificates, allows a message to be authenticated as to content and origin, if the signer (Alice) and recipient (Bob) have a prior relationship. This prior relationship specifically includes Bob knowing Alice’s public key and being sure it is hers.

A public key certificate allows the digital signature of a stranger to be verified. This is achieved through a chain of intermediate Certification Authorities.

“Non-repudiation” refers to the collected characteristics of the digital signature scheme which prevent Alice from signing a message and sending it to Bob, but later claiming it wasn’t her who signed it after all¹.

The core requirement for non-repudiation is that Bob and Alice have a prior agreement that Bob can rely on messages digitally signed by Alice’s private key until Alice notifies Bob otherwise. Notification can be either direct or through a chain of intermediate Certification Authorities using Certificate Revocation Lists or some other agreed process. “Non-repudiation” means that Bob can hold Alice responsible for any such messages if Alice has *not* notified Bob that her private key is no longer secure enough for this purpose.

This places an onus on Alice to protect the privacy and the use of her private key, which is the root of all the problems which follow.

The contents of this paper can be understood by simplifying the arrangement to two parties, Alice and Bob, with no intermediate Certification Authorities.

THE FALSE REPUDIATION PROBLEM

Non-repudiation is *only* as strong as Alice's commitment and ability to secure her private key, as shown in the following scenario.

The false repudiation scenario

Alice has a private key and Bob has the corresponding public key. Bob is a large organisation, a government Benefit Agency, and it was Bob who guided Alice toward using a really first-class digital signature system.

Alice's digitally signed message signature is a declaration (to claim a benefit from Bob the Benefit Agency). This is in contrast to the more commercial type of message, which is effectively some type of authorisation (perhaps to debit her bank account, Bob being a bank or merchant).

Now suppose that Alice's declaration is later exposed as a false statement, which resulted in Bob paying Alice benefits to which she was not entitled.

There are offences with stiff penalties defined under the Benefits Act for making false statements, especially to get Benefits. Bob charges Alice with such an offence.

What does Alice say next?

Three possible answers

Well, Alice might be kind enough to say: *Yes, I lied. The digital signature was mine, I was responsible, and the message was a lie. I'm therefore guilty of an offence under the Benefits Act and ask only for leniency in view of my good standing in the community.*

But then again, it might also cross her mind to say something like: *Sure, it is a false message and it's my digital signature, but someone's trying to set me up. You see, I let my personal assistant use my private key, then there was this big argument, and she left on bad terms and you see...it wasn't me who signed that message. Ok, I agreed to keep my key secret but this is the real world. I agree that I broke my contract with Bob (to keep the key secret) and I'll pay the overpaid benefits back (as agreed under a pre-transmission contract with Bob). However, I'm not a criminal and I'm not guilty of any offence. No contractual agreement can make an action like giving away my private key a criminal offence.*

On the last point, Alice shows a disturbingly accurate understanding of fundamental legal principles.

Or perhaps she will say this: *Sure, it's false message with my digital signature on it. But it wasn't me who signed it. I tried really hard to keep my smartcard locked up but it's just possible that I left it out once or twice, and anyone in the office could have seen me keying my password to activate it. So that's what must have happened. Someone knew my password and also got hold of my card when I wasn't looking. They used it to sign a false message and put it back where I'd left it. I did try to do the right thing, but this system Bob expects me to use isn't really very secure. Anyway, you can have the money back, because I agree it isn't mine, but I haven't done anything wrong.*

Either of the last two answers leave Bob with a problem. Let's call it the *false repudiation* problem.

A successful criminal prosecution requires Bob to prove beyond reasonable doubt that Alice committed an offence, which will ordinarily mean proving that Alice personally signed the false message.

It may seem far-fetched to bring in criminal law at this point but it is not. Statutory offences, such as lying to the government, are creatures of the criminal law, and this affects most transactions between governments and individuals (and many bodies corporate).

Where false repudiation is a problem

Types of Message

Where the signature is an authorisation, Alice has a built-in incentive to keep the key secret, and repudiation will be rare (assuming a conventional non-repudiation infrastructure is present). In financial transactions, an "authorisation" usually means money moves from the Alice to the Bob. A prime case is a retail customer shopping on the Web.

When the signature is a declaration or claim, the signer may well be tempted to repudiate if the declaration is challenged as false or incriminating. In financial transactions of this class, money is probably moving from Bob to Alice.

Consider also information or professional advice given by Alice for a fee. The advice is a message signed by Alice the professional. Her advice may later turn out to be faulty or negligent, resulting in death or disadvantage to Bob the client. Alice the professional may then want to repudiate the message, just like an incriminating declaration. This is really a variation on the previous case—Alice signs the advice message in order to earn the professional fee, so it's a bit like a claim, with money again moving from Bob to Alice.

Types of Law

The three cases above correspond with three classes of law, Contract, Crime, and Tort (specifically negligence). These relationships are summarised in the following table. (Assume in all cases that there is a pre-transaction agreement with Bob that Alice protect her private key and will be "responsible" for messages signed with it, unless Alice has reported a compromise or revocation.)

| <i>Type of message</i> | Authorisation | Declaration | Professional Advice |
|---|---|--|---|
| <i>Example</i> | Authorisation of bank debit to pay Bob | Statement to a government agency to claim benefits | Pathology test results |
| <i>Movement of Money</i> | Alice to Bob | Bob to Alice | Bob to Alice |
| <i>Law class</i> | Contract | Criminal | Tort (Negligence) |
| <i>Level of proof needed by Bob to show Alice responsible</i> | Minimal—the pre-transaction contract provides that it will be up to Alice to show she <i>didn't</i> sign the message. | Alice signed it, beyond reasonable doubt. | It is more likely than not that Alice signed it. |
| <i>Motive for Alice to falsely repudiate</i> | Negligible, since she gains nothing. | Very high. The penalties for admitting responsibility far outweigh those from breaking the contract. | Depends on the relative cost of admitting responsibility. The low level of evidence Bob needs is also a factor. |

In the rest of this paper, the three types of message are referred to in shorthand as “authorisation”, “declaration”, and “advice” messages.

As shown in the table above, authorisation messages are hardly affected by the false repudiation problem. At the other end of the scale are declaration messages, where false repudiation is a serious threat. Advice messages fall somewhere in between.

Security of Private Keys: Incentives

In a normal electronic commerce situation, where the message signed is an authorisation, Alice is motivated to keep the private key secure. Alice has various means of protection available, with tradeoffs between cost, convenience, and effectiveness. All security regimes require a certain diligence on Alice’s part, and she will make a sufficient effort.

In the declaration situation, Alice, who we now know to be somewhat slippery, is actually motivated the other way. It is in her interest to let the key remain somewhat insecure, or make it appear possible it was used by another person. The mere possibility that the key could have been insecure is enough to facilitate repudiation if the declaration is challenged. It is in this case in *Bob's* interest to secure Alice’s private key so its use can be clearly attributed to Alice. This can be quite difficult for Bob if Alice is a 5000km away or if there are 500,000 Alices for each Bob. The ratio of Alices to Bobs can be very high if the Alices are individuals or small businesses.

The actual private key protection requirements and mechanisms are the same in both categories. They will include such things as:

- physically securing the private key token, perhaps inside a smartcard or on a physically secured magnetic disk;
- storing the private key encrypted;
- decrypting and activating the private key only by a passphrase known only to Alice; and
- never releasing or transmitting the private key itself, only unique digital signatures for specific messages.

The difference is in who is interested in maintaining security of the private key. If Alice has a strong incentive, and Alice has an appropriate set of security protections available, the key is likely to remain secure and uncompromised. If it is mainly Bob who wants to say that only Alice could have used the key, he has a problem. Alice can be careless and even deliberately share the key with other people, without Bob even knowing about it until he challenges a message signed with Alice's key—when it's too late.

It gets worse. Alice does not actually have to compromise her own key. In order to escape responsibility for a message, Alice has to only suggest that someone once *might* have used her private key. This someone else could be either someone she trusted or a sophisticated hacker on the other side of the world. Sophisticated hacking might be more or less improbable, depending on the security measures employed, but there usually remains a sliver of possibility even if Alice is entirely diligent with securing her key. An admission of a credible slight possibility is all it takes to allow repudiation. In a criminal prosecution, guilt has to be proved beyond reasonable doubt.

NON-SOLUTIONS

Standard PKI, even with smartcards

All the discussion above assumes the full operation the full weight of standard Public Key Infrastructure, including Certification Authorities and Certificate Revocation Lists. It even assumes the private key is protected in the best possible way, on a well-designed tamperproof smartcard.

All of this, complex and heavy as it is, will not work if Alice does not have a sufficiently strong incentive to control the use of her private key.

Contracts

Bob can insist that Alice sign an undertaking that she accepts responsibility for all messages signed with her private key. In the certified world, every key certificate is issued subject to such a condition, one way or another.

The trouble with this is that such an undertaking cannot be absolute. Alice can “accept responsibility” but this does not mean that in fact every message signed with Alice's key actually came from Alice. The well-established law of contracts is such that it means only that Bob is entitled to assume all such messages came from Alice, and it is Alice rather than Bob who bears any economic or similar loss.

Such losses are limited to the value of the original transactions involved. If Alice makes a false statement, gets \$100 as a result, and Bob finds out, the most Alice will forfeit, under a contract to be responsible for the message, is giving back the \$100. It is not even legally possible to include penalty clauses in contracts. Alice and Bob cannot agree in a contract that if Alice lets someone else use her key, she must pay Bob a pre-determined penalty. The most which can be contractually agreed is that certain benefits otherwise due from Bob to Alice under the contract are forfeited by Alice if she breaks the contract (for instance, by sharing her private key). Clever drafting of the contract can increase the forfeit to more than the value of the particular false messages, but an absolute upper bound on the forfeit is the total value for all messages from Alice to Bob.

These limitations directly limit the size of the incentives for Alice to protect her private key and take responsibility its use. These incentives will, under certain conditions, be exceeded by the penalties accruing to Alice if she *admits* responsibility for incriminating messages. Those conditions include nearly all cases where Alice is facing criminal liability, and also where Alice faces liability for negligence or any other liability far in excess of the original value of the transaction.

Electronic Transactions Act

The Commonwealth Electronic Transactions Act 1999 directly addresses the status of electronic messages and electronic signatures. The general approach taken by the Act is to deem an electronic message or signature to be equivalent to a written message or signature if it was reasonable to do so in the context. It does not deal directly with the false repudiation problem as outlined in this paper.

Similar Acts are being passed for other jurisdictions, all based on the United Nations Commission on International Trade Law (UNCITRAL) Model Law on Electronic Commerce of 1996.

This legislation is necessary for larger expansion of electronic commerce, but not sufficient to put electronic messages in place of writing for all potential applications.

In particular, it does not legislate that Alice is the author of all messages signed with her private key, but at most that Bob is entitled to act as if she were. If Bob happens to be a government agency, he can prosecute Alice if a message breaches particular legislation. However, the Electronic Transactions Act does not enable Bob to bypass the requirement to prove that it was actually Alice who authorised the incriminating message. Instead, it actually draws attention to this requirement by applying a reasonableness test. The exact words referring to the *acceptability* of non-paper signatures in the Commonwealth Act are:

“If, under a law of the Commonwealth, the signature of a person is required, that requirement is taken to have been met in relation to an electronic communication if having regard to all the relevant circumstances at the time the method was used, the method was as reliable as was appropriate for the purposes for which the information was communicated.”²

And the words on the *binding effect* of electronic signatures are:

“...[U]nless otherwise agreed between the purported originator and the addressee of an electronic communication, the purported originator of the electronic communication is bound by that communication only if the communication was sent by the purported originator or with the authority of the purported originator.”³

This is exactly what we don't need in the false repudiation problem.

Legislated responsibility for key usage

No just law could have the stronger effect of deeming Alice absolutely responsible for her private key, because Alice does not have the means to be absolutely certain at all times that her key is not used by someone else.

The nearest approach to this might be that taken in the Utah Digital Signature Act, which has been the subject of criticism on this very point⁴. The Utah act imposes “unlimited” liability on Alice if she does not take reasonable care to protect the use of her private key. The critics contrasted this “unlimited” liability with the strict liability limits enjoyed by consumers who lose their credit and debit cards.

Nothing like this has been done in Australia via the Electronic Transactions Act, nor by the (Commonwealth) Crimes Act as amended to deal with electronic data and communications⁵.

PARTIAL SOLUTIONS

Simple Biometrics

Biometric authentication directly undermines Alice's claim that “it was someone else”. However, the biometrics involved must be strong and secure enough that Alice cannot reasonably claim that it really *was* someone else, who maliciously defeated the biometric security.

Biometric authentication systems vary in strength according to the level of false positives and negatives obtained. Any biometric system tweaked for speed is likely to be weak. In particular, affordable fast systems which minimise false negatives will almost inevitably achieve this by permitting a relatively high rate of false positives.

However, that is not the root problem with biometrics. All biometric authentication systems, regardless of biometric strength, have fundamental limitations directly affecting the false repudiation problem for “declaration” and “advice” type messages.

Firstly, biometric devices have to be secured against tampering. This includes tampering by or on behalf of Alice or Bob, as well as by intruders. The party with interest in tamperproofing changes according to the type of message. In regular electronic commerce it's Alice; but where the signature is a declaration, it's Bob. This rule is the same as for non-biometric authentication, such as passphrase-activated smartcard. However, with biometrics the challenge of remotely securing a device against being fooled or tampered with is even greater.

Secondly, biometric devices transmit to Bob (via Alice's application software) either biometric data itself or a signal indicating that Alice's body part matches a stored template. Either of these signals can be captured and re-used on other messages, and thus have an inherent level of repudiability.

There are some complex solutions to the second problem if the first can be considered solved. These are discussed later in this paper.

Signature Dynamics

I find biometric systems based on digitising single conventional pen-based signing events particularly attractive. They're cheap and not socially objectionable, being based on the established convention of signing a piece of paper with a pen. Rather than press one's retina to a laser beam, or submit to being fingerprinted by a plastic box, one authenticates oneself to a signature dynamics system by signing a piece of paper with a special pen, within range of a base unit which does the biometric assessment.

These devices normally go under the name of "signature dynamics" to distinguish them from the weaker systems which use a bitmapped or other purely visual image of someone's paper signature. Signature dynamics systems record the movement of the pen in space, including the speed and pressure used at different moments in making the signature. Current signature dynamics systems are already cordless⁶, with a pen that looks just like an ordinary pen, and the cost per unit is in the low three-digit range. This price range is probably just enough to warrant putting the cordless pen on a chain anyway.

The biometric strength of these systems may or may not be outstanding. However, for the sole purpose of defeating the suggestion that "it was someone else", a high level of *biometric* accuracy is not required. False positives of one in a thousand (a relatively weak level) would probably work well enough to convince a jury beyond reasonable doubt that it really was Alice (all other uncertainties aside).

This level of certainty does not correspond with the cryptographic strength needed for a *digital signature* in electronic commerce. We have declared the odds of a false positive for the biometrics to be 1 in 1000 or about 1 in 2^{10} . Stronger biometric systems claim false positives of one in a million, or about 1 in 2^{20} . Digital signatures conventionally have odds of a false "match" of 1 in 2^{128} , or 2^{108} times less likely. Later in this paper is a brief explanation of how weak biometrics can work with strong digital signatures to create potentially viable strong non-repudiation schemes.

Crimes Acts

New Crimes Acts (specifically the Criminal Codes) are being developed, which specifically address crime in the digital age. These Acts are based on a Model Criminal Code developed by the Australian State, Territory, and Commonwealth Governments in the early 1990's. The relevant parts of the Model Criminal Code include such helpful features as defining fraud to include "deceiving" an electronic system, and theft to include improperly causing money to be transferred between accounts. There are also some defined offences which avoid the need to prove specific personal actions in causing a deceit or loss⁷. The proposed Commonwealth Crimes Act includes a low-level offence for obtaining a financial advantage knowing you are not eligible⁸. It does not appear necessary to prove the specific act of dishonesty leading to obtaining the advantage.

Therefore, in particular relationships and transactions between Alice and Bob it may be possible to use the Crimes Act or other criminal law to ensure Alice is sufficiently deterred from sending a false message to Bob for improper gain. This might work by engineering business processes to ensure that no matter who sent the message, only Alice receives any consequent benefit and can be in no reasonable doubt about where it came from and that she was not entitled to it. Therefore, if Alice fails to return the improper benefit within a reasonable time, she would be clearly guilty of an offence which does not require proof that Alice was the author of the false message which triggered it. In such a case the digital signature is not then essential proof of the origin of the message, but only part of a context which in total puts beyond reasonable doubt that Alice knew where these benefits were coming from.

An over-controlled example

Suppose Bob is a government Benefits Agency and Alice routinely claims Benefits by digitally signing declarations of entitlement (i.e. Benefit claims) and sending them to Bob.

Before agreeing to this arrangement, Bob insists that Alice open a specific bank account to receive Benefit payments, which only Alice can operate. If there are enough Alices and Bob pays them enough Benefits, Bob can actually have an agreement with the bank about the operation of this type of account. The account is called "Alice's Benefit Account" so there is no possibility of confusion.

Bob also insists on a face-to-face meeting with Alice in which Bob explains the scheme carefully.

The scheme includes several obligations on Alice. Firstly, Alice will be responsible for the use and protection of her private key, and will immediately report any doubt about its security to Bob. Secondly, Alice agrees that each withdrawal she makes from "Alice's Benefit Account" is an assertion that she authorised all claims contributing to the balance and that she is entitled to the funds. Not only that, Alice agrees that letting someone else operate the account will cancel the entitlement. Thirdly, each declaration Alice signs with her key.

These understandings are put into a triplicate paper contract with Alice, which Alice signs in front of a witness. The witness also signs the document to certify that it was Alice who signed it, and she clearly understood the contents.

Thereafter, Alice starts claiming Benefits from Bob with her digital signature. Each declaration has a unique identifying number. Each declaration includes a clause confirming the *last* declaration, the details of which Bob's computer inserts into the text before Alice digitally signs.

Bob's computer responds to each claim by validating the signature, assessing the contents, and depositing the due amount of Benefit into "Alice's Benefit Account". Each deposit record includes an electronic annotation of the related claim identification number, to help Alice reconcile the account. On an annoyingly frequent basis—at least once per deposit—Bob's computer reminds Alice of the significance of withdrawals and the need to reconcile first.

After six months, Alice includes a false declaration in one of the claims, with the intent of getting more Benefits than she is actually entitled to. Bob's computer pays the Benefits as if the claim were true, and the matching deposit appears in "Alice's Benefit Account", complete with the claim identification number. Alice withdraws funds from "Alice's Benefit Account" several times over the next three months, and then Bob discovers Alice's lie.

Bob charges Alice with the offence of obtaining a financial advantage knowing she was not entitled to it. Alice can say that it wasn't her who used her key to sign the false claim, but this doesn't help her. She has already attested that she signed the claim, in the act of withdrawing the funds. The best Alice can do is to "admit" that not only did someone else use her key, but she also failed to reconcile the account before withdrawing the funds, disregarding the recent warning Bob gave her about the significance of the act of withdrawal.

Alternatively, Alice can say that someone else made the withdrawal from in "Alice's Benefit Account", which sounds hauntingly similar to something she said in the original false repudiation scenario. However, this time it puts Alice in a virtually untenable position. She has broken the agreement several times over. She has already agreed that she has no entitlement to the funds. More than that, she willingly let someone else use her bank account, which means she knew what was going on.

Bob's case against Alice is not absolutely watertight, but it is a reasonably strong collection of circumstantial evidence from which Alice will find it hard to escape. In effect the original digital signature is corroborated by the authentication Alice submits to the *bank* when withdrawing funds.

This approach may not be applicable to every situation in which the false repudiation is more attractive than the consequences of admitting authorship of a false or otherwise incriminating message. It will require that the specific possible offences are carefully anticipated, and it also requires a defined relationship based on at least some communications between Bob and Alice *outside* the electronic message channel. In the example, this was the initial meeting, contract signing, and special bank account that Bob insisted on, plus the element of communication via Alice's bank.

Crimes Acts are only relevant where the false message is "incriminating" in the strictly criminal sense. It is difficult to see how the approach can work where negligent professional advice has been in an electronic message and the client suffers (that is, in the law of tort rather than in the criminal law). However, one approach for Bob in that case might be to argue that sharing your private key is negligent in itself.

Specific legislation defining sufficient evidence for particular offences

This is a sort of brute-force legislative approach which might work for some government agencies with strong political support.

The method is to re-define the specific offences in the relevant legislation (in our example, the Benefits Act) to include specific combinations of circumstances, regardless of who actually did what to cause them. For instance, the Benefits Act might declare Alice guilty of an offence if she withdraws Benefit moneys and does not within a prescribed time produce a prescribed form of evidence to support the related declarations. This definition simply bypasses the question of who authored and signed the message, also the question of whether the message was actually true or false. Therefore, it greatly reduces the amount of evidence Bob needs to find in order to get Alice convicted of an offence.

Such legislative approaches might be effective, but they are liable to be controversial, cutting across generally accepted common-law criminal principles. They also go against the clear intentions of the new Crimes Acts based on the Model Criminal Code, which actually remove statutory offences from specific legislation and replace them with all-purpose offences in the Crimes Act⁹.

Multiple function keys

In the above example Alice only ever uses her key to sign declarations claiming benefits from Bob. Probably Bob issued her with the key.

Suppose instead Alice also uses her key frequently and for a wide range of other functions, such as withdrawing funds from her bank accounts (including, but not exclusively, "Alice's Benefit Account") and reading personal encrypted e-mail.

In such a case Alice can still falsely repudiate an incriminating message to Bob. However, it is much less likely that she was actually careless with it or let another person use it, because that would work directly against her interests in respect of the other uses of the key. If she said she did, this is significantly less credible than in the case where she only uses the key to sign declarations claiming benefits from Bob.

Generous Opportunities to Repudiate

The over-control example above includes another element which is useful in isolation. Alice has several different opportunities to detect and repudiate bogus messages before Bob.

1. Alice was shown the previous use(s) of the key in every new message she signed.
2. She was able and obliged to reconcile deposits to “Alice’s Benefit Account” before making any withdrawals.
3. Even if she missed that opportunity, the reconciliation could still be done after the withdrawal, before Bob challenges any messages.

These opportunities could be extended further, for instance by giving her easy access to a detailed, memory-jogging log of all signing events, complete with cross references to the source records for events leading up to the signed declaration.

The greater the opportunity and obligation to detect unauthorised messages, the less credible it is for Alice to suggest that someone else was responsible.

It is helpful if Bob queries *immediately* any message from Alice which looks out of the ordinary. This is a courtesy to Alice in the event of actual third-party intrusion, and it also encourages Alice to withdraw any fabrications before harm is done (e.g. before Bob pays her any Benefits to which she is not entitled). This may not result in a lot of prosecutions, but it will discourage Alice from trying her hand at simply constructed false declarations. She may decide that it isn’t worth trying, given the amount of effort involved in concocting a credible message which results in a worthwhile level of extra Benefits.

Witnessing

Evidence that Alice signed a document is much stronger if the signature was witnessed. The witness, named on the document, can be found, and can testify that it was actually Alice who signed. This principle is equally valid in the world of electronic signatures.

It works well in both worlds subject to a couple of fairly inconvenient conditions.

- Firstly, the witness must be a person who is independent of Alice and will neither “witness” a signature without seeing Alice sign, nor lie if Alice falsely repudiates the digital signature. For instance, the witness cannot be Alice’s personal assistant, who will presumably “witness” everything Alice leaves on her desktop, regardless of who actually saw what. Neither can the witness be someone on Bob’s payroll.
- Secondly, the witness must be physically present and observe Alice digitally signing the message. Witnessing cannot be done after the event, and doing it remotely (perhaps via video link and remote conferencing software) is highly questionable, given the ease with which this process could be manipulated.

Together these conditions tend to negate the benefits of using electronic messages and digital signatures in the first place. However, witnessing may be useful in some situations.

FULL SOLUTIONS

Combining Weak Biometrics with Strong Digital Signatures

The goal is to create a digital authenticator on a message which cannot be forged by a third party, even if Alice is careless with her private key or other authenticating knowledge or token.

The following combination of biometrics and digital signature can work but involves several unpleasant pre-conditions.

- The biometric unit is issued with its own certified public/private key pair.
- The biometric unit cannot be tampered with, that is, cannot ever be induced to disclose either a spurious “match” message, Alice’s biometrics, or any private keys stored inside the unit.
- A certified attribute of the digital certificate for the biometric unit is that it meets a recognised level of tamperproofing.
- Software on Alice’s computer is secured, to prevent unauthorised disclosure of keys, interception of data, or other irregular actions by Trojan horses or other malevolent code.

How it might work

The message is prepared by Alice’s application program and displayed on screen to Alice. At the bottom of the screen Alice is prompted to sign the message with the special signature-dynamics pen (substitute your favourite low-cost biometric device).

At the same time, the application or a utility program creates a corresponding message digest, along with a secure timestamp and Alice’s distinguished name as it appears on her Public Key Certificate. The application sends this compound message to the biometric pen unit.

Alice signs with the biometric pen. The pen captures the metrics from this signing event, and compares these with an internal specimen (or several specimens). If and only if an acceptable match is achieved, the unit makes up a message from the original message digest, the timestamp, and the match decision. The message says in effect, “I checked Alice’s signature for that message [digest] at time T and found it good.” This message is itself digitally signed by the pen unit using its own private key; which means extracting a message digest and encrypting it with the private key. Call this the “pass” message.

The biometric unit then transmits the digital signature for the “pass” message back to the application program¹⁰. This signature cannot be used with a different message (because it is derived from a near-unique message digest) and cannot be replayed (because it is also derived from the time stamp). It also proves it was Alice who signed, because her distinguished name was involved.

The application verifies the digital signature on the pass message from the biometric unit using the certified public key for the unit. If the verification is successful, the application knows the biometric authentication was successful and has a signed pass message from the biometric unit to prove it.

The original application message is forwarded to Bob along with the pass message from the biometric unit to authenticate it. Bob will then know that the certified biometric unit saw Alice sign the original message, which may be all he needs. In effect, the biometric unit in Alice's office is an *automated notary or witness* trusted by Alice, Bob, and any court which might be later involved in a dispute or prosecution.

A peculiarity of this scheme is that Alice's own private key is not used. Variations on this scheme are possible and may be more suitable for particular applications. For comparable schemes and a discussion of the issues surrounding integrating biometrics and digital signatures, see the article by Jueneman and Robertson¹¹.

In summary, combining biometrics and digital signatures can more or less solve the repudiation problem, but it is complex and the pre-conditions are very difficult to meet. Buried in the vague condition that the application software be "secured" is a requirement that no person or process can substitute the digest of an unauthorised message for the digest for the message displayed to Alice. This is in itself a daunting problem, shared with other digital signature configurations.

Paper confirmation of earlier electronic signatures

Alice and Bob transact regularly via digitally signed electronic messages. Periodically, Bob requires Alice to confirm by paper signature (or some other non-electronic authentication method) the collection of messages digitally signed with Alice's private key since the last confirmation.

To facilitate this, Bob prepares a summary of those messages for Alice to confirm, and has preferably ensured, as part of the contract, that Alice maintains adequate records to supplement her rather selective personal memory. (If Alice is a good girl who genuinely looks after her private key and always tells the truth, she can confirm all messages *without* consulting any records, because she already knows she was the only one who used the key.)

There is still paper flow involved in such a system, violating the ideals of electronic commerce. However, the volume of paper handling can be reduced by a factor of somewhere between, say 10 and 1000.

Messages digitally signed by Alice but not yet confirmed on paper are still subject to false repudiation. Bob needs to monitor this risk and adjust the frequency of confirmation events accordingly. If Bob has a lot of confidence in Alice (or is able to monitor her behaviour fairly closely), Bob might choose to demand from Alice confirmation of only on a small sample of her digitally signed messages. When Alice reveals her slippery nature or starts claiming a higher level of Benefits, Bob increases the proportion and frequency of confirmations.

Substitute contractual payments for statutory entitlements

The false repudiation arises where the substance of the transaction, and penalties for “incriminating” messages, are determined by statutory rights and statutory penalties. This puts the process in the realm of the criminal law, in which proving Alice’s responsibility involves stringent evidence requirements.

A radical line of solution is to move the entire business out of statutory and criminal law, into commercial contract law. This might work where there is an active and continuing relationship between Alice and Bob; it is less practical where Alice only talks to Bob once or twice a year.

With Alice the claimant and Bob the Benefits agency, it works like this. Alice has a statutory entitlement to Benefits if she makes a valid declaration to Bob.

This is balanced with statutory (criminal) penalties if she makes a false declaration to Bob. Realising that this precludes or complicates the efficient use of electronic commerce in Benefits administration (because of the false repudiation problem), Bob offers Alice a deal. Bob allows Alice to make declarations electronically, but Alice has to assign her statutory rights to Benefits to Bob and accept a contract-based payment instead. The amount of the contract payment may be equal to the amount of Benefits otherwise payable, or even slightly greater, sharing the benefits of electronic commerce.

Also, Alice is no longer under threat of statutory penalties under criminal law.

To balance the relationship, Alice contractually agrees that Bob can sharply reduce the total contract payment (or even terminate the contract) in the event of defined circumstances, which include any case where there is a message signed with Alice’s key which she cannot substantiate when challenged. The amount of the reduction is somewhat higher than the value of the single challenged message, but necessarily no more than the total value of all messages over the period of the contract. (This is a fundamental legal constraint on contracts.)

Bob also includes clauses about his rights to challenge any or all of Alice’s messages, and he will actually challenge a lot of them after finding the first false message. The method of determining the truth or falsity of a message is defined in the contract. Further false messages incur an increasing forfeit of contract payments on a sliding scale, up to total loss of payments when a predefined level of falsehood in Alice’s messages is identified. The terms of the original assignment of Benefits should be such that the original statutory entitlement is not re-instated automatically or otherwise on cancellation of the contract.

This will be a good deterrent against Alice generating false messages to increase her Benefits, as long as Alice has a fairly high level of legitimate Benefits entitlement and needs the Benefits on a regular basis, e.g. as her main source of income.

It is not automatically legal to assign or waive statutory entitlements. In fact, it is probably not permitted under the existing Benefits Act. However, the amendment to the Benefits Act to enable this scheme (thereby enabling electronic commerce) may be quite minor and politically acceptable, under the circumstances.

SUMMARY

False repudiation threatens many potential applications of digital signatures in electronic commerce. A quick vulnerability check is to ask whether the message is in effect an authorisation, a declaration, or an advice, and then ask which direction the money is moving. If the message is a declaration or advice, and the money is moving towards the signer, there is probably an exposure to false repudiation.

There is no single solution to the false repudiation problem which can be implemented as part of Public Key Infrastructure. All solutions are specific to the business process or application. The stronger solutions substitute or supplement the digital signature authentication by some other authentication of the message. This might be a biometric authentication, a paper confirmation, withdrawing or using the proceeds of the transaction, or just continuing to use the private key. The amount of supplementary authentication needs to be considered relative to the applicable risks and incentives.

Biometrics seem very promising, but straightforward deployment of technology is subject to comparable threats of manipulation in the sort of environments where false repudiation threatens. Countering these threats is complex and barely feasible in 2000.

Changes in the law addressing electronic commerce are a significant help, but so far do not provide any direct solutions for clearly-motivated false repudiation.

For the time being, the best bet may be to consider having digitally signed electronic messages confirmed in writing, on a summary or sample basis, thus reducing rather than eliminating paper document flow. This may be a dirt section of the information superhighway, but it may at least enable the main bulk of individual messages and data flow to be processed automatically and without paper.

So this paper is mainly bad news. I hope it's wrong and the problem of false repudiation can be solved after all. I have my fingers crossed in the digital age.

DISCLAIMER

The opinions in this paper are solely those of the author. The examples used do not reflect actual situations in the Health Insurance Commission or elsewhere.

NOTES

¹ There is also a concept of non-repudiation of messages *received*, covering cases where Bob sends Alice a message and Alice later claims not to have received it. This class of non-repudiation is not within the scope of this paper.

² *Commonwealth Electronic Transactions Act 1999*, Section 10(1)(b), available from <http://law.gov.au>.

³ *Commonwealth Electronic Transactions Act 1999*, Section 15(1).

⁴ Utah Digital Signature Act, cited in R. R Jueneman and R. J. Robertson, Jr., *Biometrics and Digital Signatures in Electronic Commerce*, Jurimetrics Volume 38 No 3, and elsewhere. Criticism is cited in Chapter 3 of *Electronic Commerce: Building the Legal Framework*, report of the Electronic Commerce Expert Group to the Attorney General, 31 March 1998.

⁵ See the *Commonwealth Criminal Code Amendment (Theft, Fraud, Bribery And Related Offences) Bill 1999*, available from <http://law.gov.au>. This Bill had not become law at time of writing.

⁶ Current cordless units use radio transmissions between the pen and a base unit, and the base unit does the biometric matching. Regardless of whether a cord or radio transmission is used, if the signals include actual signature dynamics data, these signals can be captured and re-used on other messages. The solution is to digitally authenticate and encrypt the signals between the sensing unit and the base unit. This detail has been left out of the discussion to minimise complexity.

⁷ See the *Commonwealth Criminal Code Amendment (Theft, Fraud, Bribery And Related Offences) Bill 1999*, and Explanatory Memorandum, available from <http://law.gov.au>. Particularly useful sections are 132.1, 133.1, 134.1(9) to (11), 135.1, 135.2.

⁸ Section 135.2.

⁹ In certain government activities the relevant Act supports a system of regulations defining low-level non-criminal offences and penalties. An obvious example traffic and parking regulation. My non-expert understanding of these arrangements is that the guilt and penalties for these non-criminal infringements can be disputed, and in that case they can only be enforced with backup support from the common and legislated criminal law. In other words, it is not possible for administrators to simply invent and police offences which conveniently require minimal evidence.

¹⁰ In fact it is not necessary to send back the whole pass message because the content is already predictable to the application program (in the event of a positive biometric authentication). This might be important if communications between the biometric unit and the application process are slow.

¹¹ R. R Jueneman and R. J. Robertson, Jr., *Biometrics and Digital Signatures in Electronic Commerce*, Jurimetrics Volume 38 No 3.