A networked registration scheme for enhancing trust

J Adrian Pickering and Christopher J Gutteridge
School of Electronics and Computer Science, University of Southampton,
Southampton, SO17 1BJ, UK
jap@ecs.soton.ac.uk

Abstract

Society has widely adopted use of electronic data without sufficient attention to the problems of non-repudiation (NR). A universal, transparent scheme is needed to replace the traditional paper-based model that people are familiar with. A registration scheme is proposed that uses a network of registration servers run in a way that is robust to legal and technical challenge. Any user can register potential electronic evidence with one or more of these servers. This enables a user to later assert that they had the data at the time. Applications encompass intellectual property (IP) protection, file-download-based e-commerce and corporate shareholder communications. Wide availability should induce proper behaviour between parties whether they use the scheme or not.

1. Introduction

Traditional, paper-based methods of working are being gradually overtaken by the exchange of data in electronic form. This pervades all that we do now both as ordinary citizens and as people at work. Most have not appreciated that doing business electronically is not just an analogue of the paper model. Even fax was still anchored on paper since there was a requirement that there be an original.

Society is fast embracing web-based methods of doing business. There is rarely any durable, shared, original paper document whose content determines a future decision. It works by the parties setting a trust threshold by making a risk assessment. So far there have not been any widely-known incidents where this trust has been seriously undermined. However, there is a problem that needs attention. Such socially-useful, electronic communication systems must embody easy-to-use, legally-robust, trust-supporting mechanisms.

The recent near-collapse of banking systems is a warning. Those trying to recover their assets may be finding the voids in the current trust mechanisms. Typically, the only evidence the claimant would have, if on paper at all, would be their own printout of data purporting to be from the disputing party.

A key concept in doing business is that each party declares its position in a way that is non-repudiatable (NR). The subsequent action by the counterparty is constrained if it knows that NR techniques are in use. Proper behaviour is induced because of the implicit threat of legal action using robust, NR evidence. In the past this state would be captured on paper and openly marked (signed) by those bound by its content. When the parties become virtual and timeliness is crucial, some other mechanism is required. Maintaining trust in these circumstances is the challenge. Fax worked: it emulated sealed post, but technology and society have moved on.

Intellectual property (IP) records are not so obviously the basis of a contract but they are subject to the same standards of evidence. Today, experimental data is extensive and electronic and, when created, it is not clear what elements may be important. The 'log book', electronic or paper, is not suitable anymore. There needs to be a robust and economical way of enabling NR over all of this data.

2. Trust requirements

A court is bound by its rules of evidence. Electronic evidence has had difficulty in being regarded as reliable, and demonstrating its robustness is still a challenge. Accordingly, any business transacted on the basis of electronic evidence is more at risk. Transactions still resort to paper when their level of risk demands robust evidence.

A way of creating reliable evidence is to have it widely witnessed at the time it crystalises. We are familiar with signing ceremonies to mark the acceptance of a significant contract, agreement or treaty. This exploits the mechanism of 'open declaration' to effect NR, often assisted using the media. Most of those observing will not know the content of the agreement: the trust exists because the mechanisms are executed openly and are easily understood. Future reputations and business depend on a successful outcome.

Contract and IP data often need to be confidential. When electronic, the physical control mechanisms familiar with paper are not so easy to enforce, particularly within networked, virtual communities. Encryption of electronic data is technically well-understood but still not easily useable. Encryption can support NR but, as will be shown, NR does not require the use of encryption. NR is also agnostic regarding the meaning of the data, encrypted or not.

A dispute signifies a breakdown in trust that needs to be reestablished. The evidence is required in order to openly replay the circumstances that are relevant to the dispute. All those involved (in court, jury and observers) must be convinced, beyond reasonable doubt, of each of the steps taken to reach a verdict. Before this, just the availability of robust evidence can shorten or even preempt a dispute. If the means of creating robust evidence is readily available, then society's trust foundations will be sounder. The result will be doing business confidently, faster and with fewer, costly disputes.

3. Current schemes

The idea of using cryptographic hashing algorithms to assist time-stamping a digital document was published by Haber and Stornetta [2]. The motivation was the ability to easily declare the existence of a document to a third party without disclosing its content. Cryptographic hashing algorithms are designed to produce a short 'digest' (or 'hash') of a digital document (file) which is (a) collision-free i.e. no two documents will generate the same digest and (b) it is infeasible to synthesise a collision i.e. generate another document that has the same digest as another. Since the digest is, in general, shorter than the original document, there is less information there than in the original. Together with the nature of the algorithm, this means nothing can be construed about the original document from its digest. The accepted cryptographic hashing algorithms are public and are continuously subject to scrutiny by cryptanalysts since they underlie electronic signing mechanisms.

The principles have been used in a number of registration systems, notably the digital notary service operated by Surety [5], which is based upon Haber and Stornetta's concept and patents. Also, since 1995 a UK Jersey-based company has been operating its 'Stamper' time-stamping service based on PGP signing (IT Consulting [4]). More recently, in the UK, Codel have been promoting their Codelmark service [3].

The Surety and Codel services are both subscription-based services. Subscribers need to have faith in the company and trust that their processes are rigorous since the underlying registrations are not open to users' scrutiny. Daily, they digest the registration data and publish the resulting 'master hash' in a newspaper of record (Codel publish in the Financial Times). 'Stamper' attempts to be more open by publishing its signing summaries on the web and over Usenet.

Since the users' digests are not disclosing anything, it is not clear why it should not be possible to openly declare the registrations. This reduces the trust barrier to using the system: the users can see their registrations, their context and watch the scheme function. Further, if they are concerned about the robustness of the service, they can take copies of sufficient data for safe-keeping elsewhere.

Though the digests do not declare anything interesting about the user, other data recorded with the digest could. Subscription services need to know whose data it is in order to secure their income stream. They can undertake traffic analysis on registrations and assign it to users. Even if this data is not made public, the registration service needs to be trusted not to ever misuse this data. The only sure way of avoiding this risk is to allow anonymous requests for service. Indeed, there are compelling human rights and citizenship reasons why anonymity is desirable.

Thus there is a need for an openly-available registration scheme whose only function is to accept and publish sequences of digests received from anonymous users. If the user needs to assert that they, or their company, are the only owners of the data at the time, then it is for them to incorporate some secret in the data before its digest is registered (e.g. using a signed HMAC, Eastlake and Hansen [1]). That is an optional, separate issue from registering the existence of the data.

4. The registration scheme

Open declaration of evidence is optimal. Digest hashes are of fixed length, short and are not expensive to store. Publishing these is cheap and web technology provides the ideal 'notice board' where the public can observe them.

Anyone with an electronic document that they wish to register posts its hash on a registration server of their choice using a protocol that supports anonymity. Because of the properties of cryptographic hashing algorithms, the registrand should be able to demonstrate later that only they had the means to create the registration at that point in the journal. The registration server's vital task is to journal the registrations chronologically. The server can annotate the registrations with other data. An obvious and useful choice is a timestamp. However, such timestamps are only indicative: they are not essential. The scheme fixes the time order of registrations, which is necessary and sufficient.

The power and scaleability of the scheme lies in realising that a registration server is itself generating material that needs evidential protection. Thus, periodically, it hashes a journal segment and registers that with another disinterested server.

Provided there are sufficient, randomly cross-registering, independently operated servers, the time order of registrations across the server network can be adequately resolved. Any timing information embedded in the journals is useful, supplementary evidence, particularly if the timestamps (or other form of time anchor) come from reputable sources.

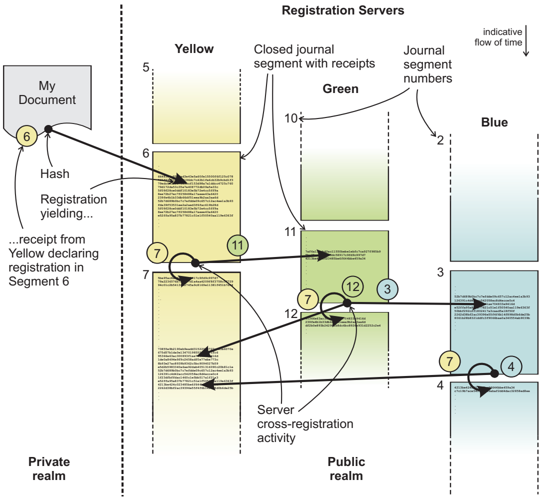

Figure 1. Registering among net-connected servers

As an example, Figure 1 illustrates what happens when a user registers ‘My Document’ with their chosen ‘Yellow’ server. The hash the user generated for My Document is stored with other registrations in Segment 6 of Yellow’s journal. It returns a receipt to the user that is kept with My Document. Shortly later, Yellow closes its Segment 6 and registers its hash with itself in new Segment 7 and the Green server. This process continues at indeterminate times among other, disinterested, cross-registering servers. Note that the effect of the registration (network ‘flow’, in graph terms) is felt twice in Yellow Segment 7. This achieves stronger connectedness between the variously-owned segment nodes and distinguishes this scheme from those that use hash trees to store registrations e.g. [5].

If the user wishes to assert they were in possession of My Document at the time claimed the registration process is replayed but comparing the registration with the evidence in Yellow server Segment 6, as indicated by the saved receipt. For clarity, the illustration shows the user making just one registration with one server. For robustness, it is wise to register with at least two servers.

Once the hash of a server journal segment is released to the network it will rapidly be subsumed within further cross-registrations. This locks all the dependent data into time order. No other data on which the hash depends can be altered without potential detection.

Many, unknown users—notably, the registrands, will observe the scheme. Further, these will be taking copies of relevant segments so that they (or their representative) can replay the algorithms later if required. Server operators will not know their users through normal use. Any inexplicable post-hoc alteration or loss of journals will cause potentially irreparable damage to their reputation.

Since there must be formal disinterest between servers and clients, there cannot be any service contract and, therefore, the service must be free. This poses some security and resourcing challenges. Fortunately, hashes are small, and networks and storage are comparatively cheap. We also have precedents for the evolution of large, mutual self-interest based systems—Internet and email.

5. The scheme capabilities

The holder of some potential evidence would now have complete freedom to register its hash with servers of their choosing. Since services might disappear, it is wise to register with several. Service operators will have a service policy that would guide a user in their choice. Features would be (a) what information is recorded with the hash (b) how often journal segments are closed (c) server cross-registration (graph connectedness) (d) time anchor policy (e) number of registrations per day accepted from a particular Internet source (f) open-source heritage (g) third-party validation (h) certified retrieval services etc.

It is always the users’ responsibility to have the means to replay the registration algorithms to a third party to prove that they possessed the data at issue at the time claimed. This is why it is important that the scheme is open and simple. In a dispute, it is still for the parties to interpret the meaning the data that was registered. What will not be in doubt is (a) document possession and (b) its time-order context.

6. Scheme usage scenarios

There are many everyday applications where voluntary registration of documents would enhance trust. Citizen-consumers are being asked to react to documents that are only published on web sites. To be sure that the document does not change without notice, either party can register the document on which they are basing their business decisions. For NR to be credible, a disinterested third party must hold the registration. And just having the ability to make registrations enhances trust between parties.

There is considerable commercial pressure to move to paperless banking and billing. Traditionally, NR relies on the infeasibility of undetectably altering or forging paper documents. Instead, the issuing party would register the electronic document when generated, possibly with some embedded secret.

Once third-party registered the document cannot be subsequently altered without giving explicit notice. The recipient would be told the locus of its registration so that, upon receipt, the registration process can be replayed. The recipient can then be sure that what they received the sender will not repudiate—the data was what the sender wanted to communicate at the time it was issued. The recipient may subsequently wish to register what they received, thus capturing the time path of the document. Of course, both parties must keep the data in a way that they are able to replay the registration process to a third party evidence reviewer. Figure 2 shows how the two parties could act to ensure fair play. The document here could be a shareholders' report, bank statement or contract.

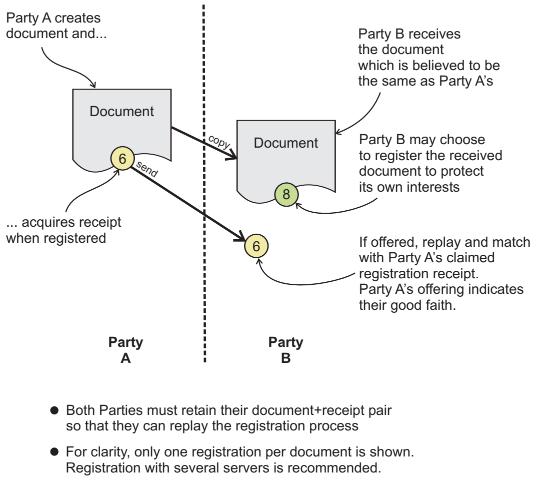

Figure 2. Confirmation of the state of a document

Figure 3 shows how a research organisation can run a local registration service to concentrate the registration traffic before using a third party service. The advantages of doing this are (a) encouraging the IP developers to use the service regularly and (b) mitigating traffic analysis by outside parties. Because the primary server is internal to the organisation, its evidential robustness would be more in question. However, the IP Administration can dynamically adapt its own registration policies to suit the risks.

Digital rights management (DRM) is where the document can only be rendered with permission of the owner (e.g. music, video, cartography). This requires the cooperation of the data renderers regards encryption system management. DRM has so irritated end-users that many schemes have been willfully broken. An alternative is to register the IP as it evolves and gets distributed rather as it is done in other IP disciplines. This establishes prior-art and any further registration will establish a time path, perhaps through licensed, selling intermediaries. Any challenge along the path will require the holder of the document to demonstrate that they acquired the necessary rights. Figure 4 shows how this may be achieved. The principle is that any holder of IP will have a matching document that shows that they have the right to hold and use it.

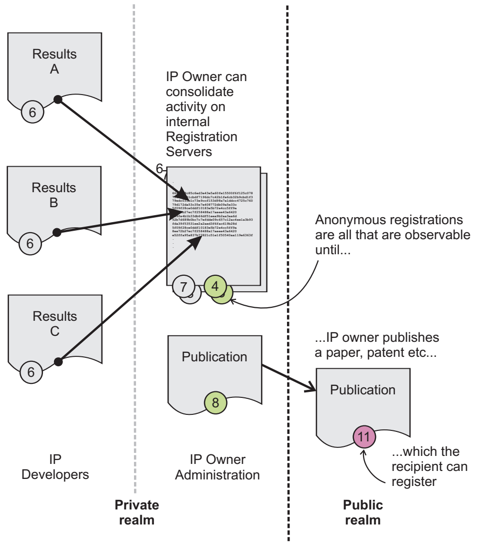

Figure 3. Gathering IP prior-art evidence before publication

Clearly, the digital product data may have many incarnations/versions with different hashes. It will still be for the reviewer to judge whether a material breach of rights has occurred. Here, registering does not change the process of disputing misuse, but it does make 'copyright depositing' very easy for everyone—corporation or citizen. Because either party can freely register their data, those who believe they rightfully own a copy will want to register it. Not registering a copy will question the individual's motives. Just knowing that anyone can cheaply and robustly assert possession of some data, be it their own or a fair-use copy, should encourage proper behaviour.

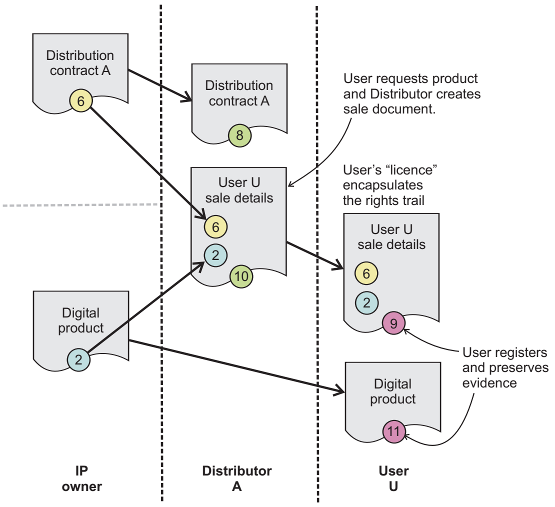

Figure 4. Selling a digital product through an intermediary

7. Conclusion and next steps

An open scheme that should enhance the ability of citizens or workers to collaborate with trust is proposed. The time is right for its widespread adoption so that it can start to enable the benefits claimed. A very simple demonstrator can be accessed at Probity [6] which readers are invited to try. This uses HTTP to effect registrations.

Work is in progress in developing open-source client and server prototypes to be used within the UK eScience support infrastructure (notably ePrints and MyExperiment). There is scope for 'added value' in embedding the registration primitives within tools where NR could be useful. Those with server resources and interest are invited to join the effort and get the scheme working for society's benefit.

8. References

[1] D. Eastlake and T. Hansen, "US Secure Hash Algorithms (SHA and HMAC-SHA)" IETF RFC 4634, July 2006

[2] S. Haber and W.S. Stornetta, WS "How to time-stamp a digital document", *Journal of Cryptography*, Vol 3 No 2, 1991, pp 99-111

[3] R. Hill, "The Codel Authentication System" Codel Technical White Paper, Codel Ltd. *www.codelmark.co.uk/white-papers*, 14 February 2007

[4] IT Consulting, PGP Digital Timestamping Service *www.itconsult.co.uk/stamper/stampinf.htm*, 2002

[5] Surety, *www.surety.com,* checked 2009

[6] Probity, Demonstrator at *www.probity.org/demo,* 2009