

Webs of Publishing and of Trust

Patrick Juola // Duquesne University // Pittsburgh, PA 15282
juola@mathcs.duq.edu

1 Introduction

Whether or not scholarly publishing is in “crisis,” the price of scholarly publishing, especially the journals and other serials, has been rising dramatically. Recent statistics published by the Association of Research Libraries show that serial prices have risen 226% since 1986, almost four times the corresponding CPI increase of 57%. Staggering under these costs, libraries have been forced to find and invent new ways to reduce costs while still providing high-quality research support. One often-named culprit is the “exorbitant subscription prices” (Suber, 2002) charged by some journal publishers, while publishers themselves cite rising costs of conventional production, both of which are prompting a significant trend away from traditional paper-based publishing, to an electronic model. In theory, this model can provide cost advantages while retaining (some) traditional advantages, such as peer review, of conventional scholastic publishing.

(Suber, 2002) points out that, although peer review itself costs money, open-access electronic journals will have few expenses beyond peer review and that costs will be far lower than for print journals. Electronic distribution also allows far wider copying (and hence distribution of ideas and potential citations). Search and information-retrieval technologies will allow researchers to find more quickly the information and relevant articles without having to flip through a thousand paper-based volumes. Even shelf and storage space, always a premium in overcrowded libraries, will be reduced by storing an entire shelf of books as a single CD-ROM.

At the same time, there are also legitimate concerns about the acceptability of electronic publication from the point of view of both producers and consumers/readers. As a perhaps silly example, electronic journal articles (and books) lack the so-called “thud factor” that impresses tenure review committees when one’s collection of offprints hits the conference table. Researchers and librarians alike are concerned about the long-term viability of electronic storage. Today I can easily read a book scribed in the 13th century, but I can’t listen to my old 8-track tapes from the 1970s, simply because the technological devices to play those tapes are no longer available. Finally, there is concern about the security of electronic publishing itself. If the US Department of Justice can find its web pages scribbled on and altered (Crume, 2000), why should one believe that any given (electronic) journal article is intact?

(Lynch, 2000; Lynch, 2002) describes (quoting Gustavus Simmons) the public perception of the digital realm as one of “pervasive deceit.” Digital information on the web is automatically suspect, a belief perhaps grounded in media coverage of security incidents and the self-aggrandizing writings of security consultants, but nonetheless real and a significant obstacle to the acceptance of electronic publication as a valid path for scholastic discourse. In brief, there is no accepted method for establishing the validity and integrity of any given article that does not appear in a mass-produced, paper-based format. Existing, standard methods for establishing authenticity in a corporate environment do not appear to translate well to the haphazard and freewheeling world of scholarship.

This paper proposes that existing methods *are* suited to handle these issues by an examination of the paths of trust necessary to confirm authenticity. In particular, although academia is not structured in a global hierarchy of the sort analyzed in corporate security theory, the alternative “web-of-trust” model fits very smoothly into the existing standards and general practice of working scholars. By deploying existing cryptographic tools and technologies, the authenticity problem can be addressed with by a security policy with minimal procedural change on the part of either authors, journals, or editors.

2 Threat Assessment

The first step in establishing a security program or policy, both in theory (Russell and Gangemi, Sr., 1991; Pfeeger, 1997; Rubin, 2001) and in practice (LLP, 2001), is to perform a threat assessment. This is

simply a codification, typically from the victim's point of view, of what can be reasonably anticipated to go wrong, in terms both of likelihood and of severity. Security measures can then be devised in keeping with the long-standing principle that "[things] must be protected to a degree consistent with their value." (Pfleeger, 1997)[p. 9]. A cheap padlock is perhaps inadequate to secure the thousands of dollars stored in a bank's vault, but by the same token, a multi-million dollar security system would be wasted defending a departmental coffee can that has never in history held more than \$35.00. There is also a traditional breakdown of the types of vulnerabilities that need be considered : threats to confidentiality, threats to integrity, and threats to access. Again, these are described from the point of view of the victim, the person or entity suffering the security loss, and focus on the effects to the victim, not necessarily on the method of operation. If the vulnerability is the loss of confidentiality of the secret formula to a crucial product, then it matters little whether the secret was leaked via telephone, mail, or fax machine. An ideal security plan would address this issue from a broad enough perspective to cover any kind of information leakage.

Key to an appropriate threat assessment is that these threats are placed in a specific context, both in terms of likelihood and of the victim's expectations. I may worry about the jar containing the departmental coffee fund being stolen (loss of access), but worrying about the possibility of people putting forged currency in the jar (loss of integrity) is probably wasted effort. Those who are paid to be professionally paranoid can continue to dream up and protect against ever more unlikely threats. For example, (Juola and Zimmerman, 1996; Juola, 1996) describes an arcane threat to a secure phone product that assumes an attacker with substantial computing power is interested enough in our telephone conversations to compromise telephone routing at the phone company's switch, along with a countermeasure to prevent his success. Most of the time, though, one simply picks up the phone and dials, without wondering whether the switchboard is deliberately misrouting calls to the Bad Guys. (Lynch, 2000) makes a similar argument about how information on the Internet is hampered by fears that are not characteristic of their physical counterparts — "For example, although forgeries are always a concern in the art world, one seldom hears concerns about (apparently) mass-produced physical goods—books, journal issues, audio CDs—being undetected and undetectable fakes."

What, then, are the (reasonable) expectations in a journal publishing environment? To a first approximation, publishing an article in a journal is a transaction between an author and a "journal," where both can be regarded as monolithic entities without interesting internal structure. Pseudonyms or even hidden group authorship are permissible (albeit unusual) gambits on the part of the authors (Columbia, 2001; Eves, 1992), while much of the editorial structure is deliberately hidden in the name of objectivity and anonymous review. From the author's point of view, the important, if not primary, functions of a journal are to provide an imprimatur of the scholastic quality of the work, and to provide a much wider distribution (in both space and time) than the author would otherwise achieve for her ideas. Implicit in this is the necessity for the journal to validate the connection between the author and her work.

From the author's perspective, then, issues of "confidentiality" are almost counterproductive; only by publishing as widely as possible can she achieve the promotion, tenure, research support, and so forth that are hallmarks of academic achievement. Her ideal audience would be every literate human worldwide from publication date to the end of civilization. Issues of "access" are equally addressed by wide dissemination — since each (legitimate) copy of her paper is equally a validated distribution of her ideas, she has little reason for concern about whether the reader has an offprint, a paper copy of the journal, a copy on CD, a web page, or a multigeneration photocopy of any of the above. Her primary issues are temporal : it is in her interest that her ideas be published as quickly as possible, and that archival copies be available for as long a time as possible, that researchers have access over the greatest possible duration to her work. Her primary concern is over the integrity of her work — that copies distributed as her work are identical in substance to the final version she submitted to the journal, and that the attribution of the ideas to her remains intact and correct.

To a first approximation, then, her primary vulnerabilities are : first, that the content of her work is altered in the publishing and distribution process. Second, that her work is misattributed in the publishing process. Third, that her work is not accessible, or accessible in too brief a (time) window. Any discussion of the security (authenticity, integrity, and so forth) of scholastic publishing on the web must address *these* concerns; any countermeasures, whether cryptographic or not, must treat these problems.

The interests of the journal editors and publishers are similar, but not identical. In particular, decisions on scholastic merit are made in-house, via the review process, and so need not be considered in analyzing interactions with the authors and public — but once an article is published, any further distribution has similar integrity concerns. Copies must be substantially identical to the accepted version and the attribution of *publisher* must remain intact. Within the broad community of subscribers, which

in the case of academic libraries can be essentially the entire world, access should be continuously and widely available over as great a time as economically useful. The primary difference is in the arena of confidentiality. Where the author simply desires that her ideas be as widely distributed as possible, the publisher desires that the ideas be distributed only to subscribers (and interested parties that will then purchase the journal), and within specific approved channels that can be monitored and paid for. Pirate access is a vulnerability from the point of view of a publisher, not an author.

In summary, the interests of a journal's authors and editorial staff overlap in a few key areas, which in turn are areas where they can be reasonably expected to cooperate in electronic (or other) publishing. Both the author and editor

- desire that an article be published as rapidly and completely as possible, subject to the scholastic demands of review.
- want the article to be distributed to all practicing scholars who might be interested in either the author or the journal itself as a source of continuing ideas.
- wish to see the article's contents be correctly attributed and contextualized.

It is because of the (expected) cooperation between author and editor/publisher in this area that we can use their shared interest to develop appropriate protocols to confirm attributions.

3 Elementary Cryptography

Central to any modern identification or authentication scheme is the notion of digital signatures. Prior to about 1975, authentication of unknown parties was most easily done via a secret, such as a password, to be given upon demand to a perimeter guard. In general, anyone who knew the secret was (supposed to be) authorized. The weaknesses of this scheme are, and were, obvious. Among them are the need to distribute passwords securely, the risk that an eavesdropper might overhear a password being used and then present it himself as a false identification, and the need to retain lists of valid passwords for the use of the perimeter guards. For these reasons among others, the development of public key encryption (Merkle, 1978; Diffie and Hellman, 1976; Rivest et al., 1978) and digital signatures was almost revolutionary.

A full discussion of the techniques and mathematics of digital signatures would carry this paper too far afield, and the reader is referred to the excellent discussions in (Schneier, 1996; Stinson, 2002). At essence, public-key cryptography separates the process of encryption (converting a simple message into a secret cyphertext) and decryption (reconverting the cyphertext back into understandable language), and making the two separate processes rely on two different pieces of information, either or both of which may independently be "secret." Formally, given a message M and the two pieces of information (which, for reasons that will become apparent, are called K_e and K_d), a person who holds both K_e and M can produce an encoded message C . He cannot, however, reverse the process—if through mischance, he loses or forgets his original copy of M , he will be unable to recover M and will remain forever M -less. A person holding a copy of K_d (and a copy of C) can decrypt it to (re)produce M . Furthermore, she knows that the person who produced C had, at that time, access to (or knowledge of) K_e . (K_e and K_d , because of their role in this process, are called the encryption and decryption keys, respectively.) In theory, it may be possible to determine K_e by examination of K_d or vice versa, but the amount of work this would entail typically requires running a network of computers the size of the earth for several universe-lifetimes of time. (Schneier, 1996)

This framework has two main effects. First, it immensely simplifies the process of communicating securely. By publishing the encryption key, a person (following long-standing tradition, named Alice) allows anyone to communicate securely with her. Her correspondent, Bob, obtains her "public" key K_e from public records, where Alice published it. He then either encodes his message directly, using Alice's encryption key, or he encodes a secondary message stating "Dear Alice, This is Bob. I want to talk to you, using a secret password X." Either way, the message sent, in transit, appears to be gibberish, and no eavesdropper can determine its content. Alice, however, is used to receiving apparent line noise, and decrypts the received message with K_d (which she has wisely kept secret) and recovers both Bob's message and his intentions.

Second, it allows Alice to preemptively confirm the authenticity of a message. Instead of publishing K_e , she publishes K_d . (One advantage of RSA encryption (Rivest et al., 1978) is that it allows either key to be used for both encryption and decryption. Hence the distinction between K_d and K_e blurs, and we can speak simply about public and private keys. Alice publishes one key and then uses it as either a

decryption or encryption key, as the needs of the moment demand.) She then produces a message, and simultaneously publishes both the message and its encryption C . Bob, concerned about the authenticity of the message, decrypts C and confirms both that M is valid and that the author of the message had access to K_e , which Alice (again) has wisely kept secret. This encryption C is essentially a magic cookie, a “digital signature” that validates both the contents and the authorship of M . Any changes by a nefarious Bad Guy to M in the process of transmission would result in a new message that did not match the decryption of C , while to recover Alice’s secret half of the key would require computational resources beyond the bounds of possibility.

The exact interpretation of a digital signature can vary, of course, with the context, just as physical ones can. Alice can sign documents she didn’t herself prepare, in the role of a witness to their preparation, as an independent attestation to their contents, or simply as an acknowledgement of receipt. Such an independent signature is often called a “certificate,” and is one of the main methods for distributing and proving knowledge in a digital environment. In particular, then, we can see one method of assuring authenticity in a digital publication environment. A known and trusted intermediary can issue a certificate for a given article, to be published in conjunction with that article, attesting that the article was published by such-and-such a journal, on such-and-such a date, written by so-and-so. Any recipient of such an article can validate the signature for herself; any article distributed without such a certificate is immediately suspect. The question of authenticity then reduces to the question of trust—trust in the legitimacy of the attestor’s keys, trust in the probity of the attestor, and trust in the attestor’s ability to assess authorship, publisher, and so forth.

4 Problems : Who keeps the keys?

There are a number of problems in this oversimplified scenario, some technical and some social. First, the process of encryption and decryption is numerically tedious. Second, keys, certificates, and encrypted messages all have a tendency to look like gibberish, and therefore not to be memorable or understandable by humans without computer assistance. Figure 1 shows the author’s key (from the widely available PGP(Zimmermann, 1995) program) as an example. Programs to automate this process are available, but scholars must be encouraged, if not expected, to use them—which in turn drives a need for easily usable and understandable programs that is not always fulfilled. Thirdly, the processes involved are time-consuming, even for computers. Again, this is a technical limitation, to which technical solutions exist. Instead of signing the document itself, Alice can sign a “fingerprint” or “digest” of the document, a piece of digital data that could not have been created without the original document. Finally, current implementations of digital signature technology do not allow any changes in the document at all, including relatively minor changes such as a transfer from Word to PDF, a change in font from 10pt to 12pt, and so forth. (Lynch, 1999; Lynch, 2000) has proposed that instead of signing documents, one signs “canonicalized” documents, documents stripped of their inessential features. This is clearly just another type of fingerprint, but has not yet been defined formally, in part because the notion of essentiality is so tricky.

A more serious social issue lies in the distribution and trust of keys. From Bob’s perspective, he has a person with which he wishes to communicate (Alice), and a piece of data *purporting* to be from that person (“Alice’s” key), but no actual reason to believe the data comes from Alice. This problem is not confined to the digital; an impersonator could easily have a set of business cards printed up identifying himself as Alice, the president of Brown University, an FBI agent, or the archangel Michael. Absent additional information, should Bob trust (and use) “Alice’s” key?

The standard solution to this problem (Pfleeger, 1997; Rubin, 2001) is, again, to involve the use of trusted intermediaries. If Alice and Bob share a mutual friend (even if they don’t know each other directly), the friend’s signature on an appropriate certificate can provide Bob with confidence. If Alice and Bob share an organization, a person elsewhere in the organization may be charged with confirming the authenticity of keys. With the increasing commercialization and corporatization of the Internet, there is a perceived market niche for sources of trust, a niche that companies such as Verisign (www.verisign.com) and RSA Security (www.rsasecurity.com) are rushing to fill. Verisign, in particular, offers various products such as “Authentic Document IDs,” “SSL Certificates,” and “Digital IDs” that, for a small fee, carry Verisign’s attestation to the provenance and authenticity of documents or keys. (Rubin, 2001) describes in detail the process of obtaining a Digital ID from Verisign. In practice, a Digital ID is a statement from Verisign that they have exercised a certain amount of diligence (the more diligence, the more money they charge) and confirmed that the the person named on the key is, in fact, the owner of

Figure 1: A sample PGP public key (belonging to the author)

```
-----BEGIN PGP PUBLIC KEY BLOCK-----
Version: PGP 6.5.1

mN+THIS+KEY+HAS+BEEN+DISTORTED+TO+FACILITATE+ANONYMOUS+REVIEW+zn
yVch22KKZQj6n+A2sIn+qLdKiK1LN0d0Bh7wIw1Jr1Yb/g6zMyw6TpWPPROPzkis
7U2eofSKZ4L19RSVw8+QeJfVHeMx89+QdTzUNXTAAthkJZporyC+v3X+p5ZhaAUR
tCpQYXRyaWNRlEp1b2xhIDxwYXRyaWNRlmp1b2xhQHBzeS5veC5hYy51az6JAJUD
BRAzw31YTBVe6iJfDROBAQmCBADJI9G8BuEvd1Iy1YbQarbC4MXssEQYr3/hE54T
x0WYvssqkd0o0sYDY68uuANGtQd9wKHn5x4cPKSPLA5ExUzWW1aeZICtos2AmIFk
AV2J0w06IorsBT8ymPFaCu73SrBfQJBj09t0RrZrK81Rgc4Gx2DxxEZA0Gmlfc1/
R0s9SYka1QMFEDGR+28gvr91/qeWYQEB7Wkd/jdbRjS9Pz+cfvpx5gytkjGiKWD1
1jRgKL3n9egdI2mL9q449/yCSby5JDRmQFJ3YegBptN6UkripcGwAA07vWBD3jFK
1JCi1JdsoU+gf6sPu9Yyo1Lj2X61ztnY2im/AM4UBzOYwT3xM7gOWeMMAjzRbt8M
A/qe4XtrDU5emAseiQcVAwUQM0Wv1mV5hLjHqWbdAQH0nQP9FUHDLtaFfDdU8eDE
yPIJPxTewWB+TxxaRdEK0AIA+DWSjd/eauy7zD2+o27LiwCnmpoOro0GVN1+Yo1
zEufHZuxewE89D0VYfe+vHZetk8jDVTAW2HN1qq7p04wxyMnH21PvAjEliROuhW
yJvMu05wMtU0yFzrx0ejixDli2+ZAIODLoruGAAAQQA1PKhUhYP3vsmqry9jOUd
yuQjXzTWTUNh3kHZFNACDs3lvadjSleWUqjUKncVnL72c3+Q+MQm3eKfcZ/1fB
BzhXIbx5jqznX303R68hjNxpawB81/xLyZwk2BG5mSIMqJN7PVSQ9Ju+rZrCd1b9
afKdup7hubJ8IU9ioLxwNK0ABRGOK1BhdHJpY2sgSnVvbGEgPHBhdHJpY2suanVv
bGFAcHN5Lm94LmFjLnVrPrQ1UGF0cmljayBKdW9sYSA8anVvbGFAY3MuY29sb3Jh
ZG8uZWR1PokA1QMFEC6xU8IImuDA8jIDyQEB6hMD/Rtj40yc+Z3jgMAdens0+Sfs
1jTVer/VdATLsDNA7SbkbvGgZKiWJO+1e1xk1FqiTQ4/Mmmw6Rn7g0pxwzKf9uSL
6H+d1A0DwwQAQWJpoKwTrzdd6DeA0hh77QFDYJpaSkimkVxKyX+o8ARBEN3vBsg0
no7Br5RZnIvJpx4nPRZu
=07n6
-----END PGP PUBLIC KEY BLOCK-----
```

the key. Verisign also provides \$1,000 (US) of warranty protection against loss or misuse of a Digital ID, and thus Bob has a certain amount of legal protection from the (financial) risk of using “Alice’s” key. Verisign’s obvious hope and business plan is to become sufficiently well-established that everyone will be able and willing to validate their signatures and to accept them at face value.

In addition to consulting directly with trusted intermediaries, one can save money by establishing chains of trusted signatures that may or may not ultimately be grounded in a Verisign-like service, but are primarily internal and usually follow the table of organization within a company or university. The president (or chief information security officer) obtains a key on behalf, not of himself, but of the university, and then certifies individual keys with the university’s key. This process can be carried out at any number of levels. For example, the president certifies the key of Alice’s dean, who certifies (the key of) her department chair, who in turn certifies Alice. Bob, uncertain of the validity of Alice’s key, confirms that it was validated by a key of someone (her department chair) with both reason and authority to validate hers. If he insists, he can confirm that the chair is validated by the dean, the dean by the president, and the president by Verisign.

The advantages of this sort of hierarchy are several. First, Verisign need only be paid for one key, instead of for several thousand student and faculty keys. Second, Alice’s key, plus the associated signature chain, validates not only her, but also her place in the organizational structure, and thus her authority within the larger university. Third, by focusing on the chain-of-command as authorization to sign keys, the structure is to a certain extent made robust against changes in personnel. Alice’s key is signed (ex officio) by her department chair, and the validity of that signature does not necessarily depend upon who holds that position. If the university feels confident of the probity of the outgoing and incoming chairs, it is not necessary to revoke the department chair’s key and to reissue new certificates. This hierarchical structure for certificates is, in Rubin’s words, “an ideal structure for a corporate environment, where there is a natural hierarchy of authority.” (Rubin, 2001)[p. 129] Ultimately, any two people within the university can trace their certification to a single (shared) trusted intermediary. Unfortunately, two people in different universities may not be able to trace their chains to such a shared trust; two universities may use two different corporate certificates that lack a common root. Absent a global root, there is no way to reliably confirm Alice’s connection to her key, and by extension to the documents signed by her.

Figure 2: Signatures on key presented in figure 1

```

RSA 1024      0xFE79661 1995/12/08 *** DEFAULT SIGNING KEY ***
                                     DELETED FOR ANONYMITY (someone@somewhere)
sig          0xFE79661      DELETED FOR ANONYMITY (someone@somewhere)
                                     DELETED FOR ANONYMITY (someone@elsewhere)
sig          0x225F0D1D     Giles Shilson <giles.shilson@balliol.ox.ac.uk>
sig          0xC7A966DD     Philip R. Zimmermann <prz@acm.org>

```

Therefore, we conclude that it is not necessarily practical for the journal’s editors to individually contract with Verisign (or its equivalent) to certify the provenance of every published article. The journal can, however, contract with Verisign to certify the journal’s editorial staff as a whole, and the editorial staff can digitally certify Alice’s role as author. Alternatively, if the journal can obtain the confirmation of a sufficiently well-regarded or highly placed individual, they can avoid using the commercial service and save on fees. Similarly, Alice can certify (as an author) that the journal’s editorial staff was the one who selected and published her article. As shall be shown later, this ability of Alice’s can place her in exactly the position of a “sufficiently well-regarded or highly placed individual.”

Another problem for scholastic publishing lies in the relative mobility of the agents, compared with the inflexibility of organizations. From Alice’s point of view, she wishes to be identified with all the papers she has written, despite having held positions at four different companies and universities in the past fifteen years. Similarly, journals can move from one publishing company to another and editors can be replaced. However, the joint mapping among Alice, her article, and her journal should remain validatable and verifiable even after all participants have changed organizations. Tying the validation of Alice’s key to her (former) place of employment is not a practical solution. Her employer is, bluntly, not a major participant in the process of journal publishing, and should not be a major participant in the process of authentication.

5 Web of Trust

There are, however, alternatives to a centralized, hierarchical distribution of keys and key certification. Pretty Good Privacy (PGP) (Zimmermann, 1995) provides a different model that its author calls a “web of trust.” The same notion of chained certificates applies in this model as well, but the certificate chains are not necessarily tied to any particular organizational or social structure. Instead, certificates are interpreted as being tied to personal knowledge of the people and keys involved in a transaction. The author’s key above has, for example, a certain set of associated signatures (see figure 2), including a signature from Philip R. Zimmermann, the author of PGP. For the readers of this article : Zimmermann is attesting, of his own personal knowledge, that the key presented in figure 1 is, in fact, my key. If you trust Zimmermann’s judgement, then you can believe that the key does belong to me and you can use it appropriately in communications. If you’ve never heard of Zimmermann, or of any of the other signatories, then you have no reason to trust the key.

The chief weakness with the web-of-trust is that it does not extend well beyond individual links. Zimmermann’s attestation is to my identity, not to my probity or good judgement, and so your trust in Zimmermann does not necessarily extend to a trust in my ability to introduce people to you. This problem exists as well in the hierarchical setting; just because Verisign certified that a given company exists does not necessarily imply that the president is honest in whom he certifies as his employees — but within the corporate environment, employees can to a certain extent be trusted to do their duties, at risk of discipline or lawsuits. This expectation does not extend well into the rest of the world.

From a mathematical/topological viewpoint, the web-of-trust only works within a social group with relatively dense social connections, and a high degree of relative regard, such that for any two people, the odds of there being at least one person known (and trusted) by both is high. If there are too many intervening introducers (friends of friends of friends of . . .), then there is a risk that one of the introducers is (unknown to the recipient) not trustworthy.

However, the structure and sociology of journal publishing itself makes the web-of-trust a practical model for confirming authenticity in scholarly publishing. Specifically, as will be shown in the next section, the assumptions of the web of trust almost exactly mirror the existing sociology of journal

selection, editing, and publishing. Thus, one can use the existing networks of journals as a network (web) of trust to confirm the validity, integrity, and authenticity of authors, journals, and articles themselves.

6 The Publishing Web of Trust

Let us consider first the case of an author (still Alice), who wishes to publish her newly-written article. To what journal will she send it? As discussed above, her primary concerns are to reach the appropriate audience in a timely fashion and with a maximum of prestige (the top-flight journal factor). If she is an experienced author, she is likely to send it to journals where she has published before, journals that have previously (and recently) published high-quality articles on her topic or a closely related topic. If she is an inexperienced author, she is likely to select journals recommended to her by her advisor or mentor. Again, these recommendations are likely to be based on articles previously published in the journal, possibly by the mentor/advisor. In other words, the ultimate validator of both the quality and the nature of a journal is the articles published. Implicitly, then, *the author of an article validates the journal in which it is published.*

Similarly, *the [editor of a] journal validates the author whose article is selected for publication.*

In the case of a new author (or an author entering a new field), where Alice does not know what the appropriate journal is, she will pick a journal based upon the articles and authors that have previously published in it. In other words, in selecting a journal, she will pick journals validated by particular persons (authors) whom she respects or has reason to trust (even if only from reputation and not from personal knowledge). These authors have validated the existence and authenticity of the journals in which they have published, and thus Alice has reason to trust these journals.

Conceptually, this is very closely related to the notion of Impact Factor (Garfield, 1979; Garfield, 2001; Buchtel, 2001), the measure used by librarians (and others) to determine the relative value and influence of a given scientific paper (or more typically journal). “A journal’s [impact factor] is based on two elements : a) the numerator, which is the number of citations in the current year to any items published in the previous two years, and b) the denominator, which is the number of substantive articles (source items) published in the same 2 years.” (Garfield, 2001) Journals are validated by being used in research written up in comparable journals, and an author deciding where to publish her research will typically choose to publish in the relevant journal with the highest impact factor. Of course, questions of relevance are hard to pin down to universal numerical factors, but the web of trust answers this question clearly on an individual basis. A journal is relevant to exactly the extent that it gathers the work of trusted authors within the area of scholarship.

Of course, not all journals are well-established, and a new editor will sometimes take a journal in a new direction. The direction a journal takes is a decision of the editorial staff, but typically in keeping with the interests and previous publications of the editor and reviewers. Therefore, *the editor(s) of a journal directly validate that journal. The articles published by the editor(s) indirectly validate the journal.* Of course, this chain of validation extends in the other way; by being appointed editor to a previously validated journal, the new editor and his other work are validated as well.

From a position of analysis of trust, then, we see that this web-like structure pervades the entire process of scholastic publishing. Only a completely isolated researcher, completely ignorant of the entire corpus of published work and of all other participants in the field, can fail to be aware of some validated journals. Similarly, the only journals with no possible validations are those that share no authors or editors with any other journals. In either case it is difficult to consider these to be part of the community of scholarly discourse.

In order to address issues of authenticity in the digital environment, it is necessary only to formalize and implement this web of trust in a cryptographic context. In particular, I assume that each participant (author, editor, and journal), has a unique public/private key pair, which has not necessarily been validated by any external source. As part of the journal publication process, when Alice receives notice that her work has been accepted for publication, she is encouraged (or required) to sign a certificate to the effect that this particular journal (including the journal’s public key) has published this particular paper (including a fingerprint) with this particular author (herself, including her public key) on this particular date. Similarly, the journal (more accurately, Bob, the editor of the journal, acting for the journal and using the journal’s key, not his personal key) will sign an identical certificate. Both Alice and Bob have every incentive to sign, and no reason not to. The article is published with both Alice’s certificate and the journal’s (Bob’s) attached.

Consider now a hypothetical reader of the article, who has no knowledge of either Alice or Bob’s

journal. However, by inspection of Alice's c.v., our reader can determine whether Alice has ever published in a recognized journal. If she has, our reader can conjure up Alice's previous paper and validate the certificate attached to the prior work. This will contain (and validate) Alice's public key, and thus Alice's authorship of the article in question will be confirmed, as will the integrity of the article.

A similar process can examine the list of authors ever published in Bob's journal—any recognized authors will validate Bob's journal and key, and hence Alice's article. Thus our poor reader will be left in uncertainty only if she knows no other authors in that area, and no other journals.

A set of related certificates can tie the journal's key to Bob's individual key, by having both the journal and Bob sign certificates to that effect. Finally, there is nothing to prevent individual authors or journals from validating each other directly, assuming you trust the probity of the individuals. It would, in fact, be a likely event that all journals published by a single company such as Elsevier, would cross-validate each other as a matter of business practice, and that low-profile journals would seek to be directly validated by high-profile ones in their field. There is, however, a subtle difference between trusting Alice's statement that a journal in which she published an article exists, and a statement that a colleague of hers exists, because her incentives, and therefore likelihood, of lying are slightly different.

Finally, we consider the case of authorial mobility. It should be clear that nothing in this structure ties Alice to a particular position or institution, or even requires that she have one (a benefit for independent scholars!) She is validated purely on the basis of her published research and the forums in which she publishes. Editorial mobility may be more problematic. If the journal itself moves, for example from Elsevier to Springer, or alternatively from the University of British Columbia to Penn State, but without adjusting the editorial board, Bob can continue to use the same journal public key without fear. If Bob decides to quit editing the journal, then the journal key may be compromised when he leaves. The new editor may be forced to create a new key for the journal, and to announce that henceforth everyone should use the new key. However, he can also validate (retrospectively) the old key for historical/archive purposes, with an appropriate certificate to that effect with the appropriate dates of validity.

There are several other advantages implicit in this authentication scheme. First, it can easily be applied incrementally, on a journal by journal or even author by author basis. Second, by encouraging the widespread use of cryptographic software, it can also have the effect of reducing the amount of paperwork necessary for the smooth running of a journal. Upon consultation with the editors and publishers, for example, copyright assignment can be handled with appropriate digital documents, via digital signatures, reducing the necessity for storing dead trees in a filing cabinet somewhere. Because every participant has their own keys, presumably including most of the reviewers, peer-review can be conducted almost exclusively on-line *securely*, preventing the pervasively deceitful Bad Guy from somehow stealing ideas or corrupting manuscripts in transit.

7 Conclusions

No cryptographic solution can solve all the problems of dishonesty and deceit, in any environment. Journal editors will still be faced with many of the same problems they face today—for example, plagiarism and misappropriation of text will not go away simply because the author produces a digital signature. A journal editor could still be bribed to accept a paper of sub-standard quality (editorial decisions are not within the realm of cryptographic control) or to reject an otherwise brilliant paper for arbitrary and capricious reasons. As long as paper journals exist, a dishonest academic could use a desk-top publishing package to produce a forged (paper) journal article and rely on the inertia of the recipient to not check its validity; the same inertia would keep the recipient from checking the validity of an electronic article. And, from a purely financial standpoint, a publisher could still overcharge shamelessly for electronic rights to a journal.

These, however, are mostly the ravings of a professionally paranoid security theorist. For practical purposes, a properly implemented authentication and authenticity scheme should introduce no new problems that did not exist before, address the existing concerns about moving scholastic discourse to a new medium, and gain the trust and acceptance of the user community. The tangled web of social connections involved in scholastic discourse and publishing fits the model of the “web of trust” key distribution scheme almost perfectly. By applying existing standards of trust and reputation, but codifying them with the addition of certificates that every party involved has positive incentive to produce and distribute, it is possible to achieve a high degree of trust in the authenticity of distributed documents. The cryptographic primitives involved are well-established and believed to be secure by most practitioners. Finally, by modeling directly the needs, wants, and existing practices of the users, this scheme should

be capable of ready acceptance and implementation within the scholastic community. Ultimately, that is the most crucial aspect to any security policy. Security is not a one-size-fits-all product, but is best used when tailored to fit a given user, community, practice, and/or need.

References

- Buchtel, H. A. (2001). Editorial : Introduction to forum on impact factors. *Cortex*, 37:455–6.
- Columbia (2001). Bourbaki, nicolas. In *The Columbia Encyclopedia*. Columbia University Press, New York, 6th edition. www.bartleby.com/65/.
- Crume, J. (2000). *Inside Internet Security : What Hackers Don't Want You to Know*. Addison-Wesley, Harlow, England.
- Diffie, W. and Hellman, M. E. (1976). New directions in cryptography. *IEEE Transactions in Information Theory*, IT-22(6):644–54.
- Eves, H. (1953/1992). *An Introduction to the History of Mathematics*. Saunders College Publishing, Fort Worth, TX, 6th edition.
- Garfield, E. (1979). *Citation Indexing : Its Theory and Applications [in Science Technology, and the Humanities*. John Wiley & Sons, Inc., New York.
- Garfield, E. (2001). Interview with Eugene Garfield, chairman emeritus of the Institute for Scientific Information (ISI). *Cortex*, 37:457.
- Juola, P. (1996). Isolated-word confusion metrics and the PGPfone alphabet. In *Proceedings of the New Methods in Language Processing 2 (NeMLaP 2)*, Ankara, Turkey.
- Juola, P. and Zimmerman, P. (1996). Whole-word phonetic distances and the PGPfone alphabet. In *Proceedings of the International Conference of Spoken Language Processing*, Philadelphia, PA.
- LLP, P. (2001). *Risk Management Forecast : 2001*.
- Lynch, C. (2002). Authenticity, authority and integrity in a digital environment. In *Inter/Disciplinary Models, Disciplinary Boundaries: Humanities Computing and Emerging Mind Technologies (COCH/COSH 2002 Meeting)*, Toronto, ON CA.
- Lynch, C. A. (1999). Canonicalization : A fundamental tool to facilitate preservation and management of digital information. *D-lib Magazine*, 5(9).
- Lynch, C. A. (2000). Authenticity and integrity in the digital environment : An exploratory analysis of the central role of trust. In *Authenticity in a Digital Environment*, pages 32–50. Council on Library and Informatoin Resources, Washington, DC.
- Merkle, R. C. (1978). Secure communication over insecure channels. *Communications of the ACM*, 21(4):294–9.
- Pfleeger, C. P. (1997). *Security in Computing*. Prentice-Hall PTR, Upper Saddle River, NJ.
- Rivest, R. L., Shamier, A., and Adelson, L. M. (1978). A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM*, 21(2):120–6.
- Rubin, A. D. (2001). *White-Hat Security Arsenal : Tackling the Threats*. Addison-Wesley, Boston, MA.
- Russell, D. and Gangemi, Sr., G. T. (1991). *Computer Security Basics*. O'Reilly, Cambridge, MA.
- Schneier, B. (1996). *Applied Cryptography : Protocols, Algorithms, and Source Code in C*. John Wiley & Sons, Inc., New York, 2nd edition.
- Stinson, D. R. (2002). *Cryptography : Theory and Practice*. Chapman & Hall/CRC, Boca Raton, FL, 2nd edition.
- Suber, P. (2002). Where does the free online scholarship movement stand today? *Cortex*, 38:261–4.
- Zimmermann, P. R. (1995). *The Official PGP User's Guide*. MIT Press, Boston, MA.