CHALLENGE-RESPONSE PARADIGM IN ELECTRONIC MAIL

Pawel Gburzynski
University of Alberta
Department of Computing Science
Edmonton, Alberta, CANADA T6G 2E8
E-mail: pawelg@sfm.cs.ualberta.ca

KEYWORDS

Electronic mail, privacy, abuse, spam.

ABSTRACT

We point out the weaknesses of many popular approaches to eliminating E-Mail abuse (spam) and argue in favor of the challenge-response paradigm that, in our opinion, is the only viable way to address the problem at its present stage. We also present an E-Mail server (available under GPL) and a publicly accessible free demo service (see http://sfm.cs.ualberta.ca). Our server exemplifies desirable features of an E-Mail handling system that completely eliminates spam while providing for reliable legitimate human contacts including acceptable e-commerce.

INTRODUCTION

Proposed solutions to the spam problem, range from drastic legislative measures to revolutionary changes in the infrastructure of electronic mail. In this paper, we discuss the nature of spam and explain what *exactly* makes it different from "legitimate" E-Mail. In contrast to many other voices, we claim that spam is easy to spot and eliminate in a way that will put reliability, decency, and respectability back into electronic mail. This will require a relatively simple change in the paradigm of electronic correspondence which, as we demonstrate, can be accomplished without modifying any critical elements of the existing infrastructure. To get our point across, we present a ready publicly available software package that can be immediately deployed at the MTA (Mail Transport Agent) level. Our tool has been in operation for over a year now and has evolved into a reliable, friendly, and spam-proof E-Mail handling system.

The cause of spam

Electronic mail has brought about the first truly free and egalitarian tool for probing the applicability of the famous maxim of H.L. Mencken.[1] This is because, for all practical purposes, the cost of spamming is zero. This makes spamming quite different from other tools used for mass marketing, and this is also what turns it into a plague. Even a token or imaginary return from the free marketing of a scam or a semi-legitimate product makes the venture worthwhile. Breaking even is not a problem. But the spammers do better than that. A sales log intercepted on the network (McWilliams 2003) revealed the magnitude of income from a blatantly phony merchandise sold through moderately massive spamming. During a four-week period, the number of orders for a $50 bottle of penis enlargement pills reached 6000 (most people ordered more than one bottle). Considering that the cost per bottle to the merchant was about $15 (including the materials and the spammer's fee), the profit was hardly insignificant. Thus, the plague will not go away on its own. It does bring revenue to its champions.

Techniques for fighting spam

Generic solutions aimed at eliminating spam can be put into the following three categories: 1) anti spam legislation, 2) filtering, 3) reorganizing the E-Mail transport system. The most popular proactive approach is filtering, which occurs in two basic variants: text categorization, and collaborative filtering based on shared databases of various fingerprints of spam. Owing to its simple logistics, text categorization receives most attention in many practical implementations as well as in academic research involving AI techniques (for numerous examples see http://www.spamconference.org/), especially Bayesian filters (Gee 2003; Hindle 2003; Robinson 2003; Androutsopoulos et al. 2000; Massey et al. 2003), which many people see as a remedy for spam (see http://www.paulgraham.com/spam.html). Some newer variations on the theme include case-based filtering (Cunningham et al. 2003) and learning systems (Oda and White 2003).

The collaborative approach is exemplified by Vipul's Razor (see http://razor.sourceforge.net) and SpamNet (see http://www.cloudmark.com). Some commercial systems, e.g., Brightmail (http://www.brightmail.com), deploy bogus E-Mail accounts intentionally exposed for harvesting by spambots. A message arriving at such an account is guaranteed to be a spam.

More drastic proposals call for a revision of the present paradigm of electronic mail. Among them is the idea of implementing a payment scheme for the right to send an E-Mail message (Carnor and La Macchia 1998; Fahlman 2002), which would bring E-Mail marketing at least up to par with the traditional (paper) mass mailing. The Tripoli project (Weinstein 2003), described at http://www.pfir.org/tripoli-overview, outlines a comprehensive solution based on public-key encryption and certified tokens used for granting sending rights and authenticating senders.

Owing to the fact that the most radical proposals are incompatible with the present infrastructure, the practical solutions being deployed today are less revolutionary. They can be jointly categorized as sender-confinement schemes, whereby to be considered legitimate a message must arrive from a demonstrably trusted source, with the trust established through some kind of sender authentication. The simplest commercial solutions, e.g., Spamex (http://www.spamex.com), allow the subscriber to create multiple aliases to be given away to different senders. Some other services, e.g., Mailblocks (http://www.mailblocks.com), maintain a single address of the subscriber, but associate with it a list of legitimate senders allowed to send E-Mail to that address. The first message from a new contact is bounced with a challenge intended to verify

1 No one in this world, so far as I know ... has ever lost money by underestimating the intelligence of the great masses of the plain people (Mencken 1926).

that the sender's address is legitimate.

Two non-commercial solutions of this kind, in addition to our system discussed further in this paper, are TMDA (http://www.tmda.net) and ASK (Paganini 2003) (http://www.paganini.net/ask). TMDA is implemented at the delivery point. Different senders are allocated different aliases of the recipient, similar to the idea presented in (Ioannidis 2003). An alias may expire on a given date or be restricted to a particular sender. The system also defines so-called *keyword addresses* (similar to addresses assigned by Spamex), which are not restricted a priori, but can be easily revoked when abused or no more needed. Unknown senders are challenged with a bounce and instructed to send a message to a dynamic confirmation address.

ASK guards the single address of the recipient with a whitelist and a blacklist. The sender's address is added to the whitelist if the sender responds to the bounced message (assuming that spambots do not reply to bounces). Owing to its simplicity, ASK can be implemented as a *procmail* script and causes little hassle to the subscriber.
A complex system functionally similar to TMDA is outlined in (Tompkins and Handley 2003). Its improvement over TMDA consists in postulating cryptographic signatures to authenticate senders (which TMDA achieves in a sense by signing the confinement attributes of its dynamic addresses) and insisting that the challenge be insurmountable to spambots.

Futility of anti-spam legislation

To people familiar with the technical aspects of the Internet, it is obvious that anti-spam legislation (Butler 2003; Weiss 2003), (also see http://www.spamlaws.com), is going to help little, if at all. First, even if declared illegal in the United States (or in any particular country), spam will continue to arrive from abroad. With the present convenience of acquiring disposable Internet domains and temporary IP addresses, whose jurisdiction is at best unclear, it is impossible to enforce a law that blocks messages with a certain content from arriving to subscribers within a given country. Many of the scams presently circulating in the network are provably illegal and punishable by law (e.g., the numerous pyramid schemes or derivatives of the notorious "Nigerian" money transfer scam), and have been so for many years with little negative consequences to the perpetrators.

Second, the trend with the anti-spam legislation in the United States is not to eliminate bulk E-Mail marketing but rather to define the framework of its legitimacy (Butler 2003). This attitude may in fact increase the level of junk mail in the network by legitimizing the kind of spam that complies with the rules. Following the Senate approval of the Can-Spam Act (http://www.spamlaws.com/federal/108s877.html), we immediately see a proliferation of new service providers specializing in laundering spam to make it conform to the law.

Futility of spam filtering

Spam filtering via text categorization is, in our opinion, little more than an academic exercise. People involved in this work assume that "spam employs a distinct tone and language that can be used to identify it" (Gee 2003). We claim that this is incidental and reflects the current intermittent stage of spam evolution rather than a fundamental property of abusive E-Mail marketing. For illustration, consider the following message:

Dear Son:

We enjoyed our visit very much, and I will shortly send you the pictures that we took on our way back home. The Shmodak 500 camera that you gave us is terrific: the pictures came out unbelievably clear and sharp.

Take good care of yourself,

Mother

and suppose that you have to decide whether it is spam or not. There seems to be something fishy about this message – it mentions the (bogus) brand name of a product – so, considering that our discussion is about spam, you may be inclined to put your bets on the latter. But the decision is not easy, even for a human being. Many TV commercials are not clearly distinguishable from the shows they interrupt, and one can argue that the best among them are the subliminal ones, i.e., least aggressive and least "commercial" in content.

Even if we agree that a spam message must sound ostensibly commercial, the spammer can always resort to encoding the commercial content in an image attachment. With this approach, the spammer need not worry about making the message itself subliminal. Moreover, it is easy to randomly disturb the image without affecting the encoded message. Such simple tricks, in addition to completely circumventing all filters based on text categorization, will additionally trick the collaborative filters driven by databases of sighted spam.

These simple ideas have not yet become overwhelmingly popular among spammers, but they will when the sophisticated (e.g., Bayesian) filters are deployed on any significant scale. The spammers will easily and quickly learn to circumvent those filters because, as we have argued, fooling them is far from posing a moderately challenging problem. They will accept the increased cost of doing their business because there will be no other way to spam. Note that the "cost" we have in mind is solely the amount of time spent by a program.

The correct definition of spam

The inescapable conclusion is that spam filtering is futile. This is because the whole concept of filtering is based on the wrong definition of spam. The definition assumed by the categorization-based filters is:

Spam is a message whose textual component includes words or phrases indicative of a commercial advertisement or offer and fitting certain patterns determined by a reasonably large corpora of messages collectively categorized as unsolicited bulk E-Mail by human recipients.

whereas the definition assumed by the collaborative filters is:

Spam is a message that has been sent in (nearly) identical copies to a significantly large number of different users.

Spam need not fit any of the two definitions, and the fact that most of it does fit them at present should be viewed as incidental. Thus, the above definitions do not cover the whole of spam. Moreover, they do not apply exclusively to spam. There is nothing wrong with people being genuinely

interested in Viagra®, refilling inkjet cartridges, or stuffing envelopes, and willing to exchange E-Mail on those topics. Also, one can think of legitimate (or even important) messages being sent in identical copies to multiple recipients, e.g., alerts, memos, bona-fide newsletters. For example, in a certain hospital in Toronto, an indiscriminately deployed categorization filter created a havoc by blocking, among others, all E-Mail that included the words "penis" and "prescription."

In our opinion, the only definition that captures the essence of spam is this:

Spam is a message with no human contact at the sending end who would be interested in the fate of its individual instances.

It accounts for the critical premise that makes spamming profitable: the sender of spam is not interested in actually contacting any single recipient, unless the recipient responds to the offer. If the sender were forced to personally (manually) send the message to every single recipient on the huge list, the whole procedure would suddenly become truly costly, and spamming would cease to make sense. Consequently, to prevent spam from entering your mailbox, you have to make sure that only human beings actually interested in contacting YOU in person can ever make it through the software guarding your privacy. Note that this may also apply to a program sending you E-Mail, as long as there is a human being behind that program that actually wants to get in touch with YOU.

How to eliminate spam

Many people believe that the key to eliminating spam is to enforce some form of sender authentication or certification, e.g., to verify the authenticity and validity of the message headers (Paulson 2003). The implicit assumption is that if the spammer is forced to reveal his/her "true" identity and operate "in full daylight," then 1) few people will be willing to put up with the shame, 2) it will be easy to track down spammers and enforce the anti-spam laws, 3) no respectable agency will want to certify a spammer's identity. We believe that this line of thought is naive and shortsighted, as all ideas relying on the decency of human race. First, there will never be a shortage of people ready to sell their reputation for not so big money. Second, as we mentioned above, the spam laws are unlikely to make a (positive) difference. Third, the "respectable" certifying agencies care little about moral issues related to the activities of their customers (or even themselves, e.g., try a Google search using the keywords "VeriSign abuse"). The spam problem is not a consequence of some minor deficiencies of SMTP (like the fact the message headers can be faked), but results from the openness of the underlying paradigm of electronic mail. Spam naturally exploits those deficiencies, but it can live and proliferate without them.

We claim that the only effective way to eradicate spam is to implement the kind of validation scheme that would put the human factor back into the operation of dispatching a message. We have to conclude that the only promising avenue is in the direction of challenge-response systems along the lines of TDMA, ASK, and Mailblocks. Their role is not to authenticate the sender or verify message headers but to make sure that the sender is a person rather than a program (spambot). A "person" can be formally defined as an entity capable of passing the Turing test (Turing 1950), although, due to the notorious incompetence of programs in certain areas (Ahn et al. 2003), the actual challenge presented to the sender can be quite trivial.

SFM: an outline

Our system, accessible at *http://sfm.cs.ualberta.ca*, is dubbed SFM for *Spam-Free Mail*. It operates as an extension to a standard (E-Mail transport agent) MTA that fully conforms to SMTP (Klensin 2001; Postel 1982). The main function of SFM is easy (typically automatic) generation of limited-accessibility, alternative E-Mail addresses pointing to the subscriber's *permanent* or *fixed* address which can belong to any E-Mail domain.

In principle, the permanent address of the subscriber need not be known outside the SFM server. This, however, is irrelevant from the viewpoint of spam elimination. In contrast to some other aliasing schemes (Hall 1998), the reliability of SFM does not depend on address secrecy. The permanent address, in addition to providing the forwarding target for legitimate E-Mail addressed to the subscriber, also plays the role of the user Id identifying the subscriber to the SFM server.

There are two main reasons why an alias created by SFM is immune to spam. First, it isn't published (exposed) but presented to a single contact (which can be a group of people). Second, it is restricted to the specific contact (a narrow population of senders). In the context of spam classification, this operation has been traditionally viewed as a tricky and unreliable component of any sender restriction scheme. What if a legitimate sender uses an alternative address? What if a legitimate sender passes the alias to another sender in a bona-fide attempt to forward your request to a more interested or competent person? What if the identity of a legitimate responder cannot be known at the time your are making the contact? To account for these issues, the confinement procedure for an alias is carried out as follows.

An alias is created *open*, and it remains open for a predetermined amount of time, e.g., two weeks. During that time, it will accept messages from everybody adding their senders to its *personalization list*, i.e., the list of authorized contacts. Then, when the open time expires, the alias becomes *closed*. From this point on, it will only accept E-Mail from the registered senders.

With this approach, the new contact is given ample time to identify himself/herself to the alias and shape its *personalization*. When that time is over, the personalization of the alias solidifies. If the alias is subsequently compromised, it cannot be used for mass mailing because it will only accept messages from a select group of senders (that cannot be known by the spammer). If those senders themselves decide to abuse the alias, the subscriber can easily delete the single alias without affecting other contacts.

A special type of address created by SFM is a *master*. The primary role of a master is to be published and exposed as a publicly available point of contact with the subscriber. A message arriving at a master is never forwarded to the

subscriber but instead treated as a *query*, i.e., request for an alias of the subscriber personalized to the sender. In response to this request, SFM sets up a new alias and bounces the message with simple instructions explaining that it should be re-sent to the alias. To eliminate automatic acquisition of aliases by spambots, the alias is shown in a CAPTCHA image (Ahn et al. 2003), as shown in Fig. 1 (also see http://www.captcha.net).

To get the message through, the sender has to reply to the received bounce substituting the presented alias for the recipient address. All the sender has to do is to copy the first segment of the address (the string *vathigof* in Fig. 1), as all the remaining components are already present in the sender address of the bounce.

Figure 1: An Alias Presented as a CAPTCHA Image

Another role for a master is to serve as a template for creating aliases. One idea behind having several different masters, which provide multiple publishable identities to the subscriber, is to associate different degrees of trust with those different identities. A typical setup involves two masters: one for contacts of a permanent nature (aliases created from this master never expire), and the other for intermittent and casual contacts (e.g., with aliases expiring after one month).

Organization of SFM

The interaction of all components in a complete SFM server is shown in Fig. 2. The three boxes circumscribed by the dashed line represent the SFM-specific components.

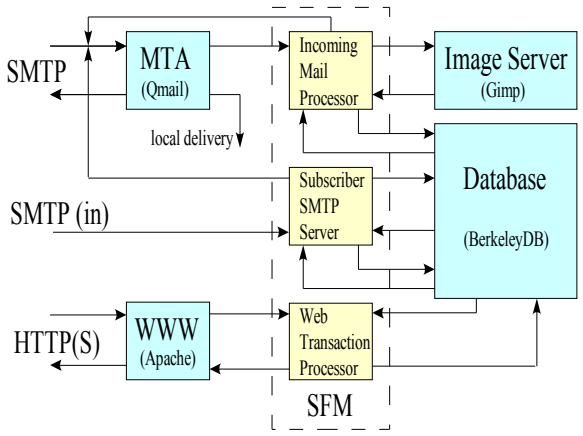

Figure 2: Organization of the SFM Server

Dynamic Web forms provide a secure authenticated interface to the subscriber's record. Through this interface, you can set up your personal attributes, create, view and edit your masters, look up your aliases and, in exceptional circumstances, create/edit them by hand.

Typically, once your SFM environment becomes set up, there is little need to access it via the Web. One exception is a quick (single-click) acquisition of a new alias, e.g., to be inserted into the Web form of an electronic merchant .

The sole task of the image server is to turn simple ASCII texts (aliased addresses of subscribers) into JPEG images. This way of presenting addresses to human contacts renders them resistant to automatic harvesting.

The alternative SFM-specific SMTP server implements only one direction of the SMTP protocol. It is available only to SFM subscribers and provides a means of dispatching E-Mail in a way that consistently and transparently presents their aliased identities to all contacts.

Reliability of contacts and e-commerce

The complete aliased address of a subscriber presented to the other party (see Fig. 1) consists of the proper alias name (the part before the first dot), followed by the master name, followed in turn by the mail domain name of the SFM server. The master name (the so-called *spice*) is not needed to reach the subscriber, but it provides a fallback measure in a situation when the alias has expired, has been removed, or is unavailable to the sender. In such a case, the message is treated as if it were addressed to the master: the sender is assigned a personal alias and informed about it via a CAPTCHA message.

This way of handling rejected messages provides for a graceful, reliable, and secure renewal of expired contacts. For example, you can create a short-lived alias and insert it into the Web form of an electronic merchant without having to worry about its possible abuse in the future. When the merchant decides to contact you after the alias has expired, SFM will make sure that the message arrives from a human being before renewing the contact.

SFM provides for an easy (one-click) alias acquisition for exactly this purpose: inserting an address into the Web form of an electronic merchant. The alias's name pops up in a new window as a complete address, which can be conveniently copied and pasted into a Web form or another document.

Mail processing

To relieve the subscriber of the responsibility for managing a potentially large and dynamic database of aliases, SFM comes equipped with an SMTP server which accepts outgoing E-Mail from its subscribers and translates their identities in accordance with the personalization of their aliases. To access this service, you must identify and authenticate yourself to SFM, which is accomplished by embedding a PIN code in your (permanent) sender address of the original outgoing message.

The authenticated sender address can be specified in one of two forms. In the simpler case, it includes the PIN code following the user name component separated from it by the + sign, e.g., *pawel+2357@dejunk.com*. Having received an SMTP request, the SFM server parses the sender address and verifies that the address identifies an existing subscriber whose PIN code stored in the database matches the PIN code included in the address. Only if this verification is successful does the server proceed with the request.

In the next step, the server collects all the destination addresses of the message and attempts to locate an existing

alias whose personalization list covers all the recipients. If such an alias is found the sender identity of the message is set to the that alias and spiced with its parent master. If no alias matching the destination list is readily available, SFM generates a new alias, adds all the recipients to its personalization list, and uses the new alias as the sender identity. Then, the message is forwarded to all its recipients.

In the second case, SFM must select a master for creating the new alias. By default the first master on the subscriber's list is used for this purpose. It is possible to explicitly indicate a different master by putting its identifier after the PIN code in the authenticated sender address, e.g., *pawel+2357+pgburzyn@dejunk.com*.

The simple PIN-based authentication scheme has been chosen as being compatible with all popular E-Mail clients, as well as the standard (unauthenticated) variant of SMTP being widely in use. It requires absolutely no effort on the subscriber's part, except for an initial configuration of the E-Mail client (MUA), and is quite secure within the framework of its application.

When the SMTP server of SFM receives an outgoing message from one of its subscribers, it looks up an existing alias whose personalization list includes the destination address. If the message is addressed to several recipients, all those recipients must be known by the alias before it is deemed suitable. If several aliases appear suitable, the one whose personalization list gives the tightest match to the list of recipients is selected. Then, the SFM server replaces the original sender address of the message with his/her aliased address (including the spice) and forwards the message to the destination.

If no alias fitting the recipient list is readily available, SFM will create one on-the-fly and initialize its personalization list with the list of recipient addresses of the message. The server avoids generating superfluous aliases; however, one should remember that each outgoing message presents a single aliased identity of its sender to all recipients. Consequently, the alias used for this identity must be personalized to all recipients. It is not uncommon that the same contact of yours will see several alternative aliases of yourself, depending on the configuration of group recipients of received messages. This poses no problem: any of those aliases can be used by your contact to reach you reliably and safely.

An incoming message intended for a subscriber may arrive on a master or on an alias. A message addressed to a master is bounced to the sender with a pertinent explanation. Two cases are possible. If an alias personalized to the sender already exists (the sender may have lost of forgotten it), the sender is reminded about the existing point of contact with the subscriber. If there is no ready alias personalized to the sender, a new alias is set up, and the sender is notified about it. In both cases, the address sent to the new contact is encoded in an image (Fig.1). The sender will have to resend the original message to the new address.

The processing of a message arriving on an existing and non-expired alias depends on whether the alias is open or closed. In the first case, the message is unconditionally forwarded to the subscriber and its sender address is added to the alias's personalization list. If the alias is closed, the sender address is verified against the personalization list. If the verification succeeds, the message is forwarded to the subscriber, otherwise, it is rejected. A rejected message whose recipient address is spiced is treated as if it were sent to the spice master. Thus, SFM generates a new alias (from the spice master) and presents it to the sender along with the bounced message.

Deployment

The primary concern of a site contemplating a transition to the new E-Mail paradigm represented by SFM will be the fate of the existing infrastructure of E-Mail addresses that have been heavily harvested and compromised. The following solutions to this problem are possible:

If SFM is installed within the E-Mail domain of an existing address infrastructure, the old addresses can be re-declared as masters. This will let them retain their official and traditional publishable status, while freeing them of spam, no matter how heavily they have been abused in the past. Owing to the fact that the MTA servicing the SFM address domain need not give up its traditional duties, this solution can be adopted gradually, as the users become convinced that they really want to switch to the new type of service.

Even if the server is installed in a different E-Mail domain, the old addresses can still be used as permanent addresses for SFM subscription. To de-spam them, the subscribers can deploy trivial and aggressive filters, e.g., blocking all incoming messages except for the ones arriving from the SFM server. This property is easily and reliably asserted via *filter cookies*, i.e., user-defined signatures inserted by SFM into message headers.

Most E-Mail clients can take advantage of the *refiltering* feature of SFM to completely de-spam an old E-Mail address while retaining its traditional and official status as a publicly known point of contact with the subscriber. This solution works with the SFM server installed in any domain, not necessarily within the domain of the old address. The necessary prerequisite on the client's side is the ability to filter messages based on keywords detected in headers and, conditionally, forward them to a special E-Mail address in a way that preserves the essential information from the original headers.

Usage

The present version of SFM has evolved significantly from its early prototype (Gburzynski and Maitan 2003). The most painful problem with the old version was the lack of SFM-specific SMTP service for processing outgoing E-Mail from SFM subscribers. To make sure that his/her identity in outgoing E-Mail was consistent with the personalization of aliases, the subscriber had to send messages to a single address within the SFM domain while passing the recipient information via special sequences in subject lines. That wasn't very friendly – the system, although useful to experts, was criticized as being cumbersome to an average user.

Another inconvenience with the old system was a large collection of arcane attributes associated with aliases and masters. The apparent need for those attributes, including complex patterns matched to the subject line and message body, was dictated by our (not always successful) efforts to account for different types of casual contacts in a way that

would make them as reliable as possible. The simple and amazingly powerful idea of keeping a new alias widely open for a limited time eliminated with a single stroke a large number of nasty problems and, at the same time, made all legitimate contacts 100% reliable. The configuration of alias attributes was reduced to a trivial number of easily understood items. As viewed by a non-expert user, the system became transparent and maintenance-free.

Summary

The general idea of a challenge-response protocol for establishing the first contact with an E-Mail recipient is sometimes criticized as cumbersome, unfriendly, unreliable, or impolite. Interestingly, we have had a chance to observe how the attitudes of people towards filter challenges evolve with time, or rather with the amount of spam that those people are forced to dig through every day. A few years ago, a message bounced with a challenge would occasionally meet with an objection from a mildly upset sender. These days, instead of objections, we are receiving words of appreciation and inquiries about our spam prevention tools. To put it in the right perspective, there is nothing wrong about a politely worded challenge after which the correspondence becomes noiseless, spam-less, and smooth.

The present author, together with many colleagues in his department, uses SFM for all E-Mail contacts, professional and business alike. We can recall no single complaint about a missing message, nor a single case of spam having sneaked through the system.

The proliferation of spam on the Internet has brought us a challenge, which we originally interpreted as a need to devise better filters in response to new spamming tricks. This is difficult and unfair. Ultimately, to carry out its duties without a mistake, a spam filter must be able to understand not only the message, but (as we argued in this paper) also the intentions of its sender. We cannot beat the creativity of spammers with mechanical filters, but we can easily reverse the problem and challenge the spambots instead. This will put the reliability, privacy, and respect back into our E-Mail contacts: the contents of our mailboxes will be shaped by people rather than programs.

REFERENCES

Ahn, L., Blum, M., Hopper, N.J., and Langford, J., "CAPTCHA: Using Hard AI Problems for Security," *Proceedings of EUROCRYPT'03,* pp. 294-311, Warsaw, Poland, 2003.

Androutsopoulos, I., Koutsias, J., Chandrinos, K.V., and Spyropoulos, C.D., "An Experimental Comparison of Naive Bayesian and Keyword-Based Anti-Spam Filtering with Personal E-Mail Messages," *Proceedings of ACM SIGIR,* pp. 160-167, Athens, Greece, 2000.

Butler, M., "Spam - the Meat of the Problem," *Computer Law & Security Report,* Elsevier, vol. 19, no. 5, pp. 388-391, 2003.

Cranor, L. and LaMacchia, B., "Spam!" *Communications of the ACM,* vol. 41, no. 8, pp. 74-83, 1998.

Cunningham, P., Nowlan, N., Delany, S.J., and Haahr, M., "A Case-Based Approach to Spam Filtering that Can Track Concept Drift," *Technical Report TCD-CS-2003-16,* Trinity College, Dublin, 2003.

Fahlman, S., "Selling Interrupt Rights: A Way to Control Unwanted E-Mail and Telephone Calls," *IBM Systems Journal,* vol. 41, no. 4, pp. 759-766, 2002.

Gburzynski, P. and Maitan, J. "A Comprehensive Approach to Eliminating Spam," *Proceedings of EUROMEDIA'03,* Plymouth, UK, April, 2003.

Gee, K.R., "Using Latent Semantic Indexing to Filter Spam," *Proceedings of the 2003 ACM Symposium on Applied Computing,* pp. 460-464, Melbourne, Florida, 2003.

Hall, R., "How to Avoid Unwanted E-Mail," *Communications of the ACM,* vol. 41, no. 3, pp. 88-95, 1998.

Hindle, R., "An Introduction to the Spambayes Project," *Linux Journal,* no. 107, 2003.

Ioannidis, J., "Fighting Spam by Encapsulating Policy in E-Mail Addresses," *Proceedings of NDSS'03,* San Diego, CA, 2003.

Klensin, J., "Simple Mail Transfer Protocol," *Request for Comments 2821,* Internet Engineering Task Force, 2001.

Massey, B. et. al, "Learning Spam: Simple Techniques for Freely-Available Software," *Proceedings of the USENIX Annual Technical Conference (FREENIX Track),* pp. 63-76, San Antonio, TX, 2003.

McWilliams, B., "Swollen Orders Show Spam's Allure," *Wired News,* Internet publication: http://www.wired.com/, August, 2003.

Mencken, H.L., "Notes on Journalism," Chicago Tribune, September 19, 1926.

Oda, T. and White, T., "Developing an Immunity to Spam," *Lecture Notes in Computer Science,* Springer-Verlag, vol. 2723, pp. 231-242, 2003.

Paganini, M., "ASK: Active Spam Killer," *Proceedings of the USENIX Annual Technical Conference (FREENIX Track),* pp. 51-62, San Antonio, TX, 2003.

Paulson, L.D., "Group Considers Drastic Measures to Stop Spam," *IEEE Computer,* vol. 36, no. 7, pp. 20-22, News Briefs, 2003.

Postel, J., "Simple Mail Transfer Protocol," *Request for Comments 821,* Internet Engineering Task Force, 1982.

Robinson, G., "A Statistical Approach to the Spam Problem," *Linux Journal,* no. 107, 2003.

Tompkins, T. and Handley, D., "Giving E-Mail Back to the Users: Using Digital Signatures to Solve the Spam Problem," *First Monday,* vol. 8, no. 9, http://www.firstmonday.dk/, 2003.

Turing, A.M., "Computing Machinery and Intelligence," *Mind,* vol. 49, 433-460, 1950.

Weinstein, L., "Spam Wars," *Communications of the ACM,* vol. 46, no. 8, p. 136, 2003.

Weiss, A., "Ending Spam's Free Ride," *netWorker,* vol. 7, no. 2, pp. 18-24, 2003.