Spam Zombies from Outer Space

John Aycock and Nathan Friess
Department of Computer Science
University of Calgary
2500 University Drive N.W.
Calgary, Alberta, Canada T2N 1N4
{aycock,friessn}@cpsc.ucalgary.ca

TR 2006-808-01, January 2006

Abstract

How can better email worms be created? How can spyware claim more victims? How can better spam be sent? It turns out that these three questions are all related. We look at the future threat posed by email worms, spyware, and spammers that use zombie machines in a new way: sophisticated data mining of their victims' saved email. The end result skirts many defenses, and could trick even experienced users. We present new defenses that can be deployed now against this threat.

1 Introduction

In both advertising and social engineering, the common goal is to convince a targeted human to perform an action. The exact nature of the action varies with the domain: email worm authors may try to have the target run an executable attachment; spyware authors may want to direct the target to a specific web site for a drive-by download; spammers may want to have the target visit their web site or that of an affiliate. We will refer to email worm authors, spyware authors, and spammers using the generic term "adversary" since they share this common goal. In each case, the goal can be accomplished by sending out bulk email – spam – to potential targets.

Adversaries have one other goal, too. They want their spam to be as effective as possible. How can they do this? A flippant answer is for adversaries to use spelling checkers and both upper *and* lower case. A full answer is more complicated. The effectiveness of an adversary's spam can be characterized using the following formula:

$$\mathit{effectiveness} = \underbrace{\mathit{sent} \times (1 - P_{\mathit{filtered}})}_{\text{pre-delivery}} \times \underbrace{P_{\mathit{read}} \times P_{\mathit{clickthrough}}}_{\text{post-delivery}}$$

This formula measures effectiveness as a function of how many targeted people eventually perform the action that the adversary wants them to perform. The formula can be broken into two parts, one part describing what happens before the spam gets delivered to the target's mailbox, one part describing what happens post-delivery. The other terms are:

sent The amount of spam email messages sent by the adversary.

$P_{filtered}$ The percentage of the spam that is intercepted *en route* by anti-spam defenses. The place where the filtering occurs is not relevant to effectiveness, and may be near the sending machine, the receiving machine, or done on a per-user basis.

P_{read} Percentage of the spam that is delivered to a potential target, and that is read by the target. By itself, this is a measure of the effectiveness of the From: and Subject: email headers used by the spam, as targeted users use the information in these fields (as displayed in their mail reader) to decide whether to read or delete the email.

$P_{clickthrough}$ The percentage of targeted users that click on what the adversary supplies in their spam, such as a URL or an attachment. This describes how convincing the spam's body is.

The bulk of effort by adversaries, when trying to increase the effectiveness of their spam, has been directed at the pre-delivery part of the formula. Spammers increase the number of messages sent, to try and reach a wider audience of potential targets, or alter the messages in bizarre ways to try and evade filters. The latter tactic is particularly counterproductive, as it often renders the message unintelligible to the target.

Some efforts have been made at improving the post-delivery percentages, too. Prepending "Re:" onto a Subject: line to suggest that the spam is really a reply to some previous message was a clear attempt to improve P_{read}. As another example, the Mydoom.bb email worm tried to improve upon $P_{clickthrough}$ by using social engineering, and claiming to be from the IT staff [12]. These efforts are not very advanced yet, however. Ironically, if post-delivery attacks were more advanced, spam volume could actually *decrease* while yielding the same overall effectiveness.

In the remainder of this paper, we present a new method adversaries could use to improve spam effectiveness. The following sections explain the improved spam technique, the current defenses it avoids, our proof-of-concept implementation, and new defenses. We then look at further possibilities for this technique, along with related work.

2 Improving Spam

Improved spam will come from zombie machines. At first blush, this isn't terribly original, but the difference is in how the zombies are used. A zombie

machine is not just a throwaway resource, a launching pad for DDoS attacks and spam; a zombie contains a wealth of data.

There are two key reasons why spam is suspicious to anti-spam filters and human targets alike. First, it often comes from an unrecognized source. Second, it doesn't *look* right. The evolution of spam zombies will change this. These new zombies will mine corpora of email they find on infected machines, using this data to automatically forge and send improved, convincing spam to others. In addition to the adversary, there are two other parties involved here: the *victim*, who owns a zombie machine, and whose saved email the adversary will be mining; the *target*, currently uninfected, that the adversary wants to click on something.

What features can be easily extracted from an email corpus? There are four categories:

1. Email addresses. The victim's email address and any other email aliases they have can be extracted, as can the email addresses of people with whom the victim corresponds. These addresses could be seen as fodder for harvesting by spammers, but we will show an alternate way to use these.

2. Configuration. This is information unrelated to individual messages, but related to the victim's email program and its configuration. For example, the User-Agent: email header added by the mail program to identify the sending program, the message encoding as text and/or HTML, and any automatically-appended signature file. The quoting style used by the victim's email program for replies and forwarded messages is a configuration issue too.

3. Vocabulary. The normal vocabulary used by the victim and the people with whom they correspond.

4. Email style. There are a number of distinctive stylistic traits, including:

   - Line length, as some people never break lines;[1]
   - Capitalization, or lack thereof;
   - Manually-added signatures, often the victim's name;
   - Abbreviations, e.g., "u" for "you";
   - Misspellings and typos;
   - Inappropriate synonyms, e.g., "there" instead of "their";
   - Replying above or below quoted text in replies.

   Email style should be correlated with the email recipient(s), because the victim may use both formal and informal email styles depending on who they are sending to.

[1] The email program and its configuration may also play a role here.

The adversary would begin by releasing malware which would infect machines and create an initial pool of zombie machines. The malware would mine any email corpora on the infected machines to determine as many of the victim-specific email features above as possible. Any of the features may change over time – a former co-worker the victim emailed may have left the company, or the victim may have switched email programs – so a prudent adversary would only use recent messages in a corpus.

Next, the adversary would supply a message to send out, with proper spelling, grammar, and capitalization. While the message could consist of multiple sentences, it would be supplied by the adversary without line breaks. It is irrelevant whether the message was built in to the malware, or whether the adversary has a mechanism to communicate a fresh message to the zombie network. The malware would transform the message using simple filters, so the message is rendered in the style the victim uses for a given target. For example, if Vic the victim normally uses an informal style to email Margaret the target, then the malware might determine the following parameters about Vic's prior email to Margaret:

Average line length	30 characters
Capitalization	none
Manually-added signature	"- vic"
Abbreviations	"u" for "you", "r" for "are"

And the message supplied by the adversary:

```
Hey, have you seen this web site yet? It has some cool vid
eos: http://evil.example.com
```

Would be transformed into:

```
hey, have u seen this web site
yet? it has some cool videos:
http://evil.example.com

- vic
```

At this point, the malware could simply spam out the message to Margaret and other targets. The transformed message still looks "spammy" by itself, however, even though it will come in the correct style from a recognized source. We think that more effective spam will not be sent like this, but in one of the following forms:

- Replies to incoming email. Here, the malware monitors legitimate incoming email addressed to the victim. When a new message arrives, its sender is targeted; before the victim has a chance to reply to the email, the malware composes and sends a reply in the victim's style:

```
Sorry, I won't be able to get to your email for a while.
However, this should keep you amused in the meantime:
http://evil.example.com.
```

This type of reply is consistent with observed human behavior in email interactions [20]. The deception can be further extended if the malware doesn't generate a reply when the victim is actively using the infected computer, because the victim may actually be replying to the email.

There are other advantages for the adversary. The malware is able to quote the incoming message, lending more credence to the spam. Also, by replying, the malware has a valid Message-Id: to use in the In-Reply-To: and References: email headers – a target using an email program that shows threads will link the malware's reply to the original mail.

- Forwarded email. The forwarding of messages, attachments, and links is a common occurrence in email, and is arguably a form of gift-giving [15]. However, forwarding by humans is done bearing in mind the recipient's interests [15]. A target's interests are not only hard for malware to automatically determine, but they are highly unlikely to coincide with the spam the adversary has to send.

  Effective spam can shy away from the problem of targeting specific interests by using humor, a common denominator. An adversary can lure targets with the promise of funny pictures, sound clips, or animations; these can be easily copied from the Internet onto an adversary's web site (which also sports a drive-by download) or attached to spam. Having such a real lure acts as a Trojan horse, and can delay suspicion by the target.

  The other problem is making the forwarded email come from an authentic-looking source. Malware can mine information from the victim's email corpora about groups of people that the victim usually CCs or sends email to; this is more precise targeting than using a victim's address book, because it reflects the victim's actual, current email practice. This information can be used for several purposes:

  - Sending supposedly-forwarded spam from the victim to an authentic social group of targets;
  - Making reference to a common acquaintance in the spam;
  - Spamming fake "e-card" announcements from the victim to a plausible group of recipients.

  As a special case, some people with a close, informal relationship simply forward URLs in the Subject: line, with a blank message body. The malware can look for these in the victim's email corpus, and exploit the trusted relationship to send URLs to the target similarly.

None of these schemes will work if the malware sends multiple copies of the spam to a target, or if the target receives one of these spams for every message they send the victim. The malware must keep a database of spam attempts, and only send one spam per target. Further suspicion can be avoided if spam is not sent to already-infected machines; a zombie can covertly hide an infection indicator in outbound email with some crude steganography [21]. For example,

the last line of all email from a zombie machine (whether a spam message or not) may have thirteen spaces appended onto it.

Malware can also avoid sending email at unusual times. As with other features, a victim's email corpus can be mined to determine those times when they usually respond to email. This also contributes to making the spam look more normal.

The path traveled by more effective spam is worth noting. A topologically-aware worm [17] is a worm that scans using information about network topology, to avoid the wasted effort in random or quasi-random scanning. More effective spam travels along the topology of *social* networks instead of computer networks, like an IM worm that only tries to infect people on "buddy lists". The "small world" hypothesis holds that any two people are separated from each other by a short chain of acquaintances [13, 19]; this idea has recently been applied to searching in computer networks (e.g., [14]). One theory is that social networks work this way due to the presence of highly-connected people [6]. Intuitively, some people are more popular than others. This would mean that more effective spam only has to fool a handful of targets, enough to turn a highly-connected target into a victim, in order to be widely dispersed.

3 Defenses Avoided

The more effective spam described in the last section avoids many defensive measures that targets may have in place. Here, we look at anti-spam defenses responsible for $P_{filtered}$; detecting any malicious content run or downloaded as a result of the victim clicking what the adversary supplies is a separate issue, and is beyond the scope of this paper. For each defense, we give a brief description of it and an explanation of how more effective spam can evade it. Specific references are given where appropriate; [24] is a good general introduction to the area.

Whitelisting. Incoming email is allowed only from those senders who are on a special list, called a "whitelist". Effective spam is sent from the victim, *as* the victim, to the target; if the victim is not whitelisted already, then no legitimate email from the victim can get through to the target.

Blacklisting. Mail is refused if the sender is on a special "blacklist". Again, if the victim is blacklisted, then legitimate mail from the victim could not be getting through to the target, and thus blacklisting is ineffective here.

Greylisting. Mail from first-time senders is initially refused, on the premise that spammers won't retry transmission of email, but legitimate senders will. Senders that do retry are temporarily whitelisted to prevent future retransmission delays [11]. More effective spam can piggyback on the victim's normal email system, which (presumably) handles retries properly.

Rate limiting. Rate limiters, or throttles, assume that spam is being sent in large quantities, and slow the rate of outbound email to suppress

spam [22]. This assumption is not valid for more effective spam. Reply-based spam will operate only as fast as legitimate incoming email to the victim; forward-based spam will not be convincing if sent in large quantities, so a smart adversary would limit the rate of forward-based spam themselves.

Tarpits. A tarpit is a special mail transport program that replies... very... slowly... to incoming spam transmissions, to slow down overall spam transmission or at least annoy spammers. There would be no way to apply a tarpit against more effective spam without harming the victim's legitimate email, and again, this defense assumes large quantities of spam transmission.

Spam traps. Spam traps are email addresses made public (e.g., posted to a web page) with the intention that they be harvested by spammers. A spam trap does not correspond to a real person, however, so any email sent to that address must be spam. More effective spam is sent to addresses from legitimate incoming email, or to legitimate addresses mined from a victim's saved email, thus avoiding spam traps.

Challenge/response systems. A first-time sender is issued a challenge, which the sender must reply to prior to their email being delivered. Essentially, the goal is to try and prove that the sender is a human being. A successful response to a challenge results in the sender being whitelisted. Our reply-based spam would not be subject to a challenge; correctly-designed challenge/response systems would automatically whitelist the victim, because the target initiated the email exchange. Forward-based spam is sent to people named in the victim's saved email, for which the victim would have already replied to the challenge.

Proof of work. Proof-of-work systems have the sender demonstrate that they have spent some nontrivial amount of computer time computing a value [4]. The idea is that this would impede spammers from sending mass email by imposing a tangible "cost" on them. Yet again, this defense assumes large volumes of spam, and is inapplicable to more effective spam.

Authentication systems. Mail is verified to be coming from a legitimate mail sender for the domain stated in the email (e.g., Sender Policy Framework [23]), *or* mail is verified to be sent by the stated domain and unaltered in transmission (e.g., DomainKeys [3]). More effective spam is sent from the victim's machine in an authorized manner, and is unaltered *en route*, so sender authentication does not help as a defense.

Checksum-based filtering. A collaborative scheme whereby a checksum, preferably one insensitive to minor text alterations, is computed for each incoming message. The checksum is sent to a central server that stores the checksums of known, reported spam; a hit from the server indicates that

the email is spam. More effective reply-based spam would not be vulnerable to detection, because the target's quoted text in the reply would throw the checksum off. Forward-based spam could be caught, if the adversary's malware only introduces minor variations to the spam. However, the low-volume, deceptive nature of more effective spam may result in either the spam (and its checksum) not being reported, or it falling short of a "bulk email" threshold used by the central checksum server.

Statistical filtering. Statistical filters, such as Bayesian filters [7], predict the likelihood of email being spam or ham based on how often the words[2] in the email appeared in corpora of previous spam and ham. Reply-based spam includes the target's quoted text to push the email in the "ham" direction. In a sense, this is a sophisticated form of "word salad", where spammers try to slip by statistical filters by including random words or phrases in their spam [9]. Forward-based spam is more vulnerable to detection, though, because its message is more exposed.

Pattern-based and heuristic filtering. Pattern-based filters detect spam by looking for a match in a predetermined set of "this is spam" patterns. Heuristic filters combine a large number of heuristic tests (which may include all of the above types of defense) to make a spam/ham judgment about email. Both cases are similar to statistical filtering: reply-based spam is likely to avoid detection, forward-based spam may or may not be caught. Because more effective spam deceives by looking like legitimate email, the risk of trying to increase an anti-spam filter's detection rate is an increased number of false positives.

Of course, there are many variations and combinations of these defenses; this should be taken as an exhausting, but not exhaustive, list. The main point is that existing defenses are not enough to combat more effective spam.

4 Proof-of-Concept Implementation

We performed some experiments mining data out of email corpora, to determine how easy it would be to mount our proposed spam attack, and to see how convincing the result would be. Although they could be forged by the adversary's malware, a number of the features mentioned previously will be automatically incorporated into outgoing mail if the malware uses the victim's email program to send its mail, including From:, User-Agent:, preferred encoding (text and/or HTML), and any automatic signature file.

For our experiments, we focused on mining capitalization, commonly used abbreviations, common misspellings, manually entered signatures, and social groups from two corpora. The first corpus we used was a small dataset which we manually generated ourselves; the second was the public corpus of Enron email. In each of the cases, except for the social groups, we only examined

[2] "Word" is used in a very loose sense here.

the victim's outgoing email folder so that we are only examining text which the victim actually wrote. The implementation was comprised of several Perl scripts, each of which mined one of the listed features from the corpus. Each script did the following for each email message:

1. Strip off the automatically-appended signature, if any. This signature was found by searching the email program's settings.

2. Search for and remove any quoted text, in case this message was a reply or forward.

3. Perform the data mining (e.g., capitalization, manual signature).

4. Save the statistically significant results for later use in generating spam.

When mining the victim's capitalization habits, the script merely split the text into individual words, counting the number of words that were completely lower case, completely upper case, or mixed case. The decision as to which form of capitalization the user used for any given message was based on a percentage of which of the three forms were encountered. Although we simply estimated that 20% of words should be capitalized before we declare that the victim was using mixed case, a proper analysis of English text could be used to fine-tune this number. Once the script had processed all of the emails, our spam generator mimicked the capitalization most used by the victim.

Mining abbreviations was quite similar to capitalization, except that our script started with a list of commonly used abbreviations such as:

u	you
r	are
ur	your
cya	see you
ic	I see
afaik	as far as I know
thru	through

Each time one of the abbreviated forms of a word was encountered, the script incremented a counter for that type of abbreviation. In the end, if we found more than five occurrences of an abbreviation, our generator also used that abbreviation, although this threshold could be fine-tuned as well.

The script for mining common misspellings sent each word to a spell checker (Aspell), and noted which words were flagged as incorrect, as well as the suggested replacement. Again, if we found more than five occurrences of a spelling mistake, and the spell checker provided a suggested replacement word, then our generator performed the reverse replacement when creating the spam email.

When mining manually entered signatures, rather than looking at each word of the text, the script examined the last non-blank line of the text. The script kept track of each unique signature, and how many times it was encountered throughout the corpus. Similar to the other transformations, if a signature was

found more than five times, it was added to the pool of signatures which were randomly added to generated spam.

The script which mined social groups was slightly different in that it examined all of the headers of an email, as well as the entire body, including quoted text. The script built a list of all of the email addresses found in the message for two purposes. The first reason was to count the number times any given pair of addresses appeared in the same message. The second was to create a map between email addresses and real names of individuals. Both of these results were also used by our spam generator.

Finally, our spam generator used all of the above results to transform any given text into a spam email. The input for the generator was an email to reply to and some text to add to the reply, which would presumably contain a URL for our target to click on. The generator then output an email message which could be sent to the target.

The results of our experimentation were very positive (for the adversary). For the Enron corpus, our scripts were able to mine almost all of the mentioned transformations successfully. In one mailbox, we found over 670 occurrences of the abbreviation "r" and over 160 occurrences of "u". Furthermore, even though the formatting of the emails was quite erratic, our scripts also successfully found the victim's manual signature "Vince", with 1470 occurrences, far above the next best guess at 216 occurrences. Mining common spelling mistakes was less successful, as the spell checker had a tendency to correct things like people's names when it shouldn't have. This result is not entirely surprising. However, the spell checker did find minor nits such as "p.m." versus "PM".

Furthermore, the generator was able to create a convincing reply to an email in our custom made corpus. This can be demonstrated through an example. Given the following data from our fictitious victim, Jane Doe:

Abbreviations	"u" for "you"
Capitalization	mixed case
Manually-added signature	"Jane D"
Social groups	`tim@bigcorp.domain`, `rick@bigcorp.domain`
Real names	`rick@bigcorp.domain` ⇒ "Rick CoWorker"

If Jane Doe had received an email from Tim Boss with the text:

```
Hi Jane,

When you get this message, I'm expecting a reply.  Please reply
to me as soon as possible.

Thanks,

T Boss

--
Tim Boss
The Big Manager
Big Corporation
```

And our intended spam message of:

```
Hi,

I just talked with [INSERT NAME HERE] and you should take a qui
ck look at http://some.bad/url
```

Our generator was able to create a reply to Tim Boss with the text below. (Note that the style of quoted text and the automatically-appended signature were added by the victim's email program.)

```
Tim Boss wrote:
>Hi Jane,
>
>When you get this message, I'm expecting a reply.  Please reply
>to me as soon as possible.
>
>Thanks,
>
>T Boss
>
>--
>Tim Boss
>The Big Manager
>Big Corporation
>
>
>

Hi,

I just talked with Rick CoWorker and u should take a quick look at
http://some.bad/url
Jane D

--
Jane Doe
An Employee
The Big Corporation
```

5 New Defenses

User education is often viewed as a key defense against malware and spam. However, this is of limited value against the more effective spam we describe, because the spam bears all the normal hallmarks that the legitimate sender would display. This implies that a strong reliance must be placed on technical defensive measures. We suggest the following:

Encrypt email corpora. Encryption will delay malware from mining email corpora, but will not indefinitely prevent it. Malware on the zombie machine can eventually steal the decryption key and operate as before.

Salt email corpora. Email corpora can be salted with false information for malware to find, such as spam trap addresses or manually-added signatures like "Help, my machine is a spam zombie!" This would provide some warning of malware infection.

Sandbox the browser. URLs followed from email should be viewed in a sandboxed browser, to limit the risk from drive-by downloads. More generally, this precaution applies to any application launched from an email reader.

Adjust anti-spam filters. Anti-spam filters can be attuned to our more effective spam in four ways:

1. Watch for known infection indicators, if the adversary's malware uses one.
2. Place extra emphasis on URLs sent via email. This is unfortunately not always a trivial check, as an adversary may use innocuous-looking jump pages in their spam that redirect to the real web site [16].
3. Canonicalize incoming email before filtering, undoing known transformations that the adversary's malware performs, before looking for known spam. The obvious disadvantage is the emphasis on known information, leaving lots of leeway for an adversary to slip new spam messages through before defenses have caught up.
4. Examine alleged reply messages both with and without any quoted text from the original message, to try and expose reply-based spam by stripping off its legitimate cloak.

For threats sent by an adversary as attachments, anti-virus software is still the best defense.

6 Other Possibilities

There are further possibilities for more effective spam. We have limited ourselves to relatively simple transformations of the adversary's original spam message. We conjecture that to scale the length of the spam (e.g., to the length of a typical 419 scam email) in a manner that would be convincing to a target would require a fairly sophisticated language model. This does not preclude an adversary from doing this, merely makes it more difficult. Furthermore, if excessive time or space were required to do this, malware could transmit necessary information from a victim's corpora to the adversary's computer for offline processing.

Zombies could also be used for targeted marketing by spammers. Currently, spammers can enter keywords describing their target market into programs like Atomic Email Hunter [1]. The program will search for those keywords on a web search engine, then mine the returned web pages for email addresses. These keyword queries could be done just as easily in a victim's email corpora, with the advantage that the social networks that the corpora describe tend to connect people with like interests. Further keyword information is available if the adversary's malware automatically locates and searches a victim's homepage [2].

Specific vulnerabilities in mail programs, either mail readers or mail transport programs, could be exploited by mining email corpora too. Malware could be instructed to look for vulnerable versions of software in saved email headers,

such as the User-Agent: and Received: lines, in order to identify targets. One defense would simply be to not save these email headers, but this defense would need to be widely-adopted to be effective, and a zombie machine could easily intercept the header information on new incoming email.

One hurdle adversaries face when evading targets' statistical spam filters is that there is no easy way of determining which words are indicative of ham and spam [8]. A victim's email corpus, however, details which words have gotten by a target's statistical filter in the past. It also documents a target's own vocabulary, which is likely a strong indicator as to which words a target considers normal, or "hammy".

7 Related Work

> 'Directed attacks can, in theory, work, but require so much work on the spammer's part that the spam-filtering community hasn't seen a true attempt of these in the wild.' [24, page 137]

The quote above refers to attacks on statistical spam filters, and is also echoed by [8] – the assumption is that an adversary will have an uphill battle, manually guessing at per-user ham words in the absence of feedback from the target. Our work suggests the contrary: that such directed attacks, using a victim's email corpus, are very likely, can be automated, and are scalable.

More effective spam can be both simpler *and* richer than predicted. The number one defense suggested by [8] is to only accept plain-text email, yet our more effective spam will work in both plain-text and HTML formats. Graham refers to one-line messages with URLs as 'the spam of the future' [7], yet we have shown that these emails can be enhanced with mined contextual information in a believable fashion, and this is only the tip of the iceberg.

The closest work to ours is the 'context-aware' phishing attacks suggested by [10]. There, a target is tricked into believing a phishing email by exploiting context information, in conjunction with a well-timed attack. While [10] gives some specific phishing-related examples, there is no scheme proposed for automating such attacks or generalizing them beyond phishing.

Some facets of our approach have appeared independently, too. Parasitic spam [18] would have zombies tag spam onto legitimate outgoing email messages; [5] mentioned zombies making use of a victim's email program settings.

8 Conclusion

More effective spam can be sent by using malware on zombie machines to mine data from email corpora. This allows spam to be sent that automatically mimics legitimate email sent by the real owners of the zombie machines, and our proof-of-concept implementation demonstrates that the result can be convincing even to seasoned users. While this more effective spam has not, to our knowledge,

been seen in the wild, there are defensive steps that can be taken now to limit its impact when this spam makes its debut.

9 Acknowledgment

The work of both authors has been supported by grants from the Natural Sciences and Engineering Research Council of Canada.

References

[1] AtomPark Software. Email extractor. http://www.massmailsoftware.com/extractweb, 2005.

[2] A. Culotta, R. Bekkerman, and A. McCallum. Extracting social networks and contact information from email and the Web. In *First Conference on Email and Anti-Spam*, 2004.

[3] M. Delany, ed. Domain-based email authentication using public-keys advertised in the DNS (DomainKeys), September 2005. Internet draft; work in progress.

[4] C. Dwork and M. Naor. Pricing via processing or combatting junk mail. In *Proceedings of CRYPTO '92*, pages 139–147, 1993.

[5] N. FitzGerald. Why 'user authentication' is a bad idea. In *Virus Bulletin Conference*, page 226, 2005. Abstract only; the comment we refer to was made during his conference presentation.

[6] M. Gladwell. *The Tipping Point: How Little Things Can Make a Big Difference*. Little, Brown and Company, 2002.

[7] P. Graham. A plan for spam. http://www.paulgraham.com/spam.html, August 2002.

[8] J. Graham-Cumming. Fooling and poisoning adaptive spam filters. Sophos white paper, November 2003.

[9] J. Graham-Cumming. How to beat an adaptive spam filter. In *MIT Spam Conference*, 2004.

[10] M. Jakobsson. Modeling and preventing phishing attacks. In *Financial Cryptography '05*, 2005. Phishing panel.

[11] J. R. Levine. Experiences with greylisting. In *Second Conference on Email and Anti-Spam*, 2005.

[12] McAfee, Inc. W32/Mydoom.bb@MM. McAfee Virus Information Library, 2005.

[13] S. Milgram. The small-world problem. *Psychology Today*, pages 61–67, 1967.

[14] Ö. Şimşek and D. Jensen. Decentralized search in networks using homophily and degree disparity. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence*, pages 304–310, 2005.

[15] M. A. Smith, J. Ubois, and B. M. Gross. Forward thinking. In *Second Conference on Email and Anti-Spam*, 2005.

[16] Spammer-X. *Inside the Spam Cartel*. Syngress, 2004.

[17] S. Staniford, V. Paxson, and N. Weaver. How to 0wn the Internet in your spare time. In *Proceedings of the 11th USENIX Security Symposium*, 2002.

[18] D. Swimmer. The spectre of parasitic spam. *Virus Bulletin*, pages S2–S4, December 2005.

[19] J. Travers and S. Milgram. An experimental study of the small world problem. *Sociometry*, 32(4):425–443, 1969.

[20] J. R. Tyler and J. C. Tang. When can I expect an email response? A study of rhythms in email usage. In *Proceedings of the Eighth European Conference on Computer Supported Cooperative Work*, pages 239–258, 2003.

[21] P. Wayner. *Disappearing Cryptography*. Morgan Kaufmann, second edition, 2002.

[22] M. M. Williamson. Design, implementation and test of an email virus throttle. In *19th Annual Computer Security Applications Conference*, 2003.

[23] M. Wong and W. Schlitt. Sender policy framework (SPF) for authorizing use of domains in E-MAIL, version 1, June 2005. Internet draft; work in progress.

[24] J. A. Zdziarski. *Ending Spam*. No Starch Press, 2005.