Methods to identify spammers

Tobias Eggendorfer
Universität der Bundeswehr München
Werner-Heisenberg-Weg 39
85579 München
Phone: +49 89 6004 2280
tobias.eggendorfer@unibw.de

Abstract **Spam has grown to become a major threat for email communication. Although spam filters' degree of sophistication has increased ever since, they still produce huge amounts of false positives and false negatives thereby reducing the reliability of email. With more and more complex filtering methods implemented, the hardware requirements for mail servers increase to avoid the risk of denial of service situations. Some already claim that mail filtering has reached its limits and ask for more preventive solutions to fight spam. One would be to significantly increase the risk of a spammer being sued for damage compensation or, if legislation permits, for criminal offence. But spammers try to hide their real identity. This paper discusses several methods to identify spammers and analyses under which circumstances they might be a valid proof in court.**

Categories and Subject Descriptors
K5.0 Legal Aspects of Computing – General. Forensics, Identification of attackers.

General Terms
Security, Legal Aspects

Keywords
Spam, Forensics, Address trading, Identification

1. Introduction

Although not anticipated by the founders of the Internet, email has become one of the most accepted and often used applications of the Internet. But with an ever increasing percentage of unwanted email, users slowly start to think about switching to other means of communication. Some use instant messaging instead, others return to the fax, albeit it is more expensive and less convenient.

Although the definition of spam seems to float, with some authors restricting it to unsolicited commercial email and others broadening it up to any unsolicited bulk email, including mass emails sent to distribute viruses, worms and Trojans, hoaxes and even chain letters, they share the observation that spam makes up for the vast majority of all emails sent worldwide, be it more than 80% in July 2007 according to [1] or even more than 97%, as claimed by T-Online, one of Germany's biggest email providers [2].

Unfortunately, spam filters only offer more or less accurate heuristics to help sorting spam and ham, as the opposite to spam is often called. Recent surveys [3][4][5] found that false positives rates of those filters might be as high as 18% and false negatives easily reach 20%. Although false negatives, i.e. spam not marked as spam, are annoying to the user, false positives are more dangerous by far, in a business environment a false positive might have been a customer ordering a product. Failing to notice this message due to an overacting spam filter might not only mean a loss in sales but also liability for not delivering the requested products, thereby increasing the potential financial losses from a false positive by orders of magnitudes [6].

Also, spam filtering increases the risk of security leaks on an SMTP server: The more complex filters are, some even implement OCR to identify image spam, the more computing power they consume, the higher are the requirements on the mail server's hardware. With each and every message taking longer to be processed, the mail server will only be able to handle less requests per second. This again increases the risk of a denial of service attack on the mail server. [7][8][9] provide anecdotal reports.

On the other hand, each additional line of code increases the risk of bugs, which in turn might lead to a remote exploitable security hole, decreasing overall system security.

Taking all this into consideration together with the limited abilities of spam filtering, it is obvious that spam filtering is only a short term solution helping to reduce the symptoms of the spam plague, but not a long term approach.

Therefore, [10][11] propose several techniques to preventively reduce spam to work around the limits of spam filtering. Those methods have their focus on technical methods to prevent spam. However, there might be non-technical ways of reducing a spammer's return on investment, where an investment does not necessarily mean a financial engagement but also other risks, such as being sentenced to prison, a spammer is willing to take in order to earn their living.

Although there is no evidence of spammers assessing their individual risks and calculate their money-worth equivalent, it is likely that they only accept certain risks because of the chance to earn enough to out weight it. Different authors estimated that a spammer's daily income exceeds 5.000 US$ [12][13][14].

If either the risk of being sued for spamming was higher or the expected revenue was less, less spammers were willing to take the risk associated to their business. This is plausible, considering the extremes: With more countries changing their law and declaring spamming as a criminal offence, if every spammer was arrested and sentenced, the risk would be too high to pay off. The other extreme is obviously the spammers income being zero because no one would buy his advertised products.

Unfortunately, in reality, the later extreme seems unlikely to come true soon. Obviously, there are enough users to buy those spamvertised products, although more and more people are aware that buying these is one of the main reasons for the spam problem. Unfortunately, with the low quality of current

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
e-Forensics 2008, January 21-23, 2008, Adelaide, Australia.
© 2008 ICST 978-963-9799-19-6.

spam filters, there is no way to prevent people from buying spam-vertised products either.

2. Organisation of this paper

This paper is organised as follows. Section 3 describes the spam business and the steps necessary to spam. It analyses how division of labour is done in spam business. Section 4 than looks on how spam senders could be identified and analyses problems related to those methods. To our knowledge, the majority of the concepts mentioned there have not been researched for their forensic use yet. The methods described in the following section 5 choice a different approach in trying to identify another player in the spam business, the address trader. Section 5.2 proposes a new method, the usage of a distributed tar pit network, to offer a safe and probative way to investigate an address trader's identitiy. In the last section 6, we conclude and give an outlook on our ongoing research.

3. How spammers work

In order to find ways to attack spammers, it is helpful to analyse how they work, because this might offer hints on how to unmask them. A first step in this analysis is to determine the different tasks involved in a spam run.

In preparation of a spam run, the spammer needs to provide the products intended for sale, acquire email addresses of – from his point of view – potential customers, needs to provide a secure and anonymous payment system and might need to install an online shop or a web site somewhere. To send out the spam, a spammer needs ideally a system not listed on any black list and a somewhat fast Internet connection to send out as many messages as possible. As soon as the first complaints about him spamming are coming in, a spammer needs to have an infrastructure allowing him to work around the ban his provider might have imposed. Also, when the first orders are placed, the spammer needs a delivery system hiding his own identity to protect him from nosy investigators.

3.1. Division of labour

All those tasks do not need to be performed by the same person, although this used to be the case in the early days of spamming. By now, it is a business based on division of labour and highly organised. This has implications on how successful forensic investigations might be in revealing the entire network or just a part of it.

Often, the following services are identified:

1. The spammer
2. The manufacturer of the product sold or the service offered
3. The address collector and seller
4. The rental agent for bot nets
5. The bullet proof hoster

Some therefore consider spamming to be organized crime, some even claim the Russian Mafia to support spam. Whether this is true or not, having to do with criminals sharing their work and often remaining anonymous towards their companions has serious effects on the effectiveness of investigation, because often only parts of a large cooperating network might be uncovered. However, increasing the risk for some involved in a crime, means that they will reassess their chances of earning enough to cover their risk of being sentenced to jail with the consequence of either bailing out or demanding a bigger share of the money earned.

Therefore, from an anti-spammer's point of view, it is enough, if some parties of the spam business are exposed, even though some will have a chance to escape without being sued.

For the following analysis of those tasks, it does not matter who performs them, but where to start investigating them.

3.2. Product provisioning

Generalising, spammers only offer products out of four categories

- tangible goods, such as drugs or coffee machines
- intangible goods, such as mortgages, sexually explicit images or software for download
- services, such as access to “adult” dating communities, email addresses to spam to or even email advertising
- stock spam, where a spammer buys stocks and later advertises them to sell them after they up ticked.

Obviously, products out of the last two categories do not require a complex purchasing processes. Provisioning them is no problem, because there is no physical product.

Intangible products are often as easily provided: Software for download is in most cases a pirate copy easily copied as often as needed, erotic pictures are available from lots of sources in the Internet and might be reproduced as required.

To sell tangible goods, a spammer has basically two options. He might act as a sales agent or sell the items as a vendor himself. The latter means stocking those products or ordering “just in time”, introducing the economic risk of overstocking and privacy risk of sending the items might identify the vendor. Acting as a sales agent however reduces the potential income to the commission the vendor offers.

3.3. Product delivery

Similar problems to provisioning tangible goods are associated to their delivery. If they are mailed with a valid sender's address, the recipient might identify the spammer. However, depending on the product's value, the spammer might have a strong interest in having it returned to him in case it is not deliverable. According to reports in de.admin.net-abuse.mail, a newsgroup, this fact let to the identification of a German spammer selling office coffee machines. To work around this, the spammer might either not give a return address or a faked one or use mail forwarding or mail box services readily available.

3.4. Email address acquisition

Email addresses to spam to are usually either collected from web pages using harvesting technologies [11][15] or from users' hard disks using Trojans. Another way is to persuade users to subscribe their email address to certain services, e.g. an adult web site offering to email daily pictures of a certain kind. Subscribers to those email newsletters might be interested in equivalent offers from other web pages and are more likely to buy related products. Targeted mails raise the response rate from 0,1% on non-target spam to up to 30% [13][14]. Unfortunately, most spammers still use harvesters and Trojans to acquire addresses.

3.5. Payment systems

Dependant on the money earning scheme chosen by the spammer, in most cases a secure and anonymous payment system is a requirement for him. Only if there is no direct customer contact, spammers do not need a payment system. Stock spam being an example, where the spammer buys stocks later advertised and than sells them at an higher price. But in every other case, spammers need to accept payments made by their customers or a third party. A third party is involved, if the spammer acts as a commission paid sales agent or promoter for an online shop. In this case, anonymity requirements might be less of an issue, because the shop operator might be trusted. In all other cases, a person buying a spamvertised product might be an investigator trying to identify the anonymous spammer. Therefore, the spammer needs an online payment system maintaining his anonymity to avoid prosecution. The system however needs to at least look safe to customers, i.e. it should operate on HTTPS or implement anything else a customer might have trust in.

In most cases, spammers want to offer their customers credit card payment. This means, spammers need to have some kind of bank account to where the amount is paid. To avoid being tracked using this account, they often use anonymous debit cards as reference account or offshore bank accounts.

3.6. Anonymous online shops and web sites

Similar anonymity requirements exist for online shops or web pages used to sell or promote the product, because there the server's IP might be traced to the spammer. To avoid this, spammers either order their servers at a so called “bullet proof hoster”, who for a surcharge usually ignore spam complaints and do not ask for identification thereby maintaining their customers' anonymity.

Instead of trusting a bullet proof hoster, some spammers prefer to use cracked servers, where they host their web pages and even shops. This has the advantage, to be almost untraceable, but also the disadvantage of an unexpected interruption of service, either because of the cracked machine's provider disconnecting it due to spam complaints or because the machine's administrator locked the cracker out again. To work around those risks, spammers usually have more than one cracked server ready and use special DNS servers with very short time outs so they might easily change the IP a name points to. From a technical point of view, this is similar to the techniques used for dynamic DNS services, some users use to run servers on their DSL line with a dynamic IP address.

Often those DNS services are offered by spammer friendly providers, to reduce the risk of an DNS entry to be removed due to spam complaints.

A rather new method is to use bot nets to host a web site on. In this case, cracked and remote controlled home PCs are turned into web servers publishing a spammer's web site. As those machines might go off line at any time and might also change their IP, a dynamic DNS-solution is again needed. Often, it also needs to support multiple A records to offer an DNS round robin address resolution [16][17].

3.7. Sending spam

Although bulk mail software is readily available from major download sites [18], it is rather inefficient to use those programmes, because they send their messages from the spammer's Internet connection, allowing to black list his IP and thereby reducing the spam run's effectivity, because the message is filtered out by more spam filters. Additionally, using their own IP, spammers risk their anonymity. Spammers therefore try to send their messages through multiple computers to both distribute their mailing faster to avoid black list updates and to hide their identity. To do so, they rent bot nets.

4. Identification of spammers

A very exposed party in the spam business is the spammer himself. If they were at a higher risk of being sentenced to jail or loosing all their earnings, they might decide to choose a more profitable business model.

Methods to identify spammers are as old as spam, starting with a simple mail header analysis to identify the sender's IP. Considering the increased usage of bot nets, mail headers become less and less useful in tracking a spam mail's source. But observation of the bot nets and who uses them might be helpful in the identification process.

Other methods more oriented on the spammer's work flow include the observation of him buying the goods sold, the payments made by customers or affiliates and the servers used to host spammer's web pages and online shops. To identify stock spammers, several governmental organisations like the federal trade commission (FTC) in the United States started investigating orders placed in context with a spam run.

All those methods are described in more detail in the following subsections and discussed with a view to their efficiency.

4.1. Mail analysis

Each email message consists of a body, where the message meant to be read by the mail's recipient is stored, and a header, containing several technical information on the message, such as the To- and Cc- addresses, the date and the alleged sender's email address. As a mail message might be relayed through several servers in the Internet, each mail transfer agent (MTA) relaying it, adds a header line. In those “received”-headers, an MTA logs the name the remote machine sent during the SMTP's HELO-command, the remote IP-address and often also the reverse DNS entry for this IP. Also, a time stamp is logged [19].

Those headers allow to trace back from where a mail message was sent. This information was used to identify spammers' providers and request them to ban those senders and / or giving out their names and addresses to allow legal action [19][20].

To reduce their risk of being discovered, spammers use either cooperating providers or send their spam from bot nets. Due to this, header analysis has become inefficient.

4.2. Observing bot nets

As bot nets became the major source of spam and are under the control of spammers at the time spam is sent, it is feasible to try to observe bot nets to identify spammers. [21] described how Microsoft tried to support law enforcement in investigating who abuses bot nets for illegal action. According to [21] they installed an out of the box Windows XP system in a monitored network and secured the network in a way, the

machine could not harm any other system. They then waited for the machine to be infected with several worms and Trojans. By monitoring the IPs from where the bots were controlled and logging the commands sent to the bots, they were able to identify the bot net users.

Besides the irony involved of Microsoft using security holes in their own products to identify attackers that would not have been successful if Microsoft had built their software with security in mind, there are a few considerations to be made in order to have a proof accepted in court.

First, the person contracting the Internet access provider must not necessarily be the attacker. If the computer is shared among several persons, each of them is suspicious. As there are thousands of unsecured WiFi access points world wide, an attacker could use any of these to control the bot net. The same is true for Internet cafés: Most of them do not require any proof of identification to use their services, offering a perfectly anonymous access point to bot nets.

Adding the possibilities of computers with remote back doors to this, an attacker has plenty of possibilities to hide his identity. The more systems he adds in between him and the bot net, the better he covers his traces. If the machines he used are located in several countries, law enforcement has to deal with different legal systems and agencies more or less willing to cooperate. Due to privacy laws, in several countries data needed to identify someone based on his IP and time of usage is impossible to get or only available for a very short time, even to law enforcement agencies.

Taking this into account, it is likely to accuse an innocent instead of the real attacker. A professional attacker would take care of hiding behind a few owned systems. Albeit the method is simple, straight forward and easy to implement, its precision in identifying the target is not high enough.

However, the method offers a starting point for further investigation. This investigation should be unintrusive due to the high risk of accusing the wrong person.

4.3. Surveillance of purchases

An approach used by the FTA to identify stock spammers is to look who invested into those stocks prior to them being spamvertised. This scheme might also apply to tangible and intangible goods sold by spammers. However, with those, most of the time, it is harder to track them, even though some are not freely available, such as medications. But just because access to those items is made difficult, this does not imply control and traceability, because there is a black market for restricted goods.

In some cases, spammers might also try to work around those restrictions by selling counterfeited products. This is common practice if boxed software is sold, but also possible for watches or garment. Spammers also started to develop their own drugs, such as “generic Viagra” and “herbal Viagra”. Depending on what those products are based on, their consumption might by a life threat for their consumers. A risk, spammers are still not willing to take. Therefore they often resort to herbal or allegedly homoeopathic drugs, because they believe those to be less dangerous. According to [22], Viagra-substitutes are often made of apricot kernels, which are biologically equivalent to almonds and therefore might cause severe anaphylactic reactions to nut allergics. Also, different species of apricots contain different amounts of cyanide, up to poisonous levels [23]. Even though those “herbal” substitutes are dangerous on their own, their trading is often not controlled. Therefore, it is almost impossible to track persons acquiring those products in quantities needed for spamming.

Just by comparing figures, it is obvious that spammers only account for a very small amount of apricot kernel consumption. The ever increasing world wide production was 2.8 million tons in 2000 [24]. Spammer's buys are unlikely to account for a substantial amount. Therefore, tracking sales of products seems not to be a promising approach, even though it might work for stock spam.

However, if an investigator is able to track down a vendor's sales channels, this might offer a very good starting point, but this is then real world, “off line” investigation, which is beyond the scope of this paper.

4.4. Partner shops

Some spammers try to avoid the somewhat dangerous process of interacting with customers and use partner shops of big web shops. There, they earn a commission on each transaction initiated with their affiliate id and often a bonus for refering new costumers. Most of those shops have an anti spam policy, most of the time saying that commissions earned with spamming will not be paid.

Therefore spammers try to subscribe only a few days prior to commission payment at the web shop, then start their spam run and collect the money before complaints start pouring in [13][14]. By doing so, they evade the risk of not being paid.

Partner shops could prevent this by waiting a certain time between their customer making its purchase and cashing out their affiliate. Although serious shops implement several security measures, some web pages, mostly in the red-light districts of the Internet, are said to be less offended by being spamvertised and therefore have lower security measures set up, thus offering spammers a certain income.

In those cases, the web shop being spamvertised might be liable as accomplice, according to German and some other countries' civil laws. The web shop could then be filed for injunctive relief. To compensate its damages, disclosing the spammer's identity is a possibility: Because the web shop needs to pay the commission, at least a certain minimum of information needs to be known, e.g. a bank or credit card account or an address the spammer has.

This might be a starting point to investigate the spammer's identity.

4.5. Payment process

The same information is needed, if a spammer's customer pays his bill. But again, spammers found ways of working around the risk of being identified by using anonymous bank accounts or credit cards, i.e. the spammer could have his costumers pay to his anonymous credit card's account and then withdraw the money he feels he needs from any ATM.

However, if the spammer's credit card data is known, identifying him might become possible if he withdraws money at an ATM by using the surveillance cameras usually installed to monitor ATMs. However, not all ATMs are secured that way and often, the quality of the pictures they deliver is not good enough to identify someone. Also, a cautious spammer

could try to avoid surveyed ATMs.

All in all, although at first the payment process might seem to be a method of identifying spammers, it is not.

4.6. Server owner

Often, spammer's have a web page dedicated to the product they are currently advertising. This page might contain a web shop, but might also only contain a redirection to some other web page, e.g. if they try to take advantage of a third parties affiliate programme.

Those web pages might be hosted on a proper server, a cracked machine or on a bot net. If the later is the case, identifying the spammer might be possible using the methods described above to identify a bot net's user.

On a cracked server, the cracker might have left traces that identify him, but their analysis is again beyond the scope of this document.

A rented server might be located at a so called “bullet proof” hoster or at any regular provider. As usually most hosters only authenticate their customer's payment details, but hot their claimed identity, spammers could use their anonymous credit cards to hide their identity, making it virtually impossible to track them down without the help of the provider.

However, his log files might help in identifying the spammer, because to install or update his web page, the spammer needs to connect to the web server. This would reveal the spammer's IP, which is a first step in identifying him. However, the same restrictions apply as mentioned above for the bot net's user's IP.

4.7. Discussion

Although there are a few options to identify a spammer, there are work-arounds. Most only provide a first step in identifying him. However, in most real world situations, criminals do not think of all possibilities of hiding their identity, e.g. if they purchase an anonymous credit card, they might do so from their own computer and thus leaving their IP in the log files of the credit card provider, if they do not use an anonymisation service such as JAP [25]. Therefore it is worth investigating each step. But, if a spammer thinks ahead, chances are, he is able to hide his identity.

5. Identification of address traders

Spammers need their potential customer's email addresses to spam to. While acquiring them, traces might be generated leading to address traders or even to spammers. At first glance, the probability of this approach to be effective seems to be higher, as spammers are used to anti spam measures, but anti address harvesting is new and defensive, e.g. by obfuscating email addresses [26]. Attacking harvesters with HTTP tar pits is still new [27].

5.1. Identifying mail addresses

[28] suggested to generate email addresses published on a web site on the fly. Those email addresses should either contain the remote IP address and the time of access or a reference to those. As soon as an email is received, the harvester's IP would be known. Together with the access time, the user of this IP address could be identified.

[28] claims to have identified a German phone book editor as a spammer because he has received an advertisement on an email address he could track back to the company's proxy. When confronted, they denied and threatened to sue him. Unfortunately, the case was not taken to court, therefore there is no legal statement on the proof's quality.

Basically, the usual problems when trying to identify a person based on an IP address arise, i.e. it is only known from which computer the attack was made, but it remains to be investigated who operated that computer. Fortunately, it is still uncommon to use bot nets or cracked machines to run harvesters on. Therefore, the IP seems to be a good starting point for further investigation.

Another issue is the algorithm used to hide the IP address and access time in the email address generated. This algorithm needs to be bijective, that is, a given IP address and time should always generate a unique email address and a given email address should resolve to one and only one IP and time combination. Algorithms like this exists, however, to provide as an proof, those algorithms need to be proven to work as described. The MD5-algorithm [28] used, does not come up to this requirement, as MD5 is not bijective.

The algorithm should also generate email addresses that resemble regular email addresses, i.e. they should neither have a too long local part nor should the local part look like generated. Ideally, the email address is a unique combination of names and maybe a middle initial. Then, a human operator of the harvester is unlikely to notice the trap.

Another disadvantage is the amount of time required to identify the address collector. From the moment the address was harvested to when the email was received, there might be several hours up to weeks. This might give the offender a chance to cover his traces or might otherwise have an negative impact on investigation.

Taking both, the complexity of the required algorithm, requiring sophisticated explanation in court, and the time issue into consideration, this approach is interesting, but might only be of limited use from a forensic point of view.

5.2. Distributed tar pit networks

[29] suggested to use a network of HTTP tar pits to identify harvesters and use this information to block their access to other web pages based on their IP.

An HTTP tar pit is a way to trap harvesters. Simply spoken, the tar pit publishes links to itself, thereby poisoning the list of pages to visit the harvester maintains until finally the harvester is caught in an infinite loop [11].

Because the IP of the harvester is recorded while it is caught in the tar pit, the harvester is identified while collecting email addresses. Therefore we suggest a distributed network of HTTP tar pits as a new method to investigate an address trader's identity, because the HTTP tar pit is a resolution to the time problem described above.

It offers another major advantage: Because harvesters usually revisit a tar pit very often, [30] reports on hundred thousands of visits within a day, the evidence gained is better and offering less excuses to the spammer.

However, harvesting itself is not illegal in most countries, therefore, the HTTP tar pit alone is not a valid proof of spamming. On the other hand, [31] showed that a HTTP tar pit publishing email addresses is by far more effective than it would be without mail addresses. If the tar pit is modified to

publish addresses identifying a certain harvester, then both the act of harvesting and the later spamming could be tracked back to a certain IP address.

But, compared to only publishing specifically crafted mail addresses, in this case, the investigator knows beforehand, from where spamming might later occur and could establish different other surveillance methods. This new method might be useful in identifying the offender.

5.3. Discussion

By combining HTTP tar pit networks used to prevent spammers from collecting email addresses from web pages and used to identify harvesting IPs to protect other web sites, and using specially crafted email addresses, identifying address traders IP and winning a time advantage over them is possible. This seems to be a promising approach. Currently, this seems to be very effective, because harvesting most of the time does not occur from bot nets or cracked machines, but from the address trader's or spammer's own network.

6. Conclusion and further research

This paper discusses methods to identify some of the parties involved in the spam business. Section 4 gave a new insight in current approaches deficiences, section 5 presented new approaches. Although they do not help to discover the entire network of spammers, it increases the risk of being discovered for some of the spammers. Because risks and expected earnings are often strongly correlated, those exposed to a higher risk will raise their services' prices, this again has effects on the other parties, because their economic risks increase due to the higher prices resulting in some spammers to quit spamming.

Even though spammers learned how they might work around being identified while sending out spam, address traders take less precautions. Therefore, identifying address traders seems to be more likely. Our new suggestion is to combine the publication of email addresses crafted to prove that spam has been sent out due to a specific harvesting action and the advantages of HTTP tar pits in identifying harvesters as an effective way to provide court proof evidence has been presented.

Our current research is into finding an algorithm to generate email addresses that meets all requirements mentioned above, specifically, we want it to only generate unique email addresses containing a human name, and to integrate it then into the HTTP tar pit. Currently, the algorithm only provides random alphanumeric email addresses.

References

[1] spam-o-meter, spam-o-meter statistics by percentage, http://www.spam-o-meter.com/stats/index.php, 2007

[2] Kuri, Jürgen, T-Onine verzeichnet eine Milliarde Spam-Mails pro Tag, http://www.heise.de/security/news/meldung/72324.html, 2006

[3] Schulz, Carsten, Erstellen eines Konzeptes sowie Durchführung und Auswertung eines Tests zur Bewertung unterschiedlicher Spam-Filter-Mechanismen bezüglich ihrer Langzeiteffekte, Master thesis, Universität der Bundeswehr, Neubiberg, 2006

[4] Eggendorfer, Tobias, Spam slam. Comparing antispam applicances and services in: Linux Magazine (International Edition) 03/2007, Linux New Media, München, 2007

[5] Hosbach, Wolf, Test Spam-Filter. ...die Schlechten ins Kröpfchen! in: PC Magazin 10/2006, WEKA Computerzeitschriften-Verlag, München, 2006

[6] Heinlein, Peer, Genervt, blockier gefährdet: Wie sich Firmen gegen Spam & Viren schützen können in Proceedings of CeBIT 2007, Hannover, 2007

[7] o. A. (apa), Spam-Attacke blockiert E-Mail-Verkehr in: derStandard.at/Web, derStandard.at, Wien, 2003

[8] Frei, Stefan, Angriff via Mail. Mailserver als Verstärker für DoS-Angriffe in: Heise security, Heise, Hannover, 2004

[9] Schüler, Hans-Peter, Spam-Welle überrollt die TU Braunschweig, http://www.heise.de/newsticker/meldung/47575, 2004

[10] Eggendorfer, Tobias, Methoden der Spambekämpfung und -vermeidung, Dissertation, FernUniversität in Hagen, BoD, Hagen, 2007

[11] Eggendorfer, Tobias, Methoden der präventiven Spambekämpfung im Internet, Master thesis, Fernuniversität in Hagen, München, Hagen, 2005

[12] Ilgner, Michael et al., The Economoy of Spam in: , Universität Wien, Wien, 2006

[13] Spammer X, Inside the spam cartel. Why spammers spam, Syngress Publishing, , 2004

[14] Spammer X, Talk by Spammer X in Proceedings of EU Spam Symposium, , 2006

[15] Center for Democracy and Technology, Why am I getting all this spam?, http://www.cdt.org/speech/spam/030319spamreport.pdf, 2003

[16] Partridge, Craig, Mail routing and the domain system, http://www.ietf.org/rfc/rfc0974.txt, 1986

[17] Brisco, Thomas, DNS Support for Load Balancing, http://www.ietf.org/rfc/rfc1794.txt, 1995

[18] Eggendorfer, Tobias, Tweak your MTA. Spam-Schutz mit Tricks in Proceedings of 3. Mailserverkonferenz, Berlin, 2007

[19] Wood, David, Programming Internet Email, O'Reilly, Sebastopol, 1999

[20] Hochstein, Thomas, FAQ. E-Mail-Header lesen und verstehen, http://www.th-h.de/faq/headerfaq.php3, 2003

[21] Kornblum, Aaron E., "John Does" no more: Exposing Zombie Spammers in Proceedings of M.I.T Spam Conference 2006, Cambridge, MA, 2006

[22] McWilliams, Brian, Spam Kings. The Real Story Behind the High-Rolling Hucksters pushing porn, pills, and @*#?% Enlargements, OReilly, Sebastopol, 2005

[23] Suchard JR; Wallace KL; Gerkin RD, Acute cyanide toxicity caused by apricot kernel ingestion in: Annals of Emergency Medicine 12/98, Mosby, Dallas, TX, 1998

[24] Asma, Bayram Murat, Malatya: World's Capital of Apricot Culture in: Chronica Horticulturae 01/2007, ISHS, Leuven, 2007

[25] Eggendorfer, Tobias, Ghost Surfing. Anonymous surfing with Java Anonymous Proxy in: Linux Magazine (International Edition) 11/2005, Linux New Media, München, 2005

[26] Eggendorfer, Tobias, Dynamic obfuscation of email addresses - a method to reduce spam in Proceedings of AUUG 2006, Melbounre, 2006

[27] Eggendorfer, Tobias, SMTP or HTTP tar pits? Which one is more efficient in fighting spam? in Proceedings of AUUG 2006, Melbourne, 2006

[28] Rehbein, Daniel A., Adressensammler identifizieren - Ein Beispiel, http://spamfang.rehbein.net, o. A.

[29] Eggendorfer, Tobias; Keller, Jörg, Dynamically blocking access to web pages for spammers' harvesters in Proceedings of IASTED Conference on Communication, Network and Information Security CNIS 2006, Cambridge, MA, 2006

[30] Eggendorfer, Tobias, Stopping Spammers' Harvesters using a HTTP tar pit in Proceedings of AUUG, Sydney, 2005

[31] Eggendorfer, Tobias; Keller, Jörg, Combining SMTP and HTTP tar pits to proactively reduce spam in Proceedings of SAM 2006 (The 2006 World Congress in Computer ScienceComputer Engineering, and Applied Computing), Las Vegas, Nevada, 2006