

Introduction

This work is about spam e-mails, which are just one type of spam we face in electronic communication. Other types are related to SMS, chats, or Internet phone (Spam over IP Telephony). However, issues relating to these are beyond the scope of this work. In this introduction, we describe the problem that (e-mail) spam causes, and its history. We also define the goals of this work, how they are addressed (methodology), and how this work is structured (architecture).

1.1 The problem

Most of us using the Internet e-mail service face almost daily unwanted messages in our mailboxes. We have never asked for these e-mails, and often do not know the sender, and puzzle about where the sender got our e-mail address from. The types of those messages vary: some contain advertisements, others provide winning notifications, and sometimes we get messages with executable files, which finally emerge as malicious codes, such as viruses and Trojan horses. Apparently, the Internet e-mail infrastructure is widely used, as well as misused, as an efficient medium for information distribution. Senders of bulk e-mail benefit from the anonymity that is inherent to the e-mail infrastructure: sender data can be easily spoofed, and remotely controlled PCs can be used for sending e-mails. The design principles of the e-mail infrastructure, which were originally intended to provide simplicity and flexibility, have become ambivalent characteristics.

There are a number of methods in use for managing unsolicited bulk e-mail, which is termed “spam”. Many organizations employ filtering technology and construct elaborate rules that determine which senders are allowed to connect or deliver e-mail to their networks and which are to be blocked. However, even with good filters, which are the most deployed type of technological anti-spam measures, we have merely heuristics on hand, that sometimes misclassify e-mails: whereas a spam e-mail in our mailbox might not seem bad, an e-mail

that has been erroneously classified as spam and remains, therefore, unnoticed, does. In such a case, an anti-spam measure is even counterproductive. Although policies and technology measures can be effective under certain conditions and help to maintain Internet e-mail a usable service, over time, their effectiveness degrades due to increasingly innovative spammer tactics. It is humbling to note that, for many years, statistics have shown that the number of spam e-mails is higher than the number of “regular” e-mails (ham e-mails).

Today, spam has even crossed the borderline between simply being annoying for private users and causing economic harm. For example, companies invest money in anti-spam software and IT staff, and they lose productivity of employees when these spend time in opening, reading, classifying e-mails as spam, and deleting them. Private users lose money due to fraud e-mails including phishing attacks. The worldwide economic harm caused by spam is estimated at hundreds of billion USD per year. This huge economic relevance of spam has motivated the national authorities of both many countries and federal states to address spam by legislation. However, despite some spammers being prosecuted, the effectiveness is limited, because e-mail messages today do not contain enough reliable information to trace them back to their true senders.

Beside technological and legislative anti-spam measures, organizational and behavioral measures have been proposed. However, many of these approaches still fail to address the root problems: first, sending bulk e-mail is a profitable business for spammers; and second, e-mail messages today do not contain enough reliable information to enable recipients to consistently decide whether messages are legitimate or forged [9]. Moreover, today’s deployment of anti-spam measures resembles a (still open-ended) arms race between the anti-spam community and spammers. Even worse, we, generally, allocate resources of the recipients of e-mails to fight spam, instead of increasing the senders’ need for resources.

What is currently lacking is the development and deployment of long-term, effective anti-spam measures, which keep Internet e-mail alive as a reliable, cost-effective, and flexible service. However, it is not necessary to “reinvent the wheel”, the analysis of the combined application of already proposed solutions may also help in this regard.

1.2 The history

The etymology of the word “spam” is, usually, explained by using an old skit from Monty Python’s Flying Circus comedy program (for example, see Merriam-Webster’s Collegiate Dictionary): In the sketch in question, a restaurant serves all its food with lots of Spam, which is canned meat and an acronym for “**S**houlder of **P**ork and **H**am”. The waitress repeats the word several times in describing how much Spam is in the dishes on the menu. When she does this, a group of Vikings in the corner start singing a chorus of “SPAM, SPAM,

SPAM...” at increasing volumes in an attempt to drown out other conversations. As “unsolicited bulk e-mail” disturbs Internet communication likewise, it was termed “spam”.

In the literature, unwanted e-mail messages were being recognized as a problem in an Internet Request for Comments as early as 1975 ([134]) and in the pages of *Communications of the ACM* as early as 1982 ([41]).

Possibly the first spam ever was a message from a DEC marketing representative to every Arpanet (the predecessor of the Internet) address on the west coast, or at least the attempt to do so ([173]). In April of 1994, the term “spam” had not yet been born, but it did jump forward a great deal in popularity when two lawyers from Phoenix, named Canter and Siegel, posted a message advertising their fairly useless services in an upcoming U.S. “green card” lottery [20]. This was not the first such abusive posting, nor the first mass posting to be called a spam, but it was the first deliberate mass posting to commonly receive that name. Some more examples of early spam attacks are presented by Templeton [172].

1.3 Goals, methodology, and architecture

The still existing occurrence of spam e-mails in bulk proves that currently deployed anti-spam measures are low effective. However, this does not necessarily imply their inappropriateness as a matter of principle. One primary goal of this work is the methodical analysis of anti-spam measures in terms of their potentials, limitations, advantages, and drawbacks. These determine to which extent the measures can contribute to the reduction of spam in the long run. The range of considered anti-spam measures includes legislative, organizational, behavioral, and technological ones.

Legislative measures As legislative measures can vary in many regards, we provide a classification scheme for them. This scheme is based on attributes, whose instantiations determine the effectiveness of the particular legislative measure. We describe this determination on an abstract level and then analyze the anti-spam legislation of many countries with regard to the classification scheme (microscopic view). From a macroscopic point of view, we assess today’s overall legislation landscape in terms of effectiveness, we identify currently unsolved problems, and we indicate means by which some limitations might be overcome.

Organizational measures We subsume abuse systems and (types of) international cooperation under organizational measures. This part is mainly descriptive, but it also shows the possible types of cooperation between national authorities, other non-profit organizations, companies, and users.

Behavioral measures Behavioral measures aim at e-mail users' procedures in using and distributing their e-mail addresses (ex ante behavior) and dealing with any spam e-mails which they receive (ex post behavior). With regard to the ex ante behavior, we identify locations where e-mail addresses can be harvested from. In order to support the empirical analysis of spammers' behavior concerning the collection and the usage of e-mail addresses, we provide the conceptualization and prototypic implementation of a honeypot. The evaluation of the honeypot data reflects the present behavior of spammers. We present mechanisms that allow for protecting e-mail addresses from being automatically collected. Concerning the ex post behavior, we provide a description and an analysis of options that the users have, once spam e-mails have found their way into their e-mail boxes. The findings of the analysis of behavioral measures can be used for the development of e-mail user guidelines. However, this issue is beyond the scope of this work.

Technological measures The vast majority of proposed anti-spam measures is technological-oriented. In order to maintain an overview of the methods, we propose several classification schemes. We describe technological anti-spam measures by following the functional classification. For the analysis of the effectiveness of anti-spam measures, we use the classification according to whether their application only refers to particular delivery routes that e-mails take or whether the measures are applicable independently of delivery routes. Whereas the former group of measures are analyzed informally, the latter are assessed formally: we provide a formal (graph) model of the Internet e-mail infrastructure, use automata theory to derive and categorize all possible delivery routes a spam e-mail may take (spamming options) and which any holistic anti-spam measures would need to cover. Finally, the effectiveness of (route-specific) anti-spam measures is analyzed relative to covering the identified spamming options.

The analysis of the various anti-spam measures shows that no single measure is the "silver bullet" against spam, and it is doubtful whether any single, simple solution will ever be able to reduce or stop spam. Rather, it seems appropriate to look for solutions that provide a complementary application of several anti-spam measures. The second primary goal of this work is, therefore, the conceptual development and analysis of an infrastructural e-mail framework, which features such a complementary application. After the presentation of the technological and organizational facets, the framework is analyzed twofold: its theoretical effectiveness is assessed with the aid of the formal model mentioned above, its storage and traffic requirements are analyzed quantitatively. We further consider deployment issues, as the framework would have to be integrated in both the technological and the organizational Internet infrastructure.

A graphical overview of the different parts of this work and their dependencies is given in Fig. 1.1. As the description of the empirical analysis of address abuse does not need necessarily to be read in order to follow the thread of this work, we put it at the end of the book. Besides the contents described above, this work first addresses two elementary issues: (1) It provides an introduction to spam and a motivation for addressing spam scientifically. (2) It explains the technological facet of the Internet e-mail delivery process and its susceptibility to spam.

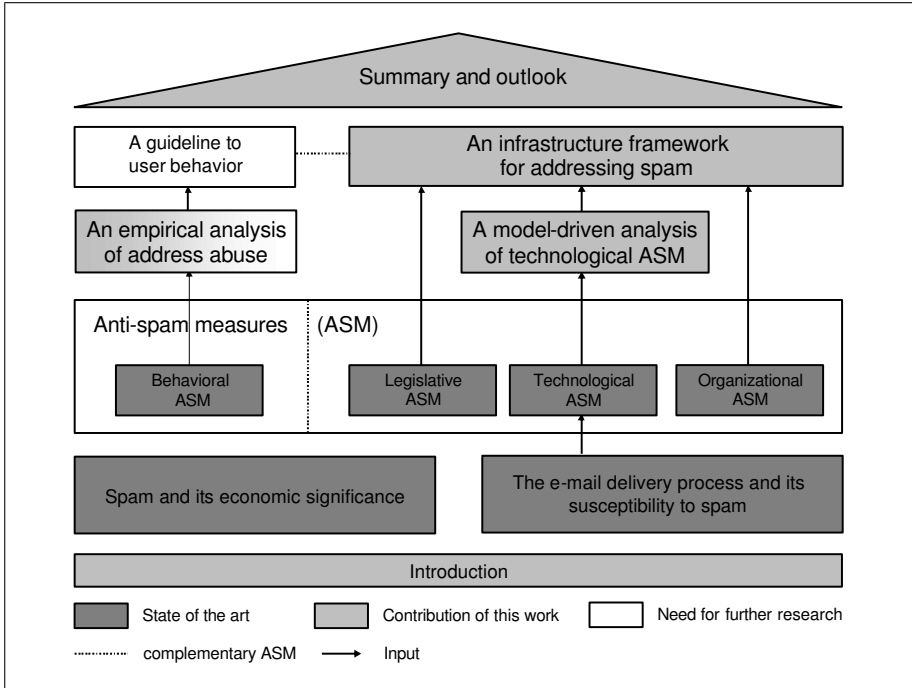


Fig. 1.1: Architecture of this work



<http://www.springer.com/978-3-540-71748-5>

Anti-Spam Measures

Analysis and Design

Schryen, G.

2007, XX, 209 p. 50 illus., Hardcover

ISBN: 978-3-540-71748-5