

# On the Detection and Identification of Botnets

Alexander K. Seewald<sup>a</sup>, Wilfried N. Gansterer<sup>b</sup>

<sup>a</sup>*Seewald Solutions*

<sup>b</sup>*University of Vienna, Research Lab Computational Technologies and Applications*

---

## Abstract

We develop and discuss automated and self-adaptive systems for detecting and classifying botnets based on machine learning techniques and integration of human expertise. The proposed concept is purely passive and is based on analyzing information collected at three levels: (i) the payload of single packets received, (ii) observed access patterns to the darknet at the level of network traffic, and (iii) observed contents of TCP/IP traffic at the protocol level.

We illustrate experiments based on real-life data collected with a darknet set up for this purpose to show the potential of the proposed concept for (i) and (ii). In (iii) we use a small spamtrap as darknets cannot capture TCP/IP traffic data, so this experiment is not a purely passive approach, but traffic moving through a network could be analyzed in a similar way to obtain a purely passive system for this step as well.

*Key words:* botnet, e-mail spam

---

## 1. Introduction

Undoubtedly spam is a problem of global impact that keeps getting bigger. Already in 2001, according to a study undertaken for the European Commission [12], Internet subscribers worldwide wasted an estimated 10 billion Euro per year just in connection costs due to spam. Moreover, the quantifiable damage caused by successful phishing attempts was estimated to be millions of Euros in Germany in 2006.<sup>1</sup> The total damage caused

---

*Email addresses:* alex@seewald.at (Alexander K. Seewald),  
wilfried.gansterer@univie.ac.at (Wilfried N. Gansterer)

<sup>1</sup><http://www.at-mix.de/news/1226.html>

is even higher, since these figures do *not* yet account for loss of reputation and reduced customer trust. The economic impact is only part of the problem—waste of time, resources and the gradual erosion of trust in e-mail communication must also be considered significant side-effects of the spam deluge.

Spam filters have become reasonably effective by now (see, for example, [21, 5]). However, new challenges, such as new forms of image spam and audio-based spam may lead to decreasing performance of today’s most widely used Bayesian spam filters. More generally, history has shown that spammers have adapted again and again to effective adversaries and countermeasures. Most of the widespread countermeasures against spam are *reactive* in nature and thus this “arms race” is very likely to continue in the future. In particular, it is a big challenge to reduce the overhead and waste of resources caused by spam [13, 10].

But spammers are already diversifying their operations, e.g. with *phishing*—faking whole online banking websites to harvest login and transaction data from a large userbase, which motivates the investigation of methods specifically targeted for this challenge [11]. Recently, a widespread vulnerability in today’s DNS system has been detected which would have allowed spammers to fake DNS entries and thus forward legitimate banking sites’ URLs to the spammers’ sites, bypassing the current more costly harvesting effort via phishing mails with embedded links to the faked sites, which depend on user naivety. Although the vulnerability has been partially fixed in the meantime [3], sufficient vulnerable servers still exist for this to be a continuing threat.

Another alternative source of revenue comes from *Pump-and-dump* spam, which is an approach to link spamming directly to the market. By sending out millions of spam mails which aim to boost a specific company’s share value, the resulting small transient peak or following valley<sup>2</sup> is used to make money directly via leveraging in form of call or put options.

Both approaches make the spammers’ revenue independent of companies paying for their services. Recent news about systematically stolen user data of large companies, to be used, e.g., for identity theft and credit card fraud, is another indication for that. By diversifying their sources of revenue, spammers’ will in time become more independent and botnets will find new uses.

---

<sup>2</sup>Once people notice they have been fooled, they are desperate to sell.

Most if not all of today's spam is sent via large networks of captured machines which are under direct control of spam operators.<sup>3</sup> These large networks have been named *botnets*, from software *robot networks*. One of the largest spamfilter companies, SpamHaus<sup>4</sup>, has estimated that already in 2004, 70% of spam was sent out via such networks [16]. Computers without the latest security updates attached to the internet may be infected rapidly depending on the security of their operating system – via known security holes, by users clicking or only viewing infected mails or via surfing past an appropriately prepared webpage, using browser or browser plugin vulnerabilities. Even secure computers may be prone to *zero-day attacks*, which rely on the delay between the publication of a weakness and the availability of a patch to correct it. There are some indications that botnet operators invest in R&D to find specific zero-day vulnerabilities, aiming at exploiting them at leisure. At least a million bots are known to exist according to the US FBI and public trackers like ShadowServer<sup>5</sup>, and the true number is likely to be much higher. Another worrying development is that although the security of newly deployed systems is clearly increasing, the number of bots is still growing at an exponential rate. This indicates that the potential to build new and larger botnets is not exhausted yet. A bot, being subverted, can harm the user in any number of ways: by capturing passwords, PIN and TAN numbers for telebanking, rerouting their online banking website to a page under the spammer's control, and so on.

By distributing their workload over several hundreds or thousands of machines, spammers are able to use combined bandwidth of gigabits per second to send out spam rapidly. But bots are not only used to send out spam. Different bots exist which attack large parts of the internet infrastructure via distributed denial-of-service attacks (DDoS<sup>6</sup>), install backdoors and spyware to harvest password and login information and allowing direct access to all files accessible from the machine, defraud ad networks with pay-per-click systems like Google AdSense by faking millions of clicks from hundreds of thousands of machines, crack weak passwords in a brute-force manner by us-

---

<sup>3</sup>According to SpamHaus, about 300-400 spammers are responsible for 80% of world-wide spam [18].

<sup>4</sup>[www.spamhaus.org](http://www.spamhaus.org)

<sup>5</sup>[www.shadowserver.org](http://www.shadowserver.org)

<sup>6</sup>These attacks are very effective and were in the past able to bring down the internet infrastructure of small countries.

ing the botnet as a computational grid to distribute workload among a large set of machines, or create specialized peer-to-peer networks for obfuscating the operator's point of origin.

The ease with which it is possible to create new malware variants for infecting machines with bots has human analysts at a great disadvantage. While botnet operators are able to automatically generate new malware variants in unprecedented numbers, out of necessity the anti-virus community has become focussed on just *detecting* malware. It clearly has become infeasible to analyze manually what each of the thousands of new variants of malware appearing each and every hour are capable of. Botnet software nowadays is capable of automatically downloading updates from the internet to stay ahead of anti-virus detection, so malware analysis would have to be done in near real-time as well. However, with only detection by antivirus software, it becomes almost impossible to fully understand the functionality of a given malware *before* it spreads on the internet. Recent approaches to characterize malware automatically such as [22] seem promising, but currently do not scale sufficiently well and should still be considered open research problems.

To summarize, botnets constitute a major threat for the internet community and attack simultaneously from several directions:

- They disable or hamper internet infrastructure even during their *normal* operation,
- they enable or facilitate cybercrime such as phishing, pump-and-dump spams, click-fraud, harvesting private user data for identity theft and fraud, etc., and
- they reduce the chances of capturing spammers by obfuscating the point of origin.

### *Objectives*

The objectives of this paper can be summarized as follows. As the botnet operators have resorted to extensive automation – in the generation of new malware bot variants, in the deployment, command and control of these variants, and in diverse profit-gaining activities – it is necessary to automate the detection, identification and tracking of botnets as far as technically feasible. As the botnets operate on a global level beyond national boundaries, so we also have to operate on a global level. As the botnet operators continually update and extend their systems, so we also have to continually update and

extend our identification systems. To be able to combat botnets proactively and reduce their effectiveness for any kind of usage – including spam – we must first learn more about them.

We distinguish three levels at which botnet analysis may take place and investigate various aspects and open issues on each of these levels. A central focus is on how to replicate human expertise – in the identification of specific bot types, specific bot net types and specific SMTP traffic patterns – by appropriate machine learning systems, which enable natural update and extension of our systems by simple retraining. A widespread worldwide sensor network using all three levels identified here should be able to achieve all our objectives.

The final goal must be to characterize specific botnet commands to enable the identification of the exact perpetrators regardless of all obfuscation tactics, to facilitate and in some cases enable swift action by law enforcement, and to develop proactive countermeasures. This remains *the* challenge for the future, incorporating not only formidable technological but also major legal issues.

## 2. Background and Related Work

Two types of approaches can be distinguished in the area of botnet detection and botnet identification: *active approaches* and *passive approaches*. Active approaches entail defanging bots or writing specialized bots to simulate the behavior of a real bot. Each way enables the bot to log on to a control center and watch the activity there, and usually enables direct interference with the operation of the botnet.<sup>7</sup> Passive approaches are those that cannot be detected by any means, because there is no flow of information back to the botnet operator. They work with more subtle information sources than active approaches, and do not allow direct interference with the operation of the botnet. Rather, the passive approach intends to be used as a tool to study botnets and serve as an early warning system.

Active approaches comprise all kinds of techniques which make the botnet operator directly or indirectly aware of observation. Examples for active sensing include capturing bot malware and deactivating its malicious parts (i.e. defanging bots), and subsequent analysis of the commands sent and

---

<sup>7</sup>This will likely lead to countermeasures. Silent bots of this kind are currently ignored, but this may change in the future.

received over the network during prolonged execution of the now harmless bot. Defanging a bot is a complex task that at present can only be done through a painstakingly manual process. While defanging need not be done when bot communication is unencrypted, it is essential otherwise. Honey pots and honeynets also use active methods to lure malware, e.g., by simulating known vulnerabilities. While active approaches seem a tempting approach at first sight, they have the big disadvantage that they can easily be detected. Once this happens, botnet operators will inevitably adapt and circumvent any measures taken against botnets. If focussed, aimed and applied in a short time period, active approaches might be helpful for a once-in-a-lifetime shot at botnet operators, but this is not likely given the diversity of approaches currently followed without global supervision. The strong manual component also make analyzing the majority of malware infeasible, and therefore this approach is used only to track the most prevalent bots.

Passive approaches analyze traffic which the botnet generates and secondary effects of its usage (e.g. broken packets resulting from a distant DDoS attack). An example of a completely passive approach is the use of darknets in packet capture instead of using a machine which appears vulnerable (low interaction honeypot), or actually is vulnerable (high interaction honeypot) for attracting botnet attacks, malware, or spam. Both honeypot types can be detected by botnet operators: the first one for example by testing the functionality of the emulated service which is often very incomplete, and the second one for example by fingerprinting the operating system after it has been successfully compromised and then excluding known honeybot configurations. The analysis of existing traffic, e.g., from Tier 1 network as in [14], without introducing significant latency, also is a purely passive approach in this sense. Purely passive methods are more complex to set up, test and adapt, but have the big advantage that their detection is completely impossible (for example, in the case of darknet data capture without SYN packet reply, where the capturing machine is indistinguishable from an unused IP address).

The approach investigated in this paper is based on passively observing network traffic arriving in a *darknet*, i. e., an IP address range which is not used for any active machines. The term “darknet” is sometimes also used in a different meaning, namely for a closed private network of computers, for example, used for file sharing. In Level 3, we needed to augment this data with received spam mails from our spam trap. While the analysis presented there is thus not a passive system, the same data can be collected in a passive

manner by analyzing SMTP traffic moving through a network.

Contrary to intrusion detection and prevention, which is concerned with analyzing attacks and cutting attackers off before their attack has succeeded, our approach is concerned with analyzing passively observed network traffic to unused IP ranges. The quantity and quality of our traffic data is significantly different from the one seen by intrusion detection and prevention systems.

In fact, in preliminary experiments we found that standard intrusion detection and prevention tools such as snort<sup>8</sup>, p0f<sup>9</sup>, fl0p<sup>10</sup>, honeysnap<sup>11</sup>, nmap<sup>12</sup> and xprobe2<sup>13</sup> are not suitable for our traffic data as they are geared towards active attack patterns which simply do not appear in our data.

Reference data was provided by Marshal Ltd. (<http://www.marshall.com>) from their semi-automated TRACE system. We matched their IP and timestamp data to ours, taking into account that dynamic IPs are a reliable indicator of the underlying machine only for a short time period.

We will now present related research, grouped into active and passive analysis approaches as outlined above.

### *2.1. Active Analysis*

In [6], a model of botnet propagation using time zones is proposed. It is based on the assumption that potential bots – being machines of private internet users – will be switched off during the night. They used a DNS redirection technique to redirect known IRC C&C<sup>14</sup> servers to IP addresses under their control. They note that over a six month period, they redirected 50 botnets. Their diurnal model of botnet prediction works reasonably well at predicting botnet population growth (better than basic Susceptible-Infectious-Removal

---

<sup>8</sup>Snort is an open-source tool and the de-facto standard for intrusion detection and prevention. [www.snort.org](http://www.snort.org)

<sup>9</sup>p0f is a passive operating system fingerprinting tool. <http://lcamtuf.coredump.cx/p0f.shtml>

<sup>10</sup>fl0p is a passive L7 flow fingerprinting tool. <http://freshmeat.net/projects/fl0p/>

<sup>11</sup>honeysnap is a tool to analyze significant event in recorded packet files. <https://projects.honeynet.org/honeysnap>

<sup>12</sup>Network Mapper - an active scanner for vulnerabilities. <http://nmap.org>

<sup>13</sup>xprobe2 is an active operating system fingerprinting tool. <http://xprobe.sourceforge.net>

<sup>14</sup>IRC bots use a central server to get their commands, which communicates via the IRC protocol and is called Command & Control center.

(SIR) models from epidemiology [8]) and may be used to prioritize botnets by propagation potential. However, their type of data collection is a slow manual process and very disruptive to botnet operation and therefore easy to detect. Obviously, it would be possible to combine alternative data collection approaches (such the one discussed in this paper) with their approach.

Various graphs of botnet connectivity are analyzed in [7], and a taxonomy of botnets based on topological structure is proposed, with appropriate targeted responses in each case. While their test case is trivial, their work has some merit as a guideline to attack future botnets with more robust communication architecture than a central C&C server (as in IRC).

## 2.2. Passive Analysis

[9] propose a framework to guess protocol types based on classical traffic analysis, without reference to the content transferred. Such a system may be able to recognize even encrypted botnet communication traffic for peer-to-peer botnets by its characteristics.<sup>15</sup> This would be similar to our Level 2 approach discussed in Section 4.2. A set of relatively simple features easily computed from a traffic stream in real-time was able to distinguish normal UDP traffic, Skype traffic, NetBIOS and Real-time Transfer protocol (RTP, used for Voice-over-IP). Similar to their results, we also found that simple features work reasonably well.

[4] propose a network quality measure based on spatial and temporal *uncleanliness* (i.e. the proportion of bots among all IP addresses measured along time and space (subnet) axis) that tries to predict future botnet addresses based on previously known botnet addresses. While interesting in principle, the work is hampered by not taking into account dynamic IP addresses, which influence both temporal and spatial uncleanliness and have not been properly addressed.

[17] propose to use DNS blacklist queries to detect bots who are not yet on the blacklist. Their detection is a purely passive approach, and even their countermeasures are covert and may not be detected by botnet operators for some time. It should be noted that in the meantime most DNS blacklists have resorted to only return blacklist hits to query hosts with DNS MTA entries (i.e. verified mailhosts). As such, DNS blacklist lookups can no longer be used by most bots as they will always return NXDOMAIN. Queries may still

---

<sup>15</sup>Steganographic approaches would still remain undetected.



be sent until botnet operators have adapted to this change. In any case, we found that more than 97% of our bots are not on blacklists so explicit lookups are perhaps no longer needed. The high prevalence of dynamic IPs may also harm DNS blacklists' effectiveness (see Section 5).

A more promising approach has been discussed in [14], who proposed using traffic summaries of data flow collected on a large Tier-1 ISP network. In our terminology, their approach would also be classified as a Level 2 approach (see Section 4.2). They propose an anomaly-based passive analysis algorithm, which has been used to detect IRC botnet controllers with less than 2% false positive rate. This still needs to be combined with a secondary more detailed analysis to reduce false positive rates further as otherwise legitimate services might be tagged as botnet controllers. However, very few institutions are likely to get direct access to a Tier-1 network, which strongly hinders the further development of this line of research. Non-centralized communication between bots was not considered.

[1] report results from a multi-year project, *Internet Motion Sensor* in darknet traffic analysis, for an almost passive distributed darknet<sup>16</sup>. They demonstrate the usefulness of the darknet in tracking the spreading of the Blaster Worm, the Bagle backdoor and the SCO denial of service attacks in 2003 and 2004. This darknet has been discontinued since then.

### 3. Synopsis

In this paper, we investigate a mostly passive botnet defense approach which proceeds at three hierarchical levels. Based on network traffic data observed in a darknet and on reference data, machine learning techniques are applied for automatically detecting and identifying existing botnets.

For clarity, we will first describe the general framework within which we have chosen to summarize our approach in general terms without too many concrete examples. From Section 5, we will describe work we have done along these lines on each of the three levels to be introduced.

---

<sup>16</sup>They sent out ACK packets in response to TCP connection attempts via SYN, which could easily be detected but increases available traffic significantly.

#### 4. A Three-Level Approach

There are several meaningful levels at which the identification of botnets can take place.

**Level 1:** At the level of a *single packet*, we can try to distinguish whether its payload is malicious or spam, whether it corresponds to a remote check for vulnerabilities, or whether it follows unusual conventions with respect to flags and TCP options.

**Level 2:** At the level of *network access*, we can try to distinguish, for example, the pattern of access to a given darknet (sequential or random access, time between accesses to a single or multiple darknets, distribution of access per type and time etc.), pattern of communication activity between bots (in case of peer-to-peer botnets) or with the command center (in case of centralized botnets) even without analyzing the content. *Traffic analysis* in the usual military meaning is exactly this.

**Level 3:** At the level of *TCP conversations*, we can try to distinguish legitimate TCP conversations from illegitimate TCP conversations – e.g., SMTP connections sending spam and malware vs. those sending regular e-mail, SMB transactions known to trigger vulnerabilities in the SMB protocol, etc. Contrary to the previous level 2, here we also analyze the content of many packets and entire conversations and are thus limited to unencrypted or weakly encrypted communication channels.

While the activities at each of these levels will be different, it makes sense to view these levels as incremental. For example, feature descriptors for single packets which allow for identifying spambot types – as in our approach later on – might be reused to compute summary statistics on the level of network access over a large set of packets, and even on the third level of entire TCP conversations. Patterns in network access can be reused to speedily prioritize suspect TCP conversations, e.g., when one machine initiates connections to the same destination port of many other machines. The reuse of patterns and data from lower to higher levels makes a machine learning approach for all levels feasible. Combined with feedback on known good and bad examples on each level, a learning passive botnet detection and tracking system no longer hampered by restrictions of human processing speed becomes a real possibility.

As mentioned before, the big advantage of our approach is that it will remain undetected by botnet operators. Honeynets, the analysis of malware in sandboxes, and other active techniques (such as IRC detectors which play bot and record botnet commands, and can even be used to take down botnets) can potentially be detected, ignored and even subverted by botnet operators. Not only virtual sandboxes are vulnerable as no simulation is 100% accurate – even physical sandboxes are vulnerable, as the necessary reuse of operating system images allows to determine fingerprints for known sandbox environments or the special-purpose hardware to replicate harddisc images on startup, and to integrate these fingerprints into new bot versions. On the other hand, such a functionality could also be used to immunize certain operating system variants against the infection with bots. Additionally, the active approach needs deep knowledge of bot malware which is not available for most circulating variants.

#### *4.1. Level 1: Analysis of Single Packets*

The challenge in single packet analysis is to get the most information out of a single packet sent out from a bot-infected machine. Combined with darknets, which can be converted into silent detectors by traffic forwarding, it is possible to construct completely passive detectors for bots.

Alternatively, instead of using a darknet, traffic moving through a network could be analyzed on packet level with similar methods. Robust models of botnet-specific packets may be built to detect any botnet activity within the system, based not only on traffic concerned with propagation (e.g. packets with malicious payload), but also on communication between bots and malicious bot activity addressed to machines within or without the network. Nevertheless, the available data which can be obtained without any risk of being detected by botnet operators is severely limited.

#### *Open Issues*

One major concern is that botnet operators could potentially randomize packet payloads, and use open-source TCP stacks resp. the operating-system TCP stack. This would limit the applicability of fingerprinting and restrict the scope of our proposed passive darknet approach to UDP packets. It would then be necessary to switch to more active approaches, e.g. sending a SYN packet as response to a TCP SYN packet which asks for establishing a connection packet. This strategy has been reported by [1] to increase available data significantly, but also increases the chance of being detected.

Clearly, receiving just a single SYN packet from every IP address within a large subnet, and never receiving anything else, makes this darknet approach highly suspicious to appropriate darknet detectors. This may be the reason why previous darknet approaches were not very successful.

Another option is to analyze packets *in transit*, i.e., when they are traveling to other vulnerable machines, at firewalls or at the network routing level. The same methods could be used, but far more care would be required to distinguish good packets from bad packets. Whereas in darknets almost every arriving packet can be assumed to come from a malicious source, this is obviously not true here. In fact, changes in machine behavior as evinced by sending unusual packets after receiving unusual packets may even be used to automatically train such a system, but this is almost a Level 2 approach (see Section 4.2). [14] proposed such an approach, but they only analyzed very simple summary statistics of traffic data rather than doing a full analysis on the packet level.

A single packet only offers limited information. Still, analysis on this level works surprisingly well and should be more often implemented (see Section 5). The lack of data on correlations between spambot types and the corresponding packets sent out is an issue, though. Here, the availability of automatically generated data (e.g. via sandboxes) from known spambot types could create sufficient training data to test more complex approaches.

Another more long-term approach, especially in the view of the potential vulnerability of sandboxes, would be a larger darknet for data collection, or even analyzing a significant portion of packets in transit.

#### *4.2. Level 2: Analysis of Network Traffic*

The challenge in network traffic analysis is to correlate patterns of connectivity (such as who sent a packet when and where) with botnet activity. As an example, the obvious fact that all bots within a single botnet can be expected to show the same behaviour in very short time frames can be used to identify botnets quite reliably with minimal information about the actual data sent out. In military applications, even with strong encryption that is practically unbreakable, knowing who communicated with whom gives a strong strategic advantage. Botnets need communication and central control from a single source – the human operator, who needs to be at one definite physical location, but may obfuscate his position through a variety of techniques.

Bots are only efficient to use when they access many different machines. Thus, a sufficiently large sensor network which records source and destination IPs and ports of network traffic would be sufficient to detect probable bots with high accuracy, although the a significant false positive rate would have to be expected. As mentioned in Section 2.2, [14] presented just such an approach, and reported a false positive rate of about 2%. This can clearly be done as a purely passive analysis without any tip-off to the botnet operators traced. A Level 1 model which tags packets corresponding to known spambot types could be used to reduce computational load and latency interference.

### *Open Issues*

A major problem is that most bots run from dynamic IP addresses, which change at relatively short intervals. A longterm record of activity based on IP addresses alone is thus useless for detecting longterm bot activity. Still, the observation that the same activity pattern is shown by two different IP addresses within a short period of time might help to tie the dynamic IP addresses together and thus generate a long-term partial model of dynamic IP address change for a certain subnet, thus making the identification of IP addresses much easier. It should be noted that internet service providers can of course trivially track the dynamic IP addresses of their own users.

At this level, data privacy issues are a major concern. The techniques used to track dynamic IP addresses may, for example, be used to track specific machines and thus specific users over long time periods. Care must be taken to handle this data in accordance with legal restrictions. This is a major problem as there are no world-wide standards available. Again, for a proactive analysis, indications from Level 1 that packets exchanged can be traced to specific spambots may be helpful to convince law enforcement agencies to actively support such techniques, and to follow up on their results.

The target population of botnet attacks is likely to be quite large, thus getting a comprehensive overview is likely to be a major problem, and the benefits are widely distributed. An alternative to a target-oriented analysis of botnet attacks, especially in the light that identification of dynamic IP addresses is trivial in this case, would be to analyze outgoing traffic from internet service providers. This has not only the advantage that ISPs could arrange the legal details with their own users in the context of appropriate contracts, and thus keep their internal networks free of bots without legal problems, but should also be considered in the context of the EU Data Retention Directive 2006/24/EC [25, 23, 24]. The benefits of a botnet-free ISP

service are clear and immediate.

#### *4.3. Level 3: Analysis of TCP/IP Traffic*

In addition to just analyzing traffic patterns without respect to contents (which is already quite effective) it is also possible to analyze patterns in content, treating, e.g., TCP streams and timing of responses as features for a learning system. This provides a lot of data not available to simple traffic analysis.

While it is theoretically feasible to build a fully passive traffic content analysis system, the practical challenges are enormous. First, only very fast learning systems are able to cope with very large traffic volumes without slowing traffic down significantly, which is undesirable as it increases the chance for detection. Downsampling the data may be helpful here, but may reduce coverage. Second, tracking connections means a high memory load only for keeping connection data until definitive connection closure, timeout or a specific event which serves as checkpoint. This is feasible for small networks, as TCP connection firewalls have proved, but cannot be done at a scale similar to what would be feasible at Level 2. Selection mechanisms might be needed to reduce the traffic which needs to go through such a system.

#### *Open Issues*

Far more so than in Level 2, data is very sensitive here. This is because content of TCP connections will include, among others, private e-mail data, social security-, bankcard- and creditcard-numbers, and personal data from literally billions of innocent people. An automated system which analyzes just a small part of this data is a large operational risk if it can be subverted, controlled or even just inspected by outside parties. While the technical difficulties could be overcome, we believe that such a system cannot be run outside government control with appropriate checks and balances. Although it would prove the most effective weapon against botnet operators, the false positive rate and potential for misuse is something to consider very deeply in this case.

One possible use case might be in law enforcement, provided that a set of potential botnet operators has been identified with simpler systems. Automatically tracking the list of potential suspects for a few days until one of them can be linked to a known botnet operator might then be an acceptable

use for such a system with manageable risks, provided data collected from innocent people is discarded as soon as possible.

## 5. Experiments

In this section, we discuss concrete realizations of the abstract three level approaches introduced before for a special type of botnets, for which most reference data could be obtained: spam botnets, whose main activity is to send out spam e-mail.

As proof-of-concept for Level 1 (single-packet analysis), we have built and deployed a purely passive darknet packet analysis system which identifies spambot type by header, contents and flags of a single incoming packet. As not even SYN packets are returned by our darknet, TCP conversations just yield one SYN packet, while ICMP and UDP packets yield a bit more information. Surprisingly, even these relatively small packets yield enough information to identify most spambot types quite reliably. We think this is because botnet operators have not yet identified these obvious vulnerabilities in their systems.

Our darknet currently contains 256 IP v4 addresses – equivalent to an /24 subnet – which are distributed among four different /26 subnets. We recorded all incoming packets since the start of the project and continue to do so as long as possible.

We also recorded all traffic for one IPv6 /48 subnet. However, no non-trivial traffic was obtained during the whole duration of the project.

As feature set we use protocol type (ICMP, TCP, UDP), TCP option types and their order, ICMP and UDP payload data (represented as 2-gram word vector (16 bit per token)), TCP and UDP destination port, and the IP Don't Fragment field. Because of the passive nature of our darknet, we could not capture full SMTP conversations – only the connection attempt was registered via TCP SYN packet. Spam mail contents were taken from a separate small spamtrap built in a different project.

In preliminary experiments we found that standard intrusion detection tools, such as snort, p0f, fl0p, honeysnap, ettercap, nmap and xprobe2 (cf. Section 2), were unsuitable for our analysis – partially because their operation mode is active (i.e., it involves sending packets which may be used to detect our presence), and partially because their patterns only fit known attacks, which the darknet effectively prevents (except for attacks purely via UDP)

as no TCP conversations can be successfully set up. We therefore had to build our own feature descriptors.

Simple single packet analysis already allows a test of effectiveness of current DNS blacklist systems under reasonable assumptions. We assume that each packet sent to the darknet has been sent from a bot-infected machine. This slightly overcounts systematic port scanning activities by the administrators of our darknet, which however happen from a single IP, can be clearly identified, and can therefore be neglected.

We found an overlap of only 2.75% with the Spamhaus XBL, which specifically targets bot-infected machines, in online lookup (i.e., within one minute of packet reception the IP is looked up). As typically dynamic IPs are reassigned on average once per hour, this is not very surprising. A faster approach to identify bots is clearly needed.

### *5.1. Analysis of Single Packets: Identifying Spambot Type*

Here, we describe just such an approach to identify bots based on a single packet which is captured by our passive darknet detector. This is an example of a Level 1 approach.

Reference data was provided by Marshal Ltd. (<http://www.marshall.com>) from their automated TRACE system. We received monthly alignments of our IP data with the entries of their database, and matched them via maximum time difference of plus/minus one hour, which is a reasonable estimate of the time a dynamic IP remains the same on average, see Fig. 7(b) in [26]<sup>17</sup>. On average only about 5% of our IPs could be assigned to an entry and thus spambot type from Marshall with these restrictions, yielding about 2000 packets which were highly likely to have been sent by a known spambot.

Another issue was the opening of a connection from a single bot to different destination IPs, reusing the same source port. The source port is therefore correlated with spambot type and must not be used. Reassigning the right spambot type in this case is trivial based on rote learning the source port - spambot mapping. So we only took one packet of each pair (source ip, source port) to prevent a positive bias. This reduced our dataset further to 308 samples of eight spambot types.

As feature representation we have used the following set:

---

<sup>17</sup>The mentioned figure shows that 97% of IPs change once an hour or less often.



Spambot Name (number)	Precision	Recall	F-measure
Srizbi (3)	0.847	0.981	0.909
Rustock (7)	1.000	0.263	0.417
Pushdo (2)	0.000	0.000	0.000
Hacktool.Spammer (6)	0.500	0.313	0.385
Mega-D (1)	0.000	0.000	0.000
<i>Unknown</i> (4)	1.000	1.000	1.000
<i>Unknown</i> (9)	1.000	0.143	0.250
<i>Unknown</i> (11)	0.947	0.973	0.960

Table 1: Precision, recall and F-measure for spambot types as determined by Marshal

- flags to indicate ICMP, TCP or UDP traffic,
- whether or not the DF flag in the IP packet is set,
- the IP destination port (ICMP, one of (20,25,1433,5900,7557,16818,23790,34357,35999,37370,37713,49110) or UNKNOWN), and
- a 2-byte gram word vector for TCP options (TCP packet), or the ICMP payload (ICMP packet), or the UDP payload (UDP packet)

The set of IP destination ports was taken from training data. We have chosen to use all information within each packet, only removing those information which cannot be meaningfully interpreted on a single packet level (e.g. sequence numbers) or which would have tainted the test data (e.g. source port).

Because of the small amount of data, we have evaluated our system by leave-one-out crossvalidation. For simplicity, because of its robustness to irrelevant attributes and in face of small training data, we have chosen a linear Support Vector Machine for training, using just the default cost parameter and no parameter optimization.<sup>18</sup> Table 1 shows the results.

Detection turned out to be very good for the bot types Srizbi, Spambot 11, and Spambot 4 with high precision and high recall; Rustock and Spambot 9 have high precision and low recall, so at least some of circulating variants can be clearly identified, and the other types are performing rather badly. The

---

<sup>18</sup>[www.cs.waikato.ac.nz/~ml/weka](http://www.cs.waikato.ac.nz/~ml/weka), weka.classifiers.functions.SMO, default parameters (-C 1)

latter might be due to the fact that no clear patterns exist (e.g., randomized payload, too many different variants, or too little training data to detect patterns – note that for Hacktool.Spammer (6), Mega-D (1) and Pushdo (2) we only have 16, 5 and 6 training examples, respectively, while Srizbi (3) has 215 examples. *Unknown* (4), (9) and (11) are distinct patterns observed by Marshall researchers which are thought to belong to distinct spambots not yet identified by antivirus researchers.

Nevertheless, we can detect bot activity and in some cases reliably assign spambot types based on our model.

### 5.2. *Static vs. Dynamic IP Addresses*

Since most bots are hosted on dynamic IP addresses, a reliable assessment of bot infection must be coupled to an incoming bot packet within a short time period. To enable a long-term lookup for non-dynamic IPs, we have also looked into recognizing static and dynamic IP addresses from reverse DNS entries. Additionally, this allows for a long-term study of botnet activity for a small subset of the bots – those which are hosted on machines with static IP addresses. This is another example of a very simple Level 1 approach, which uses just the source IP address from each packet and rDNS lookup. Presumably, at the moment this cannot be detected easily by botnet operators.<sup>19</sup>

For distinguishing dynamic and static IP addresses, we have created a training set of known dynamic and known static IP addresses. For static IP addresses, we have used all entries from `www.dnswl.org` referring to single machines (/32) with high trust level (med and hi). For dynamic IP addresses, we have used an overlap of incoming IP addresses with the PBL from `www.spamhaus.org`, which has been initialized from the now obsolete `dynalist` which was intended to capture dynamic IP addresses without reference to their spamminess. While `dynalist` was explicitly a list of dynamic IP addresses, the PBL also includes non-MTA IPs and may thus be considered slightly tainted. However, a better resource for known dynamic IPs was not available (note that [26] used the same resource). We thus obtained 8359 dynamic IPs with

---

<sup>19</sup>However, injecting IP addresses where the rDNS entry is only available on a DNS server under the botnet operators' control is not impossible and would allow detection of the application of such a classification system. This is clearly a worst case scenario which would also allow for non-user-detectable phishing, so we do not take it into account for now.

existing rDNS entry and 10387 static IPs with existing rDNS entry. 33.96% of IPs did not have a rDNS entry and were discarded. An additional test of spam-sending IPs by our local spamtrap found that 36.33% of the IPs sending spam did not have a rDNS entry, which agrees well with our sample, so bot IP addresses and spam-sending IP addresses show a similar fraction of entries with existing rDNS entry – about two thirds.

We downsampled the list of static IP addresses to the same number of entries as the list of dynamic IP addresses and created a set with equal class distributions. We removed entries present in both sets (there were none) and used only unique entries. We randomly selected half of this set to create a same-sized training and a test set. Experiments were run as is, and also with training and test set swapped, reporting the average error rate (similar to a two-fold crossvalidation).

In a first experiment using word vectors with delimiter “.” (dot), where word vectors were computed on the training set and applied unchanged to the test set, an error rate of 2.91% with a standard deviation  $\sigma = 0.174\%$  using a multinomial version of Naive Bayes<sup>20</sup> was achieved. In a second experiment using  $n$ -gram characters, we tested  $n = 1, 2, \dots, 9$ . 4-grams achieved the smallest error rate of 0.848% (with  $\sigma = 0.083\%$ ) using the same learning algorithm. A linear SVM was competitive on the word vector data with error rate 2.18% ( $\sigma = 0.16\%$ ), and on 4-grams with error rate 0.864% ( $\sigma = 0.136\%$ ). Thus, we obtained very good results in differentiating dynamic and static IPs.

When applying this to our data, we found that approximately 4.38% of the IP addresses from our darknet detector could be predicted as static IP addresses, and thus could be tracked over longer time periods, giving a more comprehensive view of botnet activity. However, the overlap of this IP set with Marshal’s reference data was too small for further analysis.

### *Open Issues*

While the current model performs very well, the fraction of clearly static IP addresses within bots is rather small at only 4.38%. There are types of dynamic IP addresses which change less often than once an hour – according to [26] another 3% of IPs might be of this type. We could thus at least double the amount of usable IPs for medium- to long-term tracking at the

---

<sup>20</sup>[www.cs.waikato.ac.nz/~ml/weka](http://www.cs.waikato.ac.nz/~ml/weka), `weka.classifiers.bayes.NaiveBayesMultinomial`

cost of a more complex model which tries to predict the approximate length of dynamic IP turnover, and do a more complex analysis with a mix of dynamic IPs with different turnover distributions. This might be an interesting venue for further research, provided sufficiently high-quality reference data can be obtained.

### 5.3. Analysis of Network Traffic

Here, we describe an example of a Level 2 approach. We have analyzed the access pattern of bots to our darknet within a six month period. For simplicity we have focussed only on the order of accesses within a short time span, chosen to minimize the risk of source IP change.

More specifically, we have extracted all accesses to our darknet from any source IP from the first occurrence until one hour later. After this time, a *cooling-off period* of six hours was chosen to prevent new accesses from the same IP in the next window. Beyond the cooling-off period, the same IP was again admitted as start for a new access sequence. The period of one hour was chosen to minimize the risk that the (probably) dynamic source IP changes, and another machine would seamlessly take over the old source IP, mixing access patterns from two distinct machines. While this means that we lose some data for IP addresses which do not change as often as we assume in the worst case, we felt that such a precaution was necessary to get sufficiently clean data.

The access pattern sequences are astonishingly varied. On the whole six months data we received 25 588 access sequences with a length of at least three accesses on at least three unique target IP addresses. 98.14% of these access patterns were unique. There were 15 627 unique source IP addresses, so 38.93% of the access patterns came from IPs which appeared more than once. None of them appeared in consecutive time windows, so the window of six hours was sufficient to isolate the accesses from one source IP address. The average length of an access pattern was  $67.95 \pm 107.53$  accesses, the average number of unique target accesses was  $47.44 \pm 69.09$ . Systematic port scanning activities are quite common – the ten most common start sequences are accesses to consecutive IP addresses within one subnet.

A sensible choice to determine distance between access pattern sequences is the edit distance, where the character set contains all our target IP addresses. We chose to use the Levenshtein distance [15]<sup>21</sup>, and visualized a

---

<sup>21</sup>This edit distance determines the cost of insertions, deletions and substitutions as a

random sample of 288 accesses (about 1% of the total number of accesses) via Sammon mapping [19, 2].<sup>22</sup> Figure 1 shows the results. A cluster which contains a significant portion of access pattern sequences is clearly visible near (0,0) on the right side. The observed spreading-out pattern reflects most likely different access lengths, since a higher length will incur a higher edit distance.




Figure 1: Sammon mapping of 288 access patterns

For a more detailed analysis, we reused the reference data by Marshal and aimed to assign a definitive spambot type to each access pattern sequence. Again, the small overlap resulted in only 82 matches by source IP and timestamp (access within plus/minus one hour of the reference) with 45 unique access pattern sequences. However, 13 of those had ambiguous spam-

---

uniform cost of 1.

<sup>22</sup>*Sammon mapping* is a non-linear multi-dimensional scaling technique which directly minimizes stress and tries to keep the same distance matrix between examples as far as possible, thus aiming for a *distance-preserving* visualization.

bot type (i.e., more than one type for the same access pattern) and had to be discarded, leaving us with 32 usable sequences. As we will see later, the unambiguous mapping of this information to spambot type may not be the most appropriate use of the access patterns. When using a Sammon mapping as in Figure 1, the majority of sequences were clustered within a small area and significantly overlapped (data not shown), so we have chosen a different visualization technique for clarity.

Figure 2 shows a distance matrix of all 32 sequences, where dark refers to small edit distance (i.e. similar access sequences) and white refers to high edit distance (i.e. dissimilar access sequences). For clarity, we have used logarithmic scale followed by histogram normalization. The columns and rows are sorted by edit distance values from small to high which facilitates the search for clusters. Columns and rows correspond to the access patterns and the diagonal corresponds to the distance of each pattern to itself, i.e. zero. The order of access patterns for rows and columns is the same. A cluster is therefore a submatrix containing one specific set of columns and the same set of rows, which is uniformly dark (e.g. rows 2,5 and 6 and columns 2,5 and 6 give a 3x3 submatrix which seems to be a cluster). Top left is column 1, row 1.




Figure 2: Complete edit distance matrix of 32 access pattern sequences with known spambot type (white = high distance, black = low distance)

By visual inspection we can discern three clusters in Figure 2: (A) second, fifth and sixth column/row; (B) first, eighth, and tenth column/row; (C) 13th, 16th, and 17th column/row. If the corresponding columns and rows from a cluster are selected, the resulting submatrix is almost completely dark (i.e. similar) relative to the background. Again, one effect of the edit distance's dependency on sequence length is that the longer sequences are to the right bottom corner of the figure and that the edit distance tends to get

higher in that direction.

A detailed analysis shows that cluster (A) contains three different sequences with a length of one which therefore all have an edit distance of one between them. This is of course rather uninteresting. However, (B) shows a variable number of accesses to a specific target IP – twice for spambot Srizbi, once for spambot Rustock. (C) also shows a variable number of accesses to a different specific target IP – twice for Srizbi, twice for Rustock. Two weaker clusters can also be discerned. (D=9,15,22) shows again a variable number of accesses to a different specific target IP, all for Srizbi. In one case, a small set of other IPs were also accessed. (E=12,18) shows again a variable number of accesses to a different specific target IP, once for Pushdo and once for Srizbi.

Immediately, we began to wonder why different spambots show extremely similar access patterns. Could this point towards a common control mechanism, e.g., a botnet consisting of multiple spambot types which were all asked to check out the same target IPs? It is highly unlikely that almost half of our sequences should be so similar if the process of choosing IP addresses to scan would be either random, or depend on spambot type. Neither hypothesis is supported by our data at present. However, we must leave a more detailed analysis for future work due to lack of data.

The six most dissimilar sequences (bottom-right of figure) cannot be fixed to a specific spambot type, either. These access sequences are long, seem to be random, and scan a significant portion of our address space ( $84.84\% \pm 18.17\%$  on average, with  $1.9 \pm 0.10$  accesses per target IP). These seem to be spambots set to a scanning mode, while the previous clusters may be reconnaissance probes intended to search for vulnerable subnets before an exhaustive scanning takes place. Note also that the access pattern sequence here is significantly longer than for an average pattern from the whole dataset.

All in all, the access pattern sequences may not tell us too much about the specific spambot type, but rather about the mode in which a botnet operates, and show tentative connections between source IPs and spambot types which may hint at a common operator, membership of the same botnet, or other factors which are not yet well understood.

#### *5.4. Analysis of TCP/IP Traffic*

Here, we describe an example of a Level 3 approach. As our darknet does not allow TCP/IP connection set-up, no continuous TCP/IP traffic could be recorded. Therefore, we aligned data from a local spamtrap addresses with Marshal's spambot data in a similar way as described earlier. This allowed

us to map a subset of SMTP conversations to a specific spambot type. We have visualized spam mail content versus spam bot type, and found some interesting patterns which may indicate that different spambots are used for sending out different kinds of spam.

This specific experiment does not result in a completely passive system, as the receipt of SMTP mail needs a TCP conversation which has to be simulated by our spamtrap, and would thus be amenable for detection attempts. However, by analyzing SMTP traffic moving through a network the same analysis could be done in a completely passive way as well. This could be done e.g. with an external firewall. Strictly speaking, the analysis of TCP/IP traffic is obviously not possible with a purely passive darknet.

By analyzing about 200 000 spam mails from our spam trap, we found that most spam data is indeed sent out in a single packet. However, not all packets sent to port #25 are spam content. In fact, we found e.g. DATA commands on top of the spam mail and QUIT afterwards, multiple spam mails in a single packet with appropriate SMTP commands between them and so on. Only an analysis of the whole TCP conversation gives enough information to distinguish the content of the mail which was sent from status and informational data transferred.<sup>23</sup>

Reference data was again provided by Marshal from their automated TRACE system. We restricted the data and obtained two selective alignments, each for a period of 24 hours. Again, we matched them via maximum time difference of one hour. We obtained 192 pairs with (received spam mail, spambot type). Because of earlier experiments reported in [20] we chose a 6-character-gram representation for each mail instead of the more widely used word vector representation. This yielded 7008 unique 6-gram tokens. For visualization we again chose Sammon mapping. We reduced the 7008-dimensional binary vector to two dimensions, and applied the known spambot types for tagging afterwards.

Figure 3 shows the results. As you can see, Rustock (7) and Srizbi (3) are responsible for the majority of incoming spam. The clusters of spam mail content are clearly distinguishable. Note that one spambot type which could not be assigned to any darknet IP (i.e. no reference data for this

---

<sup>23</sup>In fact, we found a few messages containing malware executables where TCP fragments were used to distribute the content of the message into multiple packets, presumably to confuse packet-level firewall virus scanners.






Figure 3: Sammon mapping of 192 spam mails based on content with attached spambot types.

spambot type could to matched to any of our darknet packets – this means that no verified activity of this spambot was observed by our darknet) was nevertheless found sending out spam, namely Grum (8). A few less active spambots can also be discerned.

Although spambots can potentially be used to send out any kind of spam, there seem to be clear patterns in the actual spam sent out. These might reflect usage patterns of spammers (such as preferring simpler spambots to send out cheaper spam), technical constraints (e.g., specific interfaces between some text generators for spam and some spambots, which facilitate combined usage) or some other factors, and thus could give valuable information about the source of the spam, as well as of the underlying botnets and their operators.

## 6. Outlook

Given our prototype system, and reasonable assumptions about the number of bots, it is possible to roughly compute the necessary size of our darknet to catch each bot at least once per day.

We have taken the estimate by ShadowServer<sup>24</sup> of around 600 000 bots, multiplied it by three and rounded up to 2 000 000 bots. Our present system sees 350 unique IPs each day from analyzing 30 000 packets, so we need on the order of 1,5 million unused IP addresses to see each bot on average once per day, and need to analyze 170 million packets per day (2 000 packets per second). This might be difficult to obtain because of IPv4 shortage, but we remain optimistic that a sufficient number of unused IPs will be donated via IP-over-UDP forwards.<sup>25</sup>

Concerning processing speed, our current system is reasonably efficient: 5 000 packets could be analyzed each second, so theoretically a single server would be sufficient to analyze this data at level 1. The technical limit for packet capture seems to be around 170 000 packets per second, but this does not include analysis and storage of this data is only feasible for short time periods.

Higher level analysis will need more computational power, but will also enable more complex detection and identification of botnets. The present system is not yet able to identify whole botnets. Although the spambot type

---

<sup>24</sup>[www.shadowserver.org](http://www.shadowserver.org)

<sup>25</sup>Please contact the authors if you are interested.

is one weak indicator of botnet membership, a reconstruction of botnets is not feasible with only this information. Reference data on known botnets would be very helpful for creating such a system.

Some of the patterns we mentioned could be removed quite easily. Even a complete passive analysis does not prevent the awareness of such patterns by papers such as this, and countermeasures are usually available. However, all such tasks will make the botherders' job harder and more complex, thus increasing costs and making botnets less attractive. Also, through our focus on machine learning approaches, regular validation with data obtained from other systems enables us to update our systems and note when some features are no longer working.

While technically feasible, there is a legal problem to automatically deactivate bots. Also, for obsolete operating systems which can no longer be made safe, removing one bot will yield another infection quite soon. We therefore believe that ISPs are best placed to solve this problem for their own customers. Forcing ISPs to keep their nets bot-free, e.g. by disconnecting known longtime offenders from the internet, would be a way to speed up this process.

## 7. Conclusions

We have presented a framework for passive tracking and identification of botnets based on analysis on three distinct but related levels:

- Level 1: Single Packet Analysis
- Level 2: Network Traffic Analysis
- Level 3: TCP/IP Traffic Content Analysis

We have also presented working systems at each of the levels which are, or can in case of Level 3 be made so with minimal adaptation, purely passive and thus cannot be detected by botnet operators directly.

We are interested in further development of our system towards more comprehensive tracking of bots, and seek collaboration with providers of reference data on botnet activity, owners of unused IPv4 addresses, law enforcement and other groups who take an active interest in addressing the botnet problem.

## 8. Acknowledgements

This work was partly supported by Internet Privatstiftung Austria (IPA). We also deeply acknowledge the support of Marshal Ltd., New Zealand, and especially our contact Phil Hay, who provided the reference data which was essential to our experiments.

## References

- [1] Bailey, M., Cooke, E., Jahanian, F. et al.: The Internet Motion Sensor: A Distributed Blackhole Monitoring System. In Proceedings of the 12th annual Symposium on Network and and Distributed System Security (NDSS 2005), San Diego, California.
- [2] Borg, I., Groenen, P. (1997): Modern Multidimensional Scaling. Springer-Verlag, New York.
- [3] US-CERT Vulnerability Note VU#800113, <http://www.kb.cert.org/vuls/id/800113> (visited 10th March 2009).
- [4] Collins, M. P., Shimeall, T. J., Faber, S., Janies, J., Weaver, R., De Shon, M., and Kadane, J. 2007. Using uncleanliness to predict future botnet addresses. In Proceedings of the 7th ACM SIGCOMM Conference on internet Measurement (San Diego, California, USA, October 24 - 26, 2007). IMC '07. ACM, New York, NY, 93-104. DOI <http://doi.acm.org/10.1145/1298306.1298319>.
- [5] G.V. Cormack: Email Spam Filtering: A Systematic Review, Journal on Foundations and Trends in Information Retrieval (1)4 2006, pp.335–355.
- [6] Dagon, D., Zou, C., Lee W.: Modeling Botnet Propagation Using Time Zones, 13th annual Symposium on Network and and Distributed System Security (NDSS 2006), San Diego, California, USA, 2006.
- [7] Dagon, D., Gu, G., Zou, C., et al: A Taxonomy of Botnets, unpublished report, [http://www.math.tulane.edu/tcsem/botnets/ndss\\_botax.pdf](http://www.math.tulane.edu/tcsem/botnets/ndss_botax.pdf), 2005.
- [8] D.J. Daley and J. Gani: Epidemic Modeling: An Introduction. Cambridge University Press, 1999.

- [9] Dhamankar, R., King, R.: Protocol Identification via Statistical Analysis (PISA), Black Hat 2007, Las Vegas, USA.
- [10] Gansterer W.N., Janecek A.G.K., Lechner P.: A Reliable Component-Based Architecture for E-Mail Filtering. In Proceedings of the ARES International Conference on Availability, Reliability and Security 2007, pp.43–52, 2007. IEEE Computer Society, Washington D.C.
- [11] Gansterer W.N., Pölz D.: E-Mail Classification for Phishing Defense. In Proceedings of the 31st European Conference on Information Retrieval (5478) 2009, pp.449–460.
- [12] S. Gauthronet and E. Drouard. *Unsolicited commercial communications and data protection*. Technical report, Commission of the European Communities, 2001.
- [13] Janecek A.G.K, Gansterer W.N., Kumar K.A.: Multi-Level Reputation-Based Greylisting. In Proceedings of the ARES International Conference on Availability, Reliability and Security 2008, pp.10–17. IEEE Computer Society, Washington D.C.
- [14] Karasaridis, A., Rexroad, B., Hoeflin, D.: Wide-scale Botnet Detection and Characterization, First Workshop on Hot Topics in Understanding Botnets (HotBots '07), Cambridge, MA, USA, 2007.
- [15] V. I. Levenshtein, Binary codes capable of correcting deletions, insertions, and reversals. Soviet Physics Doklady 10 (1966):707–710.
- [16] National Infrastructure Security Co-ordination Centre (NISCC). Botnets - the threat to the Critical National Infrastructure, NISCC Monthly Bulletin, London, October 17 2005.
- [17] Ramachandran, A., Feamster, N., Dagon, D.: Revealing Botnet Membership using DNSBL Counter-Intelligence, In Proceedings of the 2nd Workshop on Steps to Reducing Unwanted Traffic on the Internet (SRUTI 2006), San Jose, California, USA, 2006.
- [18] Register of Known Spam Operations (ROKSO) database, operated by the Spamhaus Project. <http://www.spamhaus.org/rokso/index.lasso> (visited 10th March 2009).

- [19] Sammon, J.W. (1969): A nonlinear mapping for data structure analysis. *IEEE Transactions on Computers*, C-18:401-409, 1969.
- [20] Seewald A.K., Kleedorfer F.: An Approximation of the String Subsequence Kernel for Practical SVM Classification and Redundancy Clustering, *Journal for Advances in Data Analysis and Classification*, Vol. 1, Number 3 / December 2007, pp. 221-239, DOI: 10.1007/s11634-007-0012-1.
- [21] Seewald A.K.: An Evaluation of Naive Bayes Variants in Content-Based Learning for Spam Filtering, *Intelligent Data Analysis* 11(5), pp. 497-524, 2007.
- [22] Seewald A.K.: Towards Automating Malware Classification and Characterization. In *Konferenzband der 4. Jahrestagung des Fachbereichs Sicherheit der Gesellschaft für Informatik (German-language proceedings)*, Saarbrücken, April 2008, pp. 291-302.
- [23] Stampfel G., Gansterer W.N., Ilger M.: *Data Retention – The EU Directive 2006/24/EC from a Technological Perspective*, Medien und Recht, 2008. ISBN 978-3-900741-53-2.
- [24] Stark K., Stampfel G., Gansterer W.N.: A distributed data warehouse for e-mail data retention. In *Proceedings of the 4th International Conference on e-Government (ICEG08) 2008*, pp.413–422. Academic Publishing Limited, ISBN 978-1-9066938-20-7.
- [25] Stampfel G., Gansterer W.N., Ilger M.: Implications of the EU Data Retention Directive 2006/24/EC. In *Proceedings of Sicherheit P-128 (2008)*, pp. 45–58. Springer Lecture Notes in Informatics.
- [26] Y. Xie, F. Yu, K. Achan et al.: How dynamic are IP Adresses, *Proceedings of SIGCOMM '07*, August 27-31, 2007, Kyoto, Japan.