Internet Bad Neighborhoods Temporal Behavior

Giovane C. M. Moura*, Ramin Sadre†, and Aiko Pras‡
* Delft University of Technology
Email: g.c.moreiramoura@tudelft.nl
† Aalborg University
Email: rsadre@cs.aau.dk
‡University of Twente
Email: a.pras@utwente.nl

***Abstract*—Malicious hosts tend to be concentrated in certain areas of the IP addressing space, forming the so-called Bad Neighborhoods. Knowledge about this concentration is valuable in *predicting attacks* from unseen IP addresses. This observation has been employed in previous works to filter out spam. In this paper, we focus on the temporal behavior of bad neighborhoods. The goal is to determine if bad neighborhoods strike multiple times over a certain period of time, and if so, when do the attacks occur. Among other findings, we show that even though bad neighborhoods do not exhibit a favorite combination of days to carry out attacks, 85% of the recurrent bad neighborhoods do carry out a second attack within the first 5 days from the first attack. These and the other findings here presented lead to several considerations on how attack prediction models can be more effective *i.e.*, generating both predictive and short neighborhood blacklists.**

I. Introduction

The Internet Bad Neighborhood concept emerged [1] from the observation that malicious IP addresses are not evenly distributed over the IP addressing space [2]–[4]. Bad Neighborhoods can be seen as subnetworks having higher concentrations of malicious IP addresses than the measured average [1], [5]–[8]. For example, we found in a previous study that 50% of spamming IP addresses on the Internet can be traced back to only 20 Autonomous Systems (ASes) [8].

Knowledge about the concentration of malicious hosts is valuable in *predicting attacks* from unseen IP addresses. Traditionally, blacklisting has been the approach of choice to predict attacks, in which sources involved in previous attacks are flagged as malicious [9]–[12]. The bad neighborhood concept furthers traditional blacklisting and improves attack prediction by *extending the reputation* of malicious IP addresses to their immediate neighbors – i.e., by blacklisting their neighboring IP addresses, which are, in turn, more likely to carry out attacks due to the typical concentration. As shown in recent works [1], [7], this idea has proved to be effective in predicting and protecting network from attacks.

As for traditional blacklists, bad neighborhoods blacklists should be both predictive and short [13], that is, be able to forecast attacks from most bad neighborhoods (and having low false positives) and list only significant bad neighborhoods, especially if they are employed to filter out traffic in real-time.

To achieve these goals, several aspects of bad neighborhoods have been scrutinized. In [1], van Wanrooij and Pras developed a bad neighborhood-based spam filter, while in [5] we have put spamming bad neighborhoods under scrutiny. The issue of the bad neighborhood size was investigated in [6], while the performance of third-party spam blacklists was evaluated in [7]. Ultimately led to the Ph.D. dissertation of one of the authors [8], [14].

In this paper, we focus on another facet of bad neighborhoods (BadHoods hereafter). The goal is to reveal their *temporal behavior*, i.e., to determine if they strike multiple times over a certain period of time, and if so, when do the attacks occur. By scrutinizing their temporal behavior, a network administrator can determine *how often* bad neighborhood blacklists should be updated in order to better protect targets. Most importantly, any observed temporal pattern can be *employed in the design of attack prediction models* to counterattack attacks (or avoid damage from the attacks). Moreover, it may allow to generate shorter blacklists, since entries not likely to attack again may be removed.

With this in mind, we address two research questions:

1) *Given a certain monitoring period, in how many days a bad neighborhood is observed carrying out attacks? And on what days do these attacks occur?* The answer to this question will show if a network administrator can expect bad neighborhoods to attack again and when that is expected to happen.
2) *Given a single monitoring day, how many bad neighborhoods that carried out attacks in this day can be traced back to previous days (recurrent)? And how long does it take for most of them to strike again?* The answer can be used to develop models that predict attacks from bad neighborhoods, based on historical past.

To carry out this investigation, we have obtained real world data sets listing IP addresses involved in attacks from different applications (we refer to those data sets as raw blacklists). Then, we have aggregated [6] each of these blacklists into bad neighborhood blacklists, which are obtained by aggregating the raw blacklist into a /24 prefixes blacklist. Finally, we then scrutinized their temporal behavior.

The rest of this paper is divided as follows. In Section II, we cover the data sets used in this paper. Next, in Section III, we address the first research question, while in Section IV we address the second one. After that, we cover the related work in Section V and the conclusions are presented in Section VI.

II. Evaluated Datasets

In this section we describe the raw blacklists used to generate the bad neighborhood blacklists. Then, we show how the number of BadHoods vary on a daily basis for BadHood blacklist.

A. Choosing the Raw Blacklists

We have obtained data from three raw blacklist sources. Many data sources can be employed when investigating Internet BadHoods. In this work, we have chosen a subset of these satisfying the following two criteria: (i) the blacklist has been previously analyzed by both scientific and Internet security communities; and (ii) the organization hosting the list provides bulk-access to the blacklist data, to ensure we have a complete view of the malicious IP addresses. These criteria led us to choose two raw sources:

- Composite Block List (CBL): CBL is operated by "a group of computer security, spam and virus professionals, dedicated to developing and maintaining an anti-spam and anti-virus DNSBL (DNS blacklist) of the highest possible quality and reliability, that large organizations can use with confidence" [10]. It lists /32 IP addresses that have reached their spamtraps. The number of traps and their location is not disclosed, but it is distributed over different networks and countries. CBL has been employed in a number of studies, including [5], [6], [15]–[17].
- DShield.org (Dshield) [12]: DShield is a community shared firewall log system. Volunteers submit their firewall logs from more than 600 contributors, which encompasses more than "500,000 IP addresses (firewalls) in over 50 countries" [18]. It is maintained by the SANS Institute [19], and contains security logs from many applications. As for Spam blacklists, the blacklists IP addresses are aggregated from many different sources. The DShield dataset has been investigated by the research community in papers including [13], [20]. We have focused on the two most commonly attacked applications:
  - TCP 445 (T-445): Microsoft-DS Active Directory and/or Windows shares.
  - TCP 3339 (T-3389): Microsoft Terminal Server (RDP).

In addition, we have obtained a raw spam blacklist from the Electrical Engineering, Mathematics, and Computer Science Faculty of the University of Twente (UT/EWI). CBL and Dshield datasets, as discussed in [7] and in Chapter 7 in [8], can be seen as public blacklists, which are usually generated using a large number of honeypots/traps distributed over many networks, which increases the probability of blacklisting more sources. UT/EWI data set, on the other hand, was generated based on the incoming traffic to two mail server. However, there are some disadvantages of employing blacklists from public sources. The main one is the fact that the dependability of the security solution designed to protect the target is put at stake, since it relies on the availability of third-party no-warranty freely distributed blacklists. Such blacklist sources might fail for various reasons – a disruption in the service can occur (e.g, PSBL users experienced a 4 day period outage in November 2011 due to bad weather conditions [11]), the public source might become victim of DDoS attacks, or change their business model and charge for access, or even stop providing blacklists overnight.

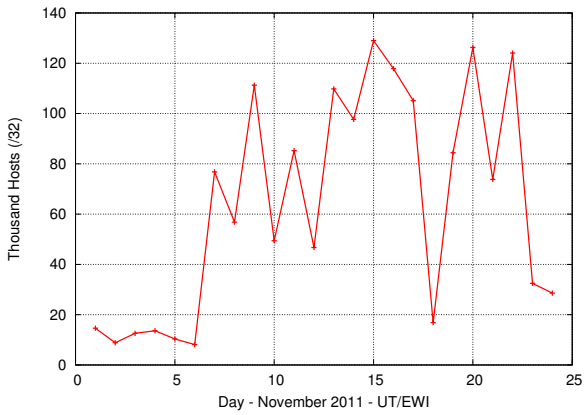

Fig. 1. Daily Spammers observed by UT/EWI – November 2011

In order to have fair comparison conditions, we have considered datasets obtained for the same monitoring periods. We have therefore considered two monitoring periods: April 2010 (19th to 26th, 8 days) and November 2011 (11th to 17th, 7 days).

After obtaining the data sets, we have, for each day and data set, generated a /24 BadHood blacklist, i.e, aggregated it under /24 prefixes [21] (since it is the smallest "routable" prefix on the Internet and incurs smaller aggregation errors [6]). These BadHood blacklists were then employed to answer the research questions presented in the introduction.

B. Daily Number of Bad Neighborhoods

We have several reasons to expect that the BadHoods distribution over different days is far from being static. The main one is a consequence of the behavior that individual hosts (/32) exhibit, trying to be as stealthy as possible (e.g., spamming only once a server and not coming back). For example, Figure 1 shows the daily number of unique spammers (/32 hosts) for UT/EWI throughout November 2011. As can be seen the values range from less than 20K to more than 120K individual hosts per day, over a period of 24 days. Another reason for expecting that the BadHood distribution changes over time is that DNS Blacklists [9], such as CBL [10], which contain many malicious /32 IP addresses, have to be constantly updated in order to keep up with the dynamics of individual hosts and be effective in the mail filtering.

Taking these into account, Table I presents the daily number of BadHoods, for each individual dataset. As we expected, for all data sets, the number of BadHoods changes on a daily basis. In addition to that, we observe that:

- The variation on the number of daily BadHoods is more significant for UT/EWI and DShield data sets (T-445 and T-3389) than for CBL (`Max Variation` row, which is the ratio between the day having most entries divided by the day having the least entries, or $100 \times Max/Min$.

April 2010				
Day/Dataset	**CBL**	**UT/EWI**	**T-445**	**T-3389**
1st Day	955,036	66,759	141,527	752
2nd Day	958,258	58,344	146,051	817
3rd Day	954,019	61,804	143,379	731
4th Day	954,522	60,045	142,531	759
5th Day	949,167	46,892	142,105	834
6th Day	957,583	48,828	142,422	832
7th Day	961,573	45,351	138,426	773
8th Day	956,410	59,739	141,512	895
Max. Variation:	∼ 1%	∼ 47%	∼5%	∼22%
November 2011				
Day/Dataset	**CBL**	**UT/EWI**	**T-445**	**T-3389**
1st Day	812,217	56,030	79,258	818
2nd Day	809,268	32,612	77,286	25,228
3rd Day	798,345	62,769	76,210	25,331
4th Day	792,098	64,452	77,003	33,319
5th Day	795,763	73,615	78,004	31,065
6th Day	803,126	69,760	79,259	19,742
7th Day	812,598	62,903	77,033	21,331
Max. Variation:	∼ 2%	∼125%	∼4%	∼4,000%

TABLE I. NUMBER OF BADHOODS/DAY

- Abrupt variations can occur, as can be seen between 1st and 2nd days of T-3389 (November 2011).

1) Variation between the Datasets : A reason for the fact that variations were proportionally more significant for UT/EWI and DShield datasets than CBL has to do with the way each original blacklist is generated. UT/EWI and DShield datasets are generated based only on attacks observed on a *single day*, that is, all /32 entries they list correspond to, at least, one attack observed on the very day.

CBL, on the other hand, may list entries on a random day that were not observed in the very day. CBL and other public spam blacklist sources employ large spam traps infrastructures and after blacklisting a certain IP address, they may keep it on the list for many forthcoming days, even though no more spam has been observed from that particular IP. In fact, the CBL's de-list policy is "manual" – that is, the responsible network administrator for the blacklisted IP should go to CBL web site and manually remove the address from the list [22], otherwise it might remain blacklisted for many days, as can be seen for the IP address `221.0.141.106`, as shown in Listing 1, which was obtained on CBL's website on October 2nd, 2012. As can be seen, this IP address has been kept in CBL for more than 12 days since it was last observing attacking CBL traps. This, in turn, confirms the fact that entries are expected to remain blacklisted over multiple days even if there was no malicious activity (spam, in this case) observed from its originating sources. Since CBL's aim is to provide better protection against spam in relation to future attacks, keeping it on the blacklists makes sense in case recurrent spammers are frequent.

2) Abrupt Variation in the Number of BadHoods: As shown in Table I, the number of active BadHoods changed significantly over a single day for November 2011 Dshield T-3389 data sets (from 818 to 25,228, as shown in Table I). Since we do not have the full traces, it is not easy to point the exact causes. However, we assume that it relates to the existence of new exploits or seasonal attack behaviors for the particular TCP port.

Such behavior is not exclusive for the Dshield data set.

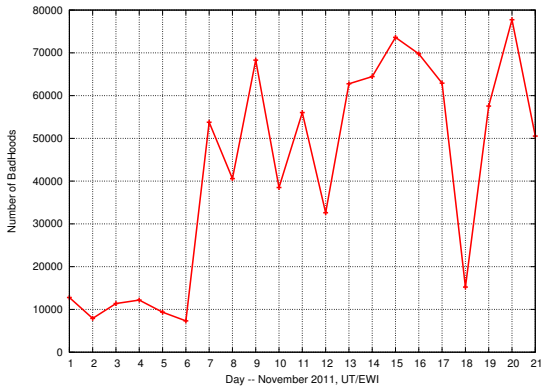

Fig. 2. Number of BadHoods - UT/EWI

Figure 2 shows the number of BadHoods for the UT/EWI dataset, for an extended period of time (November 1st -24th, 2011). As can be seen, also for Spam BadHoods, we can observe significant variations from one day to the next (from less than 10K to more than 50K, between November 6th and November 7th).

```
IP Address 221.0.141.106 is listed in the CBL. It appears to
    be infected with a spam sending trojan or proxy.

It was last detected at 2012-09-20 08:00 GMT (+/- 30 minutes
    ), approximately 12 days, 7 hours, 30 minutes ago.

This IP is infected (or NATting for a computer that is
    infected) with a spambot we have not yet been able to
    identify. For the time being we refer to it as the
    unknown66 spambot.

This IP is infected (or NATting for a computer that is
    infected) with a spam-sending infection. In other words
    , it's participating in a botnet. If you simply remove
    the listing without ensuring that the infection is
    removed (or the NAT secured), it will probably relist
    again.
```

Listing 1. CBL Lookup Result – October 2nd, 2012

III. BAD NEIGHBORHOODS ATTACK STRATEGY

In this section, we address the first research question raised in the introduction. We start by determining the total number of days that BadHoods are active.

Figure 3 shows the distribution of the BadHoods taking into account the number of days they are active (not necessarily consecutive days), that is, carrying out attacks. As can be observed, a significant part of the BadHoods are likely to attack again (for CBL, this is most prominent due to its de-listing policy). This is good news, implying that using historical data is useful to predict attacks for a new day.

On the other hand, some datasets have presented a significant percentage of BadHoods that attack a single day out of monitored days (almost 50% for T-3389 data sets). BadHoods that are active for only one day pose a challenge for BadHoods-based security systems, since such BadHoods attack on a single day and are not further observed within a short term period.

A. The Bad Neighborhood Occurrence Score

The results presented in the previous section show the number of days a BadHood is active for the monitoring data

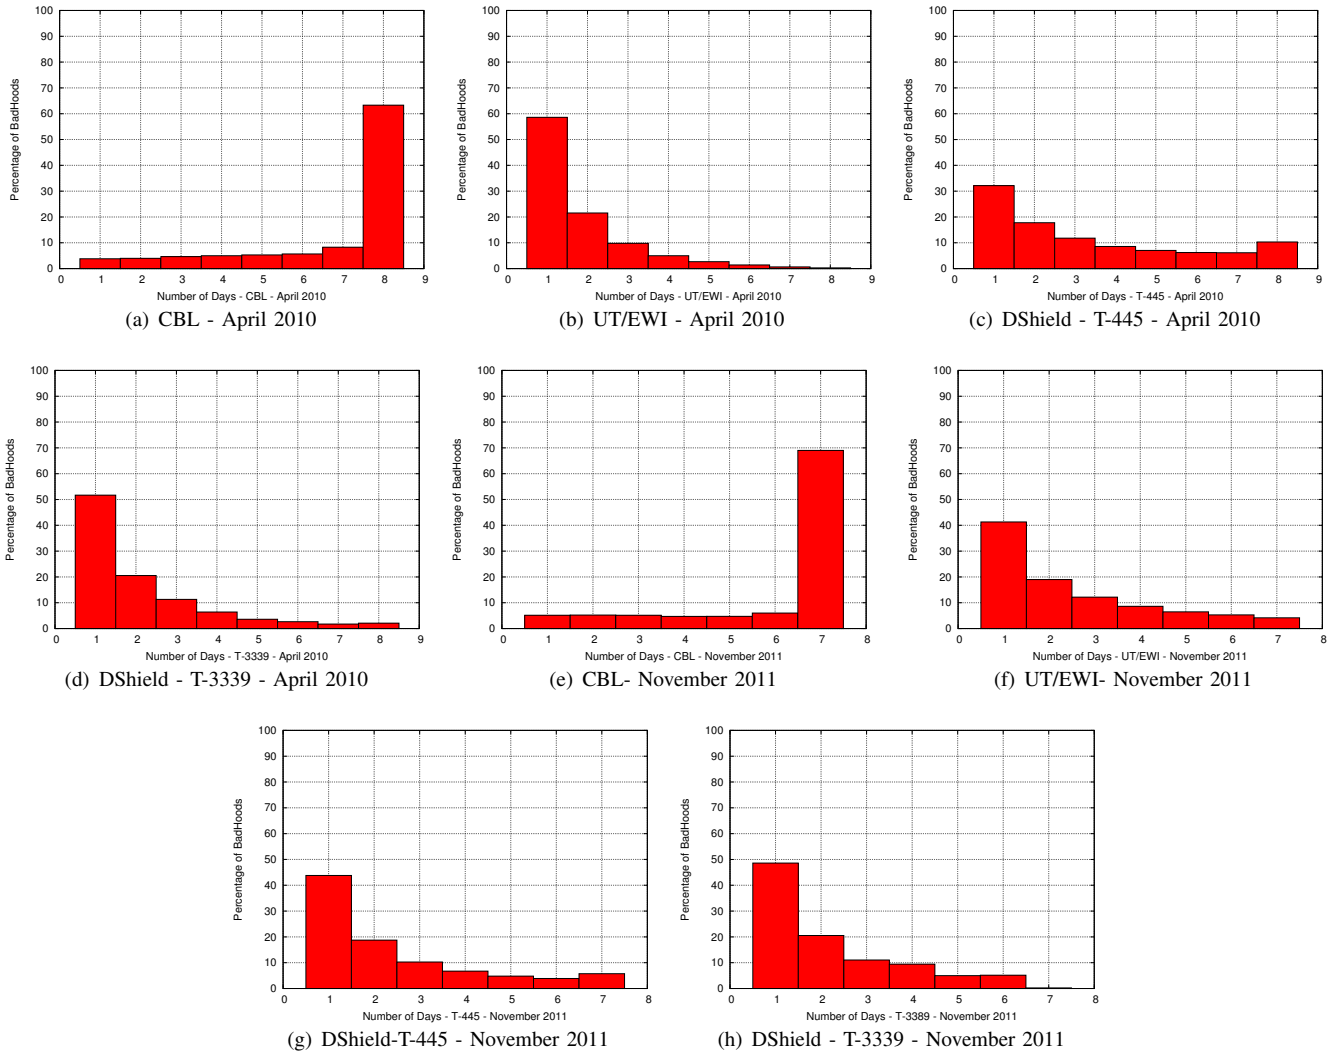

(a) CBL - April 2010 (b) UT/EWI - April 2010 (c) DShield - T-445 - April 2010

(d) DShield - T-3339 - April 2010 (e) CBL- November 2011 (f) UT/EWI- November 2011

(g) DShield-T-445 - November 2011 (h) DShield - T-3339 - November 2011

Fig. 3. BadHoods – Number of Days Active

sets. However, it does not show *which days* of the monitoring period are chosen by the BadHoods. For example, 2 days could be a combination of *any* 2 random days within the monitoring period.

In order to be able to tell what days (or combination of) the BadHoods are active, we propose in this section the *occurrence score*. For a given data set having n days of data, we define, for each /24 BadHood (B^{24}), an occurrence score as follows:

$$occurrence(B^{24}) = \sum_{i=1}^{n} s(B^{24}, i) \times 2^i \qquad (1)$$

In this equation, i refers to a monitoring day, and may vary from 1 (first day of the monitoring day, not necessarily the day of the month) until n, the last day in the observed data set. $s(B^{24}, i)$ represents if the BadHood is active on day i, therefore its value is set to 1 if the BadHood active, otherwise 0. The final occurrence score is the sum of 2^i for each n day B^{24} is active. In the end, the final number is a single integer number that can be decomposed to reveal which days from the n days a certain BadHoods carried out attacks.

To better illustrate how the occurrence score is calculated and decomposed, consider the April 2010 data set from UT/EWI. Table II shows an excerpt of the final BadHood score file that was generated after scoring BadHoods for the monitoring period. For each BadHood, an occurrence score is provided, calculated using 1. As shown in this table, a score of 96 can be decomposed into two terms. The power of each of them (5 and 6) represents the days the BadHood was active: 5th and 6th of the monitoring period. These, in turn, represent April 23rd and 24th.

BadHood	Score	Decomposed Terms	Days Active
93.105.233/24	96	2^5+2^6	5th and 6th
94.66.155/24	16	2^4	4th
94.66.154/24	12	2^2+2^3	2nd and 3rd
93.105.231/24	228	$2^2+2^5+2^6+2^7$	2nd, 5th, 6th, and 7th
71.223.131/24	8	2^3	3rd
94.66.153/24	80	2^4+2^6	4th and 6th

TABLE II. OCCURRENCE SCORES FOR UT/EWI BADHOODS (APRIL 2010)

An important property of the occurrence score is that any

score $2^i < x < 2^{i+1}$ implies that the BadHood is active on the i-th day plus any previous day(s) ($i' < i$), but never on any days $> i$. For example, a score of 32 means that a BadHood is active on the 5th day. However, there is no other combination of days that would yield to a score > 32 and < 64 that would not include the 5th day. For example, if a BadHood is active on days 1–4, it's final score is 30, which is smaller than the occurrence of a single day alone (5th day = 32).

1) Occurrence Scores Distribution and CDF: Figure 4 shows both the distribution and the cumulative distribution function (CDF) of the occurrence scores (left and right columns, respectively), for the April 2010 datasets, while Figure 5 shows the results for the November 2011 data sets.

Analyzing the figures, we can observe that, with the exception of CBL, *no occurrence score is significantly more prominent than the others*. In fact, with the exception of CBL, all the other data sets observe small spikes on scores equal to 2^i, which are BadHoods that have only attacked on a single day. CBL, on the other hand, presents a significant spike on score 254 for Figure 4(a) and 510 for Figure 5(a), (scores that represents all previous days summed up), as expected from Figure 3(a) and 3(e), which is due to their manual de-listing policy (as discussed in Section II-B1).

That leads to the conclusion that, except for CBL, *there is no day or a combination of days that is significantly more recurrent than others, even for different applications.* Therefore, our results show that a network administrator should not expect any pattern or regularity in terms of which days BadHood chose to attack – which makes the task of predicting attacks more complex.

IV. Tracing Back BadHoods: Time Since Last Attack

From the previous results, we conclude that there is no particular combination of days that BadHoods choose to carry out their attacks. Therefore, in this section, we focus on a single day of the monitoring period instead of all the monitored days. We single out the *last day* and scrutinize each observed BadHood, in order to determine *if* they can be traced back to any previous days. After that, we determine *how many days* have passed since the last attack.

To do that, we have carried out a three-step approach. First, we obtain all the /24 BadHoods of the last day of each data set (as covered in Section II). Then, for each of them, look it up on the final occurrence score file generated for the whole monitoring period (as shown in Section III-A). Those BadHoods that have been observed carrying out attacks in the last day in combination with any of the other previous days (in any combination) are filtered. Mathematically, this means that we have only considered BadHoods having an occurrence score larger than the threshold $\varepsilon > 2^i$, in which i is the number of monitoring days for each data set. For the April data sets, ε is equal to 256 and 128 for November datasets.

In our case, we are interested in the last i' day that a BadHood X^{24} has been active (the day before the singled out day). To illustrate this, consider that a certain BadHood from UT/EWI (April data set) has a score of 262. By decomposing this number into powers of two, it reveals that this BadHood has been active in days 8, 2, and 1 ($262 = 2^8 + 2^2 + 2^1$). From the days it was active, we compute the *difference between the last day (8) and the day right before it (2)*, which result in 6 days between attacks.

Table III shows the number of BadHoods on each data set, and the percentage of the recurrent ones ($OcurrenceScore > \varepsilon$). We can observe that for all the data sets (we have disregarded CBL, due to its removal policy), *the majority of BadHoods that attack a target have also been observed in at least one of the previous days.* For the April 2010 data sets, that means that 65-89% of all BadHoods observed in the last day are likely to have been observed on all previous days (7 days), while for the November 2011 datasets, 73-80% of BadHoods observed on the last day are likely to also have been active on all previous days (6 days).

April 2010			
	UT/EWI	**DShield-T445**	**DShield-T3389**
BadHoods	59,739	141,512	895
Recurrent	39,237 (65.68%)	126,057 (89.07%)	602 (67.26%)
November 2011			
	UT/EWI	**DShield-T445**	**DShield-T3389**
BadHoods	62,903	77,033	21,331
Recurrent	51,745 (80.97%)	61,159 (79.39%)	15,732 (73.75%)

TABLE III. Total and Recurrent BadHoods in Relation to the Last Day

Then, the next step was to determine when each of the recurrent BadHoods was last observed. Figure 6 shows these results as a cumulative distribution function (CDF). As can be seen, for all the data sets, the majority of the recurrent BadHoods return within 5 days (>85%), which is valuable information to determine how many days should be considered to build BadHood attack prediction models.

V. Related Work

This paper is related to research works that fall into two categories: concentration of malicious IP addresses and attack prediction models.

In regards to concentration of malicious IPs, we have seen several research works. For example, Ramachandran and Feamster addressed the network level behavior of spammers [2], concluding that most of spam comes from a few concentrated part of IP address space. Collins *et. al* [3], in turn, defined the concept of uncleanliness, which "works as an indicator for how likely the network is to contain compromised hosts". DNS blacklists [9], such as CBL [10], also suggest the same concentration. Van Eeten *et al.* [23], analyzing billions of spam mesages from the period of 2005 and 2008, also confirm similar results: they found that 50 ISPs account for half of worldwide spam sources.

Based on those works, the Bad Neighborhood concept was introduced in [1]. The authors developed a mail filter that employed Spam BadHoods to judge whether a message was spam. They have proposed a combination of several rules to classify a message based on the message's source IP address and URLs in the contents. In relation to this work, the authors have employed different spam blacklists but have only consider a period of one day in their monitoring period. However, as shown by the CBL case, monitoring a single day may reflect

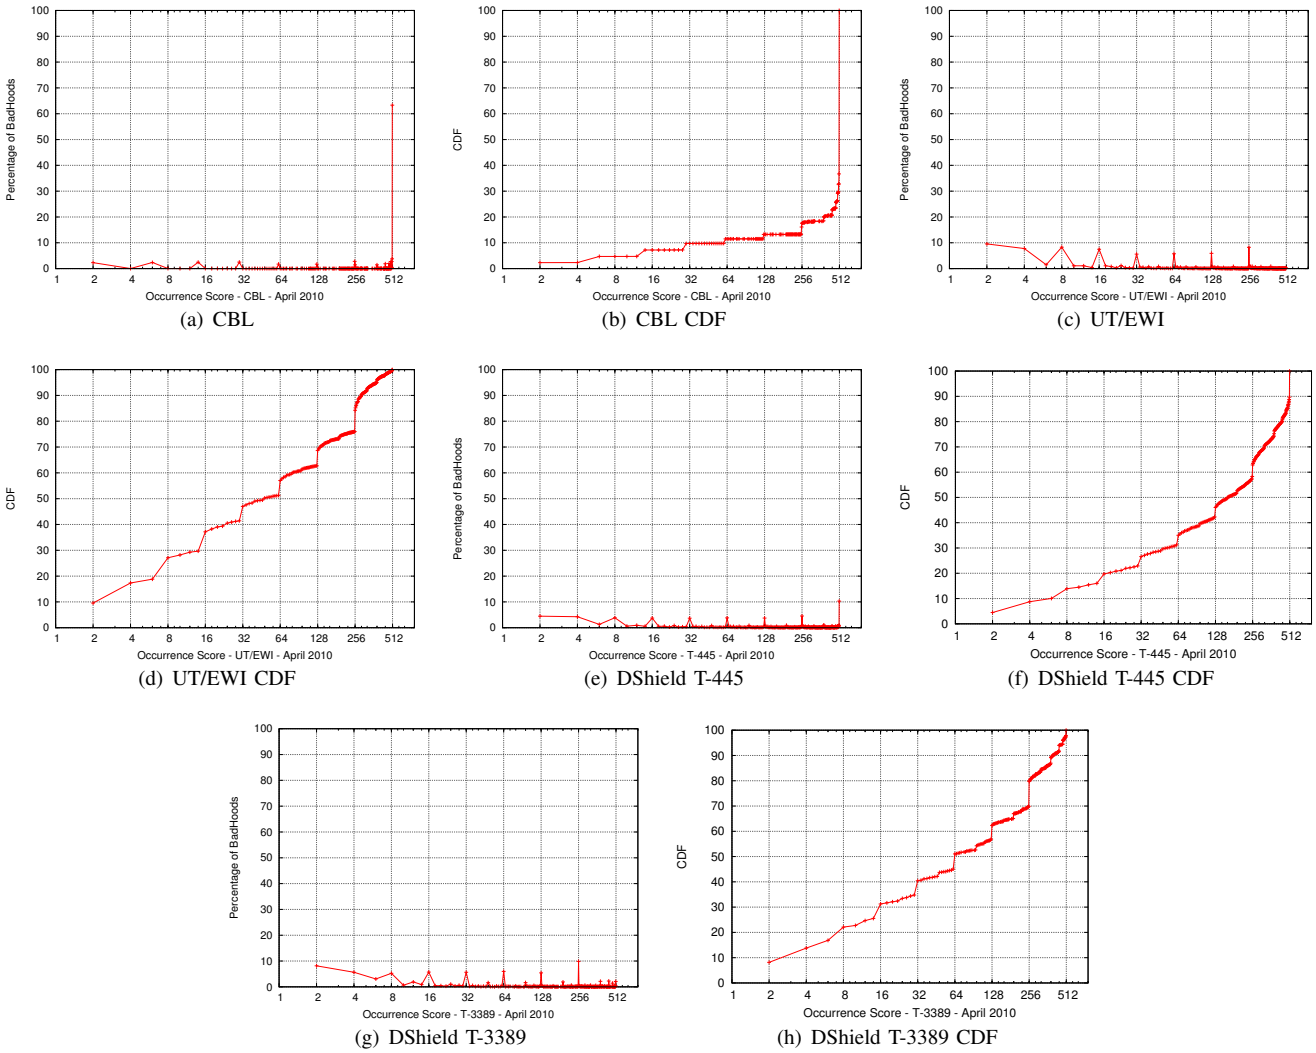

(a) CBL (b) CBL CDF (c) UT/EWI

(d) UT/EWI CDF (e) DShield T-445 (f) DShield T-445 CDF

(g) DShield T-3389 (h) DShield T-3389 CDF

Fig. 4. Occurrence Scores – April 2010

what is observed in several previous days, depending on the blacklist de-listing policy.

Taking this previous studies into account, we have investigated in [5] the specifics of spamming BadHoods. Our monitoring period was one week. We have proposed four definitions for Spamming BadHoods, each of them addressing a particular part of the "Spam picture". We have found that botnets (and individual bots) are responsible for most of Spam; however we cannot neglect the impact of high volume spamming BadHoods – that is, BadHoods having few spamming hosts but that send large amounts of spam. Following this, we proposed and evaluated two IP-based techniques to aggregate malicious IPs into network prefixes other than /24 (from /24 to /8) [6]. For this work, we used datasets for a single day period. We have found that BadHood can be viably aggregated into different BadHood sizes; however, the larger the BadHood, the larger the aggregation error. Finally, in [7], we have evaluated third-party raw blacklists to filter out spam messages. For this filter, we have generated BadHood blacklists using a single day. These findings then led to a Ph.D. dissertation [8].

Even though our work does not cover attack prediction models, it provides the temporal analysis for such models. In [13], Soldo *et. al* employed a recommendation system to predict /24 prefixes that were likely to attack "neighboring" targets (or victim networks). They evaluated the D-Shield data set [12] and employed a neighborhood model (popular approach in recommendation systems) to predict attack sources by "trusting similar peers". Differently from theirs, our BadHood definition is not a recommendation system technique and it is defined in terms of neighboring sources of attacks – and not on neighboring targets. For their model, the authors considered a five day monitoring period, since they also found that the majority of recurrent /24 prefixes attack within this time frame. In comparison to our work, the authors have only evaluated the DShield data set and they have not taken into account the exploited applications. We have evaluated also UT/EWI and CBL raw blacklists, in addition to DShield, and have evaluated separatedly the two most exploited applications from DShield.

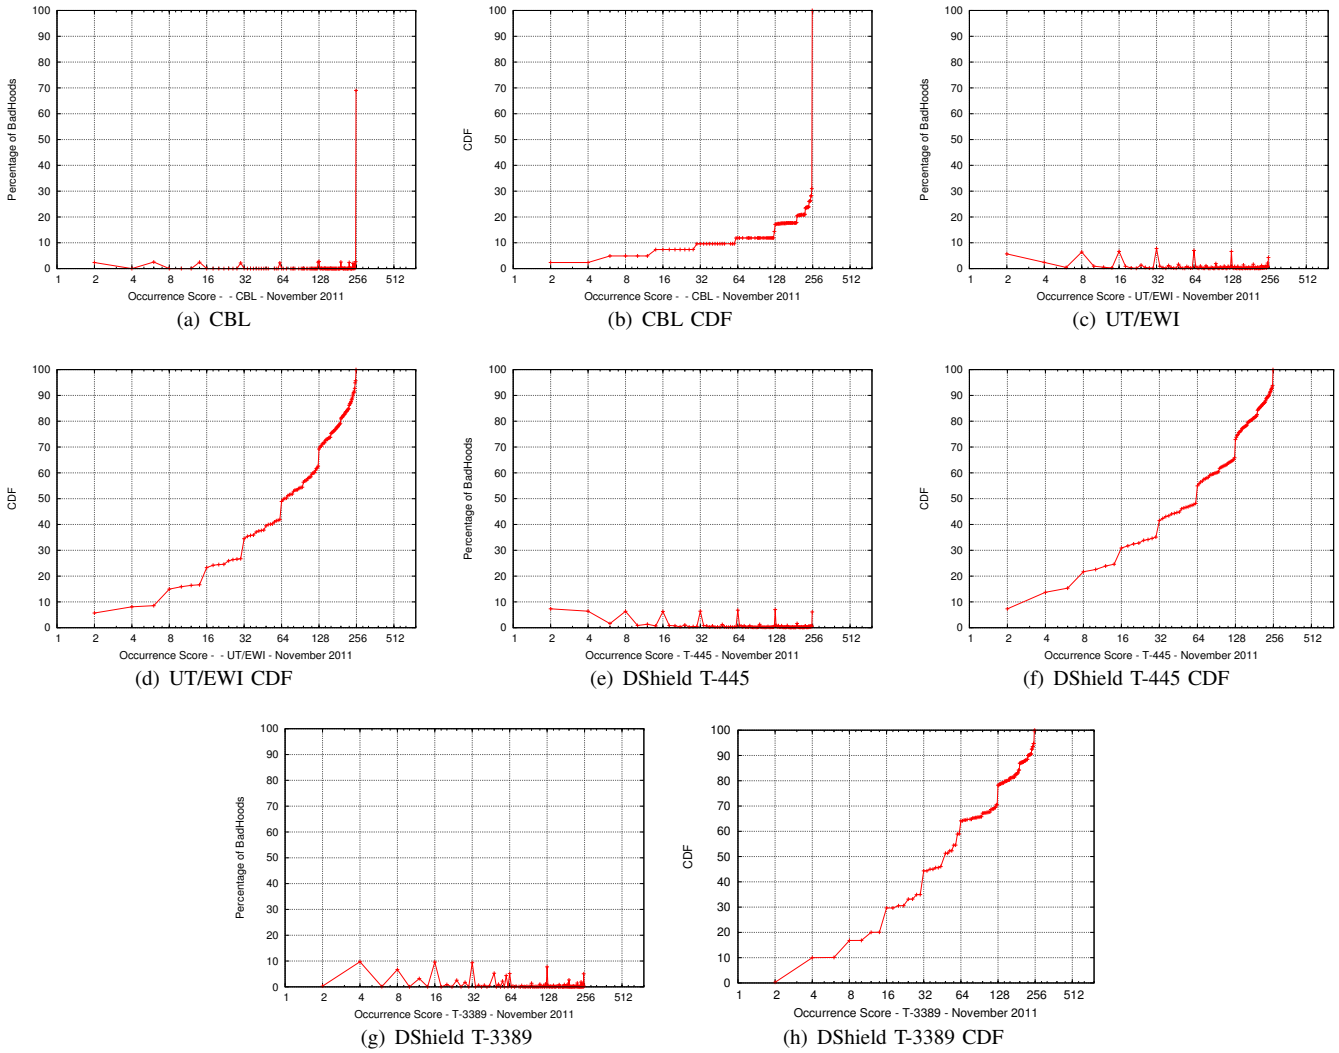

(a) CBL (b) CBL CDF (c) UT/EWI

(d) UT/EWI CDF (e) DShield T-445 (f) DShield T-445 CDF

(g) DShield T-3389 (h) DShield T-3389 CDF

Fig. 5. Occurrence Scores – November 2011

VI. Conclusions

In this paper we have scrutinized Bad Neighborhoods' temporal attacking behavior. We have raised two research questions and investigated them using real world data sets for three different misused applications: e-mail, Windows shares, and Windows terminal servers.

For the first question, (*"Given a certain monitoring period, in how many days a BadHood is observed carrying out attacks? And on what days do those attacks occur?"*), we found that a significant part of BadHoods (between 40% and 95%, depending on the data set) are likely to attack a same target on multiple days (recurrent BadHoods). This confirms that it is useful to use historical past of BadHoods to predict new attack sources. However, we also found that there is no particular combination of days that BadHoods chose to attack a target, which poses an extra challenge when predicting attacks.

For the second question ("*Given a single monitoring day, how many BadHoods that carried out attacks in day can be traced back to previous days? And how long does it take for most of them to strike again?*"), we found that the majority of the BadHoods (85%) that attack a particular target are likely to attack it again within a 5 day period, for the three applications evaluated from two different raw blacklists.

The findings presented in this paper provide information that can be used in prediction models. We showed in Section II-B that daily number of BadHoods attacking a target varies with the data source and misused application (also investigated in Chapter 7 in [8]); therefore an attack prediction model should leverage this and one could not expect an "one-size-fits-all" temporal prediction model. Also, the usefulness of the recent historical past has been proved first research question since up to 95% of BadHoods are likely to attack more than one day, which justifies its use to predict future attacks. Finally, when deciding the number to use from the historical past to design a prediction model, our results have shown that 5 days covers the majority of the recurrent BadHoods (85%).

As future work, we intend to combine the findings presented in this paper with findings showed in previous works to design attack prediction models, similarly to [13]. For example, we have shown that a 5-day long training dataset allows to predict 85% of new attacking sources, and we have

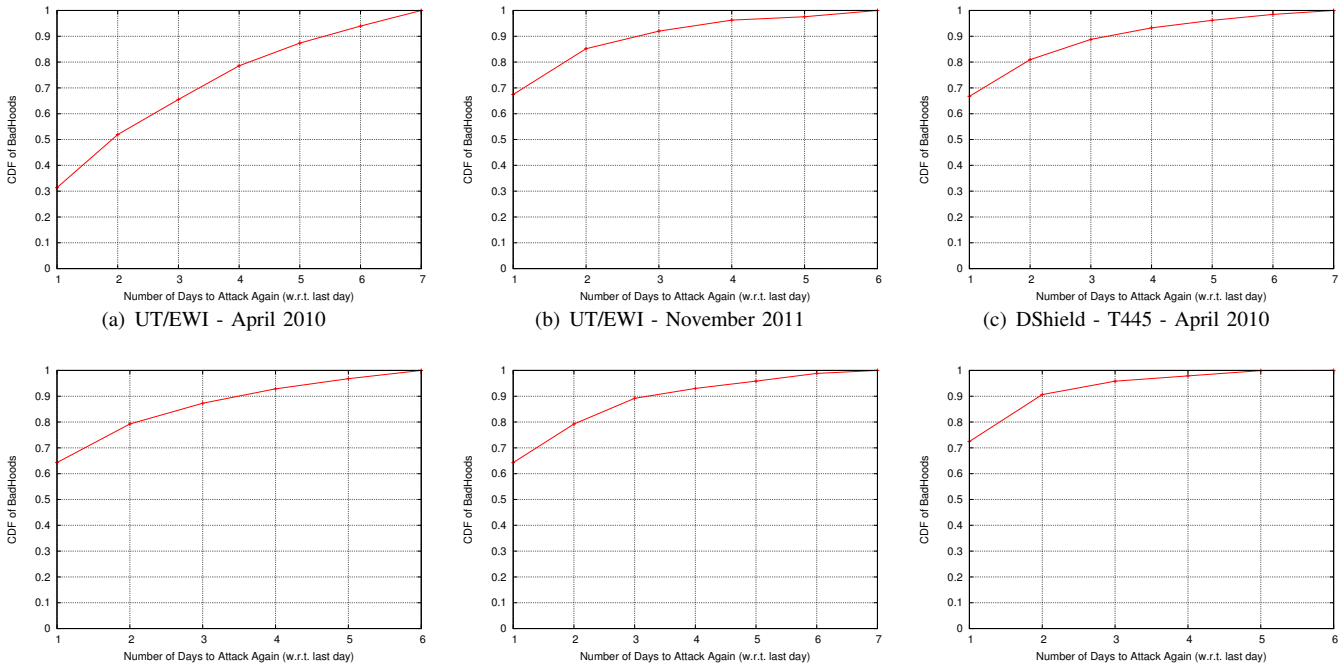

(a) UT/EWI - April 2010 (b) UT/EWI - November 2011 (c) DShield - T445 - April 2010

(d) DShield - T445 - November 2011 (e) DShield - T3389 - April 2010 (f) DShield - T3389 - November 2011

Fig. 6. Number of Days to Attack Again - CDF

shown in [8] that the model should be application-specific.

Acknowledgments: The authors would like to thank Marc Berenschot for his support for this research. Special thanks to the maintainers of CBL and DShield.

REFERENCES

[1] W. van Wanrooij and A. Pras, "Filtering Spam from Bad Neighborhoods," *International Journal of Network Management*, vol. 20, no. 6, pp. 433–444, November 2010.

[2] A. Ramachandran and N. Feamster, "Understanding the Network-level Behavior of Spammers," in *Proceedings of the 2006 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, ser. SIGCOMM '06. New York, NY, USA: ACM, 2006, pp. 291–302.

[3] M. P. Collins, T. J. Shimeall, S. Faber, J. Janies, R. Weaver, M. De Shon, and J. Kadane, "Using Uncleanliness to Predict Future Botnet Addresses," in *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, ser. IMC '07. New York, NY, USA: ACM, 2007, pp. 93–104.

[4] M. A. Rajab, F. Monrose, and A. Terzis, "On the effectiveness of distributed worm monitoring," in *Proceedings of the 14th conference on USENIX Security Symposium - Volume 14*. Berkeley, CA, USA: USENIX Association, 2005, pp. 15–15. [Online]. Available: http://dl.acm.org/citation.cfm?id=1251398.1251413

[5] G. C. M. Moura, R. Sadre, and A. Pras, "Internet Bad Neighborhoods: the Spam Case," in *Proceedings of the 7th International Conference on Network and Services Management (CNSM)*, October 2011, pp. 56–63.

[6] G. C. M. Moura, R. Sadre, A. Sperotto, and A. Pras, "Internet Bad Neighborhoods Aggregation," in *Network Operations and Management Symposium (NOMS), 2012 IEEE*, April 2012, pp. 343 –350.

[7] G. C. M. Moura, A. Sperotto, R. Sadre, and A. Pras, "Evaluating Third-Party Bad Neighborhood Blacklists for Spam Detection," in *IFIP/IEEE International Symposium on Integrated Network Management (IM 2013)*, Ghent, Belgium, May 2013.

[8] G. C. M. Moura, "Internet Bad Neighborhoods," Ph.D. dissertation, University of Twente, Enschede, The Netherlands, March 2013. [Online]. Available: http://dx.doi.org/10.3990/1.9789036534604

[9] J. Levine, "DNS Blacklists and Whitelists," RFC 5782 (Informational), Internet Engineering Task Force, Feb. 2010.

[10] CBL, "Composite Blocking List," 2012. [Online]. Available: http://cbl.abuseat.org/

[11] Passive Spam Block List, 2011. [Online]. Available: http://psbl.surriel.com/

[12] DSHIELD.org, "dshield Home — DShield; Cooperative Network Security Community - Internet Security," May 2012. [Online]. Available: http:/www.dshield.org

[13] F. Soldo, A. Le, and A. Markopoulou, "Predictive blacklisting as an implicit recommendation system," in *Proceedings of the 29th conference on Information communications*, ser. INFOCOM'10. Piscataway, NJ, USA: IEEE Press, 2010, pp. 1640–1648.

[14] G. C. M. Moura, R. Sadre, and A. Pras, "Taking on Internet Bad Neighborhoods," in *Network Operations and Management Symposium (NOMS 2014), IEEE (to appear)*, Krakow, Poland, May 2014.

[15] J. P. John, A. Moshchuk, S. D. Gribble, and A. Krishnamurthy, "Studying spamming botnets using botlab," in *Proceedings of the 6th USENIX symposium on Networked systems design and implementation*, ser. NSDI'09. Berkeley, CA, USA: USENIX Association, 2009, pp. 291–306.

[16] V. Chandra and N. Shrivastava, "Ways to Evade Spam Filters and Machine Learning as a Potential Solution," in *Communications and Information Technologies, 2006. ISCIT '06. International Symposium on*, 18 2006-sept. 20 2006, pp. 268 –273.

[17] J. Franklin, V. Paxon, A. Perrig, and S. Savage, "An inquiry into the nature and causes of the wealth of Internet miscreants," in *Proceedings of the 14th ACM conference on Computer and communications security*, ser. CCS '07. New York, NY, USA: ACM, 2007, pp. 375–388.

[18] DSHIELD.org, "About the Internet Storm Center— DShield; Cooperative Network Security Community - Internet Security," May 2012. [Online]. Available: http://www.dshield.org/about.html

[19] SANS, "SANS Information, Network, Computer Security Training, Research, Resources," May 2012. [Online]. Available: http://www.sans.org

[20] J. Zhang, P. A. Porras, and J. Ullrich, "Highly predictive blacklisting," in *USENIX Security Symposium*, P. C. van Oorschot, Ed. USENIX Association, 2008, pp. 107–122.

[21] V. Fuller and T. Li, "RFC 4632: Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan," August 2006. [Online]. Available: http://tools.ietf.org/html/rfc4632

[22] CBL, "The CBL FAQ," 2012. [Online]. Available: http://cbl.abuseat.org/faq.html

[23] M. van Eeten, J. M. Bauerb, H. Asgharia, S. Tabatabaiea, and D. Randc, "The Role of Internet Service Providers in Botnet Mitigation: An Empirical Analysis Based on Spam Data," in *WEIS 2010: Ninth Workshop on the Economics of Information Security*, 2010.