Measurement and Analysis of Global IP-Usage Patterns of Fast-Flux Botnets†

Xin Hu, Matthew Knysz, Kang G. Shin
The University of Michigan, Ann Arbor, MI 48109-2121, USA

***Abstract*—This paper considers the global IP-usage patterns exhibited by different types of malicious and benign domains, with a focus on single and double fast-flux domains. We have developed and deployed a lightweight DNS probing engine, called *DIGGER*, on 240 PlanetLab nodes spanning 4 continents. Collecting DNS data for over 3.5 months on a plethora of domains, our global vantage points enabled us to identify distinguishing behavioral features between them based on their DNS-query results. To help us analyze the enormous amount of data, we have quantified these features and designed an effective classifier capable of accurately discriminating between different types of domains. Applying the classifier on the 3.5-month DNS data allows us to reveal the relative prevalence of different fast-flux domains and conduct detailed studies on them separately. These results provide insight into the current global state of fast-flux botnets and their range in implementation, revealing potential trends for botnet-based services. We also uncover previously-unseen domains whose name servers *alone* demonstrate fast-flux behavior and a new, cautious IP management strategy currently employed by criminals to evade detection.**

I. Introduction

A botnet is a vast collection of compromised computers under the control of a botmaster utilizing a Command-and-Control (C&C) infrastructure. Among the numerous criminal uses of botnets, one of the more advantageous is the botnet-based hosting service, which proxies or redirects unsuspecting users to illegal or nefarious content. Used as a misdirection mechanism for evading detection, botnet-based hosting services often come in tandem with a variety of other criminal scams, constituting an essential portion of botnets' overall operation. For example, spam/phishing campaigns often utilize botnets for misdirection. Once victims click the URL links in the spam email, they connect to the bots, which then redirect them to—or serve as proxies for—the host of the nefarious content. This strategy grants criminals a high level of anonymity while enabling easy and centralized management of the malicious content. However, because botnets are composed primarily of compromised home computers with unreliable connectivity, it is not uncommon for them to unpredictably go offline (e.g., the computer is turned off or the installed malware is discovered and removed). Botnet-based hosting services, therefore, must be protected against the failure or disruption of individual bots, ensuring the availability and stability of the hosted service/content. As a result, they adopt fast-flux (FF) DNS techniques, which frequently change the domain-name mappings to different bots' IP addresses. Using this FF technique, botmasters effectively turned their botnets into a global Content Delivery Network (CDN), providing highly available and reliable content-hosting services despite frequent node failures/disconnectivity. This extends the lifetime of illegal activities the botnets provide, complicating disruption efforts by introducing an additional layer of misdirection.

† The work reported in this paper was supported in part by the Office of Naval Research under Grant N00014-09-11042

Continent	N. America	Europe	Asia	S. America	TOTAL
DIGGER Nodes	111	95	28	6	240
% of TOTAL	46.25%	39.58%	11.67%	2.50%	

TABLE I: *Global distribution of DIGGER nodes by continent*

Previous research mostly focused on understanding the malicious use of FF botnets in phishing scams [8] and devising effective detection systems for them [5], [9], [10]. However, little has been reported on botnets' IP-usage behavior from a *global* perspective. Because botnets are formed with myriad compromised hosts dispersed around the world, accurate characterization of how botmasters manage this vast number of IPs can only be achieved by collecting and analyzing data from a distributed and global perspective. In this paper, we attempt to achieve this goal by measuring FF botnet behavior from vantage points distributed around the world. This unique global perspective enables us to gain insight into various global IP-usage patterns of FF botnets that are inherent to their operation. The contribution of our work is three-fold. First, we build a global query engine called *DIGGER* that monitors—for an extended period of time—complete DNS behavior from 240 geographically-dispersed vantage points spanning four continents. Second, we conduct comprehensive studies on IP-usage patterns of different types of domains and propose effective methods to characterize and quantify unique features of FF botnet domains. This allows us to uncover several previously-unknown features of FF botnets and discover new, discreet IP-management strategies currently employed by criminals to evade detection. Third, to help us better analyze the current state of FF botnets and their relative prevalence, we design and implement a multi-level classifier capable of separating different types of malicious and benign domains based on their IP usage behavior. Applying the classifier on more than three months' worth of data allows us to spot potential trends of FF botnets and demonstrate the wide spectrum of their implementations.

The remainder of this paper is organized as follows. Section II presents the global IP-usage patterns for domain types. Section III describes our multi-level classifier and key findings on 3.5-month global DNS data. Section IV discusses related work and finally, Section V concludes the paper.

II. Measuring and Analyzing Global IP-Usage Patterns of Domains

In this section, we explore the DNS IP-usage patterns of different malicious and benign domains. First, we describe how we set up a globally-distributed monitoring system and give an overview of the different domain types we have observed in the gathered data. Then, we discuss various interesting features we've identified that are useful in understanding the inherent operations of the different domains types.

A. System Architecture

We created a distributed DNS-query engine called DIGGER, deployed on 240 geographically disparate nodes in the PlanetLab testbed [11]. The nodes were chosen based on the location of the DNS servers they queried, such that DIGGER would issue queries to DNS servers in different geographic locations around the world. Table. I shows the distribution of DIGGER nodes, which is reflective of the overall distribution of available PlanetLab nodes.

A rec	• The address (A) record in a DNS query on a domain. • The IP addresses of the domain's content servers.
NS rec	• The name server (NS) record in a DNS query on a domain. • The domain names (not IP adresses) of the domain's NSes.
NA rec	• The A rec in a DNS query on a domain's *name servers.* • The IP addresses of the domain's NSes.
Reverse DNS lookup	• The result of a DNS-query request for an IP's domain name. (i.e., PTR request) • When performing a DNS query on a domain, we also do a reverse DNS lookup on the domain's A and NA rec IPs.

TABLE II: *Domains' DNS record data gathered by DIGGER*

On each node, DIGGER performs DNS queries for a set of domains to gather the information shown in Table II. Based on a domain's most recently returned DNS results, DIGGER continues to dig active domains periodically based on their observed TTL, ensuring fresh DNS-query results. Domains determined to be offline are intermittently dug every 24 hours, so that DIGGER can discover if they come back online. The set of suspicious domains monitored by DIGGER is compiled from multiple sources, including online repositories of phishing [2] and malware [1] websites, and URL links embedded in the spam emails collected from a spam relay trap and recent additions to online repositories [4]. As a result, the domains probed by DIGGER tend to be malicious in nature, which is desirable for our purpose of studying malicious FF domains. DIGGER has been deployed and gathering global DNS-usage patterns for over 3.5 months in early 2009 on 5,169 active domains. Analysis of this data has revealed several distinct types of IP-usage patterns employed by malicious and benign domains. Next, we will describe these domain types whose differentiating features will be explored throughout this paper.

B. Domain Types

Before delving into the details of different domain features, we present the reader the following high-level overview of the domain-type nomenclature. To provide an intuitive view of these domain types, we have plotted the global IP usage—as seen from the DNS queries—for some representative domains in Fig. 1. In this figure, the *Time* axis represents the time since DIGGER started monitoring the domain; *Node Index* represents the DIGGER node that the IP was observed on, with positive values indicating an A rec IP and negative values an NA rec IP; *IP Index* is a unique index incrementally assigned to each newly-observed IP.

FF domains (Fig. 1 (a)-(c)) are malicious domains utilizing a fast-flux (FF) DNS-advertisement strategy, typically built atop botnets. Because bots may unexpectedly go offline, FF domains advertise numerous IPs in their DNS-query results, helping ensure some of the IPs belong to a functional bot. The TTL of the IPs used by FF domains tend to be relatively short; this permits the botmasters a finer level of control in replacing IPs advertised to the DNS servers, increasing the availability of an online bot and access to the malicious payload. It is this excessive number of constantly-changing IP addresses that qualifies a domain as "fluxy", and the domain is considered a FF domain. Domains exhibiting FF behavior in only a *single* record type (i.e., A rec or NA rec) are considered **FFx1 domains** (single fast flux). More specifically, FFx1 domains that are fluxy in their A recs (i.e., content servers) are termed **FFx1_Arec domains** (Fig. 1 (a)), while those that are fluxy in their NA recs (i.e., name servers) are termed **FFx1_NArec domains** (Fig. 1 (b)); FFx1_NArec domains are able to evade current detection strategies that focus on A recs by migrating their fluxy behavior to their NA recs, where it is less likely to be noticed. When FF domains are fluxy in *both* their A and NA recs, they are considered double fast flux, or **FFx2 domains** (Fig. 1 (c)).

CDN domains (Fig. 1 (d)) are valid, benign domains that uses a Content Delivery Network, such as Akamai, to improve the delivery of their content. CDNs—consisting of a system of computers networked together for the purposes of improving the performance and scalability of content distribution—produce DNS-query results resembling those of malicious FF domains: numerous, changing IPs per query with short TTL values. This affinity is a consequence of their similar goal to provide reliable content delivery despite node failure, as well as their shared assumption that any node can temporarily or permanently fail at any time. However, CDN domains demonstrate geographic awareness (i.e., IPs geographically close to a DNS server will be advertised with higher probability at that server) and load balancing (i.e., techniques improving performance and scalability not observed in FF domains).

Non-CDN domains (Fig. 1 (e)) are valid, benign domains that *don't* use a CDN for delivery of their content. Typically, non-CDN domains use a few stable content servers and name servers (NSes) during the entire monitored period.

MAL domains (Fig. 1 (f)) are domains that aren't fluxy enough to be considered FF domains, nor benign enough to be considered non-CDN domains. Their DNS behavior demonstrates potentially suspicious behavior often attributed with malicious domains. They tend to recruit more IPs than a non-CDN, but not nearly as many as a FF domain. For example, during a monitoring period of a few months, a FF domain typically advertises thousands of different IPs. A MAL

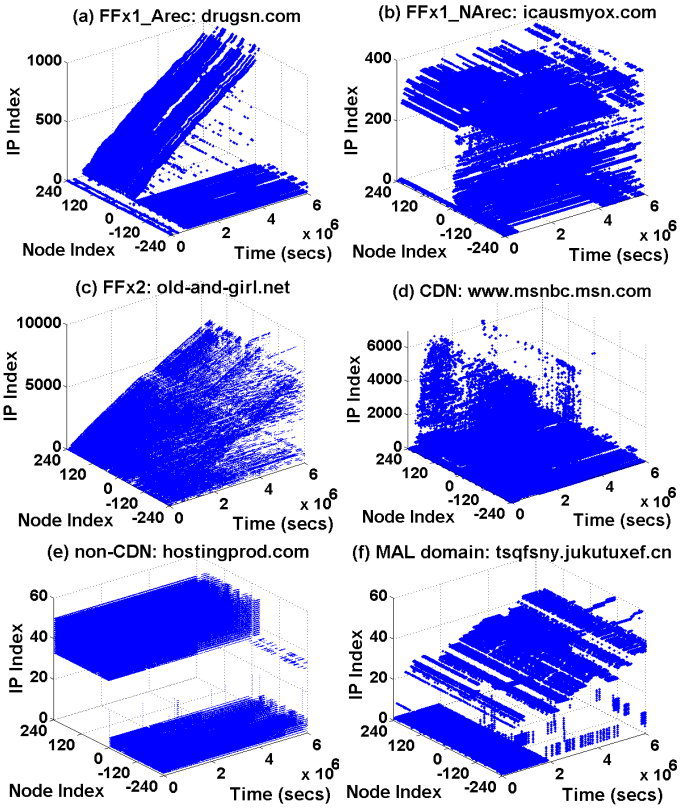

Fig. 1: *Global IP-usage patterns (in DNS query results) for some examples of the domain types*

domain, on the other hand, will advertise perhaps a total of 20-30 IPs—roughly one or two new IPs every few days. MAL domains will tend to slowly add more IPs because they will slowly lose some as their malicious activities are detected and their IPs are blocked[1]. The IPs used by MAL domains may consist of bots or valid servers being used for malicious means.

Having introduced the nomenclature we adopted to describe different domain types, we now present several interesting features of their IP usage patterns we discovered through the analysis of the globally collected data.

C. Number of Unique IP Addresses per Node

The first feature we examine is the number of unique IPs seen across the DIGGER nodes over time. Fig. 2 and 3 show the CDFs of the number of unique A and NA rec IPs observed by our 240 DIGGER nodes during the ≈3.5 month monitoring period. MAL domains have been omitted in the plots due to their similarity to non-CDN domains. Our empirical data reveals that non-CDN and FFx1_NArec domains (whose A recs behave like a non-CDN) use a small set of stable content servers. For example, in Fig. 2, neither the non-CDN nor the FFx1_NArec domains contain more than 18 unique A rec IPs per node. CDN domains are sometimes found to demonstrate a larger number of unique A rec IPs on some nodes, though

[1]Notice, websites hosted on home computers with dynamic IP addresses could be considered MAL domains by our definition. However, we consider this acceptable since most valid websites are not hosted on home computers, causing those that are to be inherently suspicious.

the number of nodes is considerably fewer than observed for FF domains. For example, for the CDN in Fig. 2, ≈2% of the DIGGER nodes observed more than 100 unique A rec IPs. On the other hand, FFx1_Arec and FFx2 domains clearly possess a much greater number of unique A rec IPs on a larger percentage of nodes —a direct consequence of the bots' unreliable connectivity. For the FFx1_Arec domain in Fig. 2, more than 35% of nodes detected over 200, and a few observed over 700. The numbers observed for the FFx2 domain are even higher, with over ≈63% of the nodes observing over 200, ≈43% more than 500, and several with more than 2,500.

While the number of unique A rec IPs per node appears a promising distinguishing feature, our data implies that this is not the case for the average number of unique *NA rec* IPs. From Fig. 3 it is apparent that CDN and FFx2 domains possess many more unique NA rec IPs per node than the other domain types. On average, FFx2 and CDN domains advertise 999 and 727 IPs on a single node respectively. It seems that, over time, CDNs can advertise numerous name server (NS) IPs. This behavior might arise from the CDN trying to ensure the availability of its NSes, affording it better control when performing load balancing. In any case, we found the behavior of the FF and CDN domains to be too similar, causing the number of unique NA rec IPs to be an indistinctive feature.

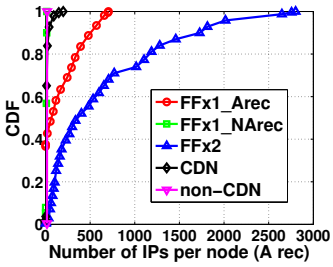

Fig. 2: *CDF: # of unique A rec IPs per DIGGER node*

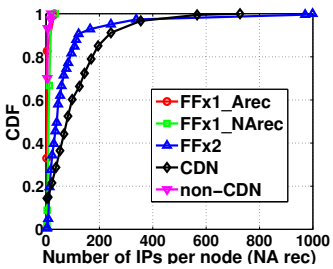

Fig. 3: *CDF: # of unique NA rec IPs per DIGGER node*

D. Overlap between IPs of A and NA Records

While analyzing our data, it quickly became apparent that FF domains tend to exhibit some IP overlap. We were seeing IPs advertised for a domain's A rec reappearing in the same domain's NA rec. Table III shows the total number of A rec, NA rec, and overlap IPs for some representative domains from each domain type . This overlap phenomenon was much more prevalent in FFx2 domains than either type of FFx1; we never observed it in valid domains. The FFx1 domains almost entirely use valid IPs for one record type and the IPs of compromised computers for the other. While the representative MAL domains have a small number of total unique IPs (like a non-CDN domain), their IP overlap is exceptionally high, with almost all of their A rec IPs also used for their NA recs, thus setting them apart from valid domains. The IP overlap we empirically observed demonstrates that valid domains use separate machines for their content and name servers to prevent a single point of failure. FF and MAL domains, on the other hand, attempt to make the most of their limited resources, reusing IPs for both the A and NA records. Clearly, IP overlap

is a useful feature for differentiating between malicious and benign domains, especially FFx2 and MAL domains.

Domain Type	Domain	A rec	NA rec	Overlap
FFx1_Arec	drugsn.com	932	33	0
	www.couldchoose.com	486	37	5
FFx1_NArec	icausmyox.com	16	370	1
FFx2	old-and-girl.com	5,227	3,047	879
	mountainready.com	4,060	2,219	2,144
MAL	duelready.com	16	32	15
	tsqfsny.jukutuxef.cn	23	42	20
CDN	www.msnbc.msn.com	1,160	5,412	0
non-CDN	hostingprod.com	18	32	0

TABLE III: *Total A, NA, & overlap IPs for diff domain types*

Table III also shows that, because non-CDN and MAL domains advertise only a few stable content and name servers, their total number of unique A and NA rec IPs is meager when compared to CDN and FF domains making it a useful distinguishing feature.

E. Continental Distribution of IPs

Next, we examine how the various domain types differ in their IP distribution (i.e., where the IPs returned in DNS queries are geographically located). We examine the IP location based on continent rather than country, because the close proximity of European countries made a country-based resolution too finely-grained. In particular, we examined two features: 1) *percentage of IPs from the wrong continent*, i.e., what percentage of IPs returned in DNS queries are located in a different continent than the queried DNS server; 2) *continental IP distribution*, i.e., from the perspective of each continent containing queried DNS servers, what percentage of IPs returned are located in each continent.

Domain Type	% wrong continent		Domain Type	% wrong continent	
	A rec	NA rec		A rec	NA rec
FFx1_Arec	83.04%	45.43%	MAL	90.38%	89.04%
FFx1_NArec	93.21%	84.54%	CDN	23.24%	17.56%
FFx2	68.17%	63.99%	non-CDN	53.77%	61.11%

TABLE IV: *Percentage of IPs from the wrong continent*

Table. IV shows the percentage of A and NA rec IPs from the wrong continent for some representative domains. From the table, it is evident that the CDN domain has a considerably smaller proportion of IPs from the wrong continent than the other domain types. The few CDN IPs from the wrong continent are due to load balancing. For example, to distribute load when traffic volume is high in Asia, CDNs may advertise some European IPs to Asian DNS servers, resulting in a small percentage of IPs from the wrong continent.

Insight into continental IP distribution can be found in Fig. 4. For brevity, we have not plotted any FFx1 domains, since their results are a subset of the FFx2 domain type; likewise, we have omitted plots for a MAL domain (since their distribution is functionally similar to non-CDN domains) and for the NA recs' distribution (since the results are similar to those for the A recs). In Fig. 4, the bars represent the continental IP distribution from different perspectives. In each domain's plot, the first bar represents the continental IP distribution from a global perspective, while the other bars are from the perspective of the different continents where we deployed DIGGER nodes. For example, the bar labeled "Asia" under *old-and-girl.com* indicates the percentage of IPs located in each continent base on queries to Asian DNS servers.

From Fig. 4, we can see that the continental IP distribution for CDN domains varies greatly across the different continents, clearly revealing the *location-aware DNS advertisement* employed by CDNs. The DNS query results for CDN domains often contain a majority of IPs located near the query issuer, providing fast, reliable services and quicker content delivery to end users by reducing the data's travel distance. Consequently, CDNs demonstrate a smaller percentage of IPs from the wrong continent and a larger variance in continental IP distribution than other domain types. On the other hand, we found that MAL and non-CDN domains operate in a similar manner. They indiscriminately advertise their small pool of a few stable server IPs around the world nearly simultaneously. This causes the continental IP distribution at each continent to be the same, and the percentage of IPs from the wrong continent will reflect the global distribution of our DIGGER nodes. Finally, our analysis suggests that FF domains adopt an advertisement strategy dictated by the unstable nature of their constituent bots, which we term *necessity-based DNS advertisement*. Since bots can go offline at any time, FF domains must rely on whichever bots are currently available and advertise available IPs to DNS servers around the globe as necessity dictates, regardless of geographic location. This results in a large percentage of IPs from the wrong continent and a fairly consistent continental IP distribution across continents.

These findings indicate that the percentage of IPs from the wrong continent and the variance of the continental IP distribution are useful features for distinguishing CDN domains from the other domain types.

F. IP Recruiting

In this subsection, we study the distinct strategies employed by FF, CDN, and non-CDN domains, when they advertise IPs to DNS servers. For a given domain, we assigned a unique IP index to each newly-seen IP in the DNS query results across all DIGGER nodes. This IP index is plotted against time for example FFx2, CDN, non-CDN, and MAL domains[2] in Figs. 5–8. The points in the graphs represent when a new IP was returned in a DNS query on a global scale. Therefore, the slope of each curve demonstrates the rate, or speed, with which a domain seems to globally "recruit" more IPs. Notice, when we discuss recruitment, we mean the *apparent* recruitment of IPs based on the DNS query results, not the actual recruitment of bots via compromising computers.

Recruitment Speed: refers to the speed (or rate) at which one observes new, unique IPs for a given domain when monitoring that domain's DNS queries over time.

Fig. 5 shows how a FFx2 domain slowly and nearly continuously accrues unique IPs over its entire online lifetime, with

[2]FFx1_Arec and FFx1_NArec are essentially specific subsets of FFx2 domains, so their plots are not included for brevity.

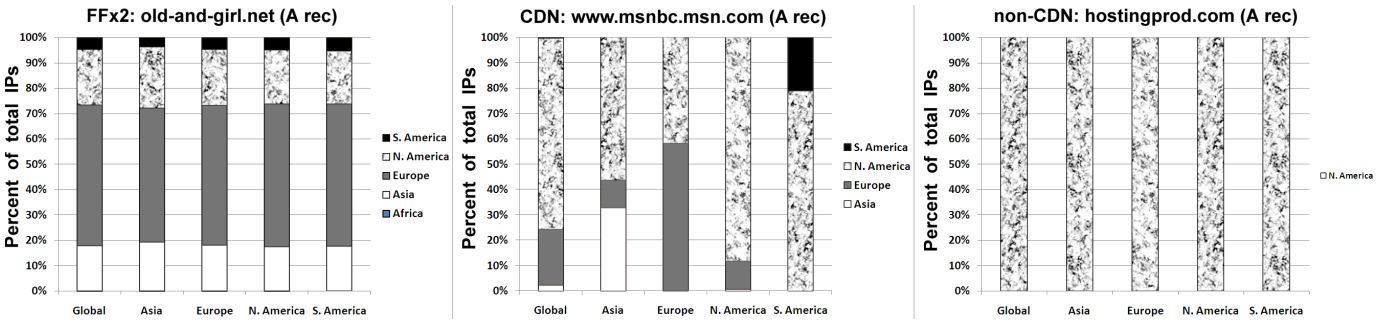

Fig. 4: *Percentage of total A rec IPs seen from each continent by DIGGER nodes globally and in each continent*

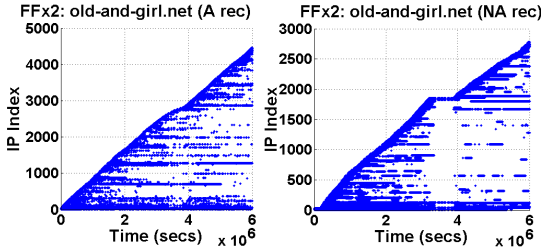

Fig. 5: *Global IP usage for example FFx2 domain*

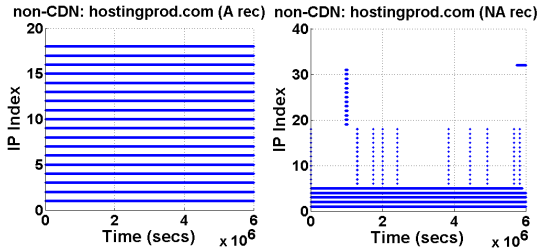

Fig. 7: *Global IP usage for example non-CDN domain*

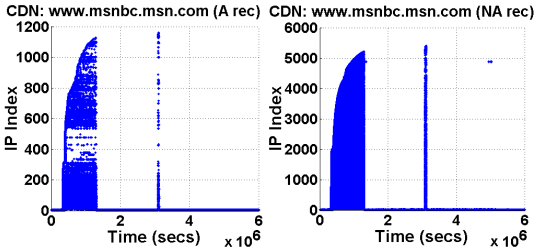

Fig. 6: *Global IP usage for example CDN domain*

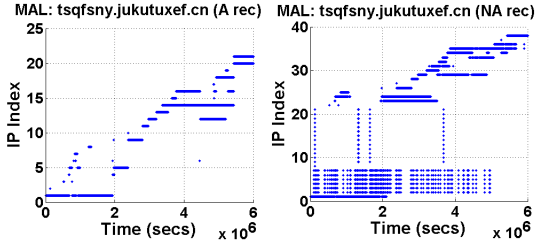

Fig. 8: *Global IP usage for example MAL domain*

short, intermittent periods of stability. These results indicated that FF domains must continually add new IPs to help ensure reliable delivery of their nefarious content. In addition, the bots used by FF domains may obtain dynamic IP addresses from their Internet Service Provider. Consequently, a bot may be assigned different IPs over time, causing our DIGGER nodes to observe the apparent recruitment of new IPs; this effect, called *DHCP churn*, is not present for valid domains using stable servers with static IPs.

Meanwhile, when viewed globally, we have discovered that CDN domains (Fig. 6) achieve a much faster recruitment speed, indicating that they advertise IPs from a large pool of stable IP addresses, which they rotate quickly and efficiently for performance purposes, such as load balancing. Since CDNs advertise their IPs in a geographically-conscious manner (e.g., a DNS query in Asia will often result in a different set of IPs than in Europe), DIGGER's global perspective observes most of the CDN's IPs in a short period of time. In contrast, FF domains use necessity-based DNS advertisement, advertising the same pool of IPs irrespective of the DNS servers' geographic location. Thus, while FF domains may change their advertised IPs as quickly as a CDN, DIGGER's global perspective doesn't allow it to observe many more IPs than at any given local vantage point, resulting in the comparatively slower IP recruitment rate.

Non-CDN domains (Fig. 7), on the other hand, hardly recruit any additional IPs over time. Rather, their IP pools consist of a small number of stable servers that are almost simultaneously advertised to DNS servers around the world.

MAL domain (Fig. 8) often demonstrates the slow and somewhat steady recruitment of IPs. This behavior is likely the result of the MAL domains' malicious activities being detected and their IPs blocked, requiring them to register fresh IPs with DNS to maintain content availability. Unlike FF domains which recruit thousands of IPs, MAL domains recruit only tens of IPs over 3.5 months. This drastic difference should prove beneficial in distinguishing MAL domains from non-CDN and FF domains.

Recruitment Period: represents the amount of time during which new IPs are seen for a given domain when monitoring that domain's DNS queries over time. Our data indicates that non-CDN domains (Fig. 7) have almost no recruitment period; a small pool of very stable IPs are advertised initially and used throughout the lifetime of the domain. On the other hand,the fast recruitment speed of CDN domains causes DIGGER to quickly observe most of their available IPs, resulting in a short recruitment period at the onset of monitoring followed by a longer, stable period consisting mainly of previously-

seen IPs. From Fig. 6, we can see that the CDN's recruitment period is smaller than its total online period. After its initial recruitment period, the CDN domain stabilizes and advertises a much smaller set of IPs before a quick advertisement spike followed by another stable period; the stable period looks like a gap in the graph, but closer examination reveals a small set of IPs with low IP indices (i.e., the earliest seen IPs). We have also discovered, as shown in Fig. 5, that the fluxy records for FF domains recruit new IPs for nearly the entire duration of the domains' online period, with only short, intermittent periods of stability. This constant IP recruitment is a result of the unreliable nature of the compromised computers serving as bots. The varying recruitment periods we have discovered for the different domain types should provide a useful metric for distinguishing between them.

G. Other Features

We also examined several other potential features such as *TTL*, *reverse DNS lookup*, *average IP online time*, etc. We found them less effective when designing the classifier for differentiating domain types. Due to space constraints, we have omitted their results and refer the interested reader to our technical report for further discussion.

III. Analysis of Current FF Botnet Threat

After introducing different types of domains and their distinguishing features, in this section we will analyze the current status of FF botnets. To aid us in analyzing the landscape of this global threat, we develop a multi-leveled classifier that uses the differentiating features identified in Section II to automatically determine the domain type distribution from the collected data. Note that the classifier is not designed to be a real time detector for FF botnets. The main purpose of building such a classifier is to help us separate different types of FF domains and expose emerging trends in their DNS management.

A. Classification Feature Quantification

Table V shows 7 features, F1-F7, that were calculated for each monitored domain over the total ≈3.5 month duration and then analyzed by the classifier to determine the domain type. The calculation is straightforward for all features except F4 (IP distribution). As discussed in Section II-E, CDNs, due to their location-aware DNS strategy, tend to have a significantly larger variance in continental IP distribution than other domain types. To quantify this effect, for each of the 4 continents with DIGGER nodes deployed, we create a vector with 8 elements, each of which represents the number of IPs observed from one particular continent[3]. Then we calculate the cosine similarity between every pair of the vectors and take the average, producing the value for F4. The closer this value is to 1, the more similar the continental IP distributions appear on each continent, and the less likely the domain is a CDN domain. The 3rd column in Table V shows how each feature will likely group the different domain types, represented by square brackets.

Classification Feature	DNS Record	Domain Type Classification Groups
F1. Avg. # of unique IPs per Node	A	• [CDN_1, non-CDN_2, MAL_3, FFx1_NArec_5] • [$FFx2_4$, FFx1_Arec_5]
	NA	• [CDN_1, $FFx2_4$] • [FFx1_NArec_5] • [non-CDN_2, MAL_3, FFx1_Arec_5]
F2. A & NA rec overlap	A & NA	• [CDN_1, non-CDN_2] • [MAL_3, $FFx2_4$, FFx1_Arec_5, FFx1_NArec_5]
F3. % IPs from wrong continent	A	• [CDN_1] • [non-CDN_2, MAL_3, $FFx2_4$, FFx1_Arec_5, FFx1_NArec_5]
	NA	
F4. Continental IP distribution's average cosine similarity	A	
	NA	
F5. IP recruiting speed	A	• [CDN_1] • [non-CDN_2, FFx1_NArec_5] • [MAL_3] • [$FFx2_4$, FFx1_Arec_5]
	NA	• [CDN_1] • [non-CDN_2, FFx1_Arec_5] • [MAL_3] • [$FFx2_4$, FFx1_NArec_5]
F6. IP recruiting period	A	• [CDN_1] • [non-CDN_2, FFx1_NArec_5] • [MAL_3] • [$FFx2_4$, FFx1_Arec_5]
	NA	• [CDN_1] • [non-CDN_2, FFx1_Arec_5] • [MAL_3] • [$FFx2_4$, FFx1_NArec_5]
F7. Total unique IPs	A	• [CDN_1, $FFx2_4$, FFx1_Arec_5] • [non-CDN_2, MAL_3, FFx1_NArec_5]
	NA	• [CDN_1, $FFx2_4$, FFx1_NArec_5] • [non-CDN_2, MAL_3, FFx1_Arec_5]

*The number by each domain type represents the level it is classified by our SVM. When selecting features for SVM-**x**, domains with a number < **x** can be ignored, since they have already been classified and removed from the unknown set.*

TABLE V: *Features to classify domain types into diff groups*

B. SVM (Support Vector Machine) Classifier

Fig. 9 shows the design of our multi-leveled classifier and the results of our separate training and test sets (test sets are all DIGGER collected DNS data minus the training set). The classifier is based on a particular machine learning approach called SVM [6], which has been successfully applied to a large number of practical problems and has the nice properties of minimizing the empirical classification error and maximizing the geometric margin. In our work, we build one linear SVM classifier[4] at each level that classifies a domain type from the test set. The classifier progressively reduces the number of unknown domains and thus simplifies subsequent classification. In Fig. 9, each oval represents a classified domain type, while each rectangle represents the remaining unknown. The values for "Train" show how many examples of a given domain type (or group of domain types) were used when training that level of the classifier. The values for "Test" indicate the number of domains that were classified (or remained to be classified) when we applied each tier of the classifier. We manually identified about 10 representative domains of each type to be used in training of SVM at each level. More difficult to detect by hand, we were only able to manually identify one FFx1_NArec domain.

[3]N. America, S. America, Europe, Asia, Africa, Oceania, Antarctic and "unknown"

[4]We choose a linear SVM because it is simple and effective. Previous work [5] has demonstrated its efficacy in detecting FF botnets from a single vantage point, and in our experiments, it produced results of comparable accuracy to other SVM kernels and decision trees.

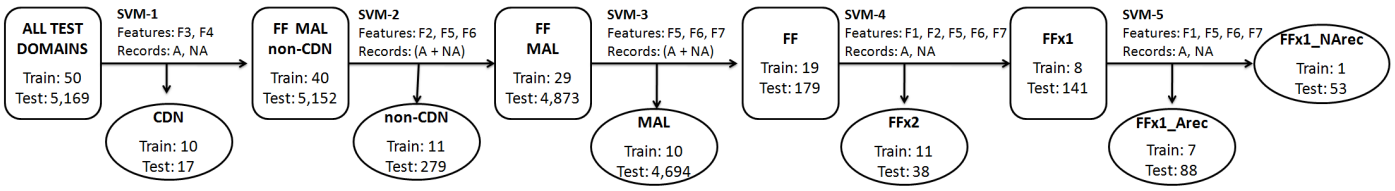

Fig. 9: *SVM flowchart*

	b	F1		F2	F3		F4		F5			F6			F7			Result > 0
	(bias term)	A	NA		A	NA	A	NA	A	NA	(A + NA)	A	NA	(A + NA)	A	NA	(A+NA)	
SVM-1	284.69				-108.10	-88.90	-124.83	-160.60										*CDN*
SVM-2	128.26			-217.20							47.23			-2,072.45				*non-CDN*
SVM-3	192.04										1.32E-06			-4.06E-03			-1.77	*MAL*
SVM-4	-390.75	3.13	12.54	0.27					-1.14E-03	-0.03		0.42	0.21		-0.37	0.38		*FFx2*
SVM-5	933.52	0.42	-0.03						2.28E-06	0.02		7.30E-04	-4.01E-03		1.74	-5.31		*FFx1_Arec*

TABLE VI: *Linear SVM equations. ((A + NA) means combined IP pool of the A and NA recs)*

Table VI shows the bias and feature weights for each level of trained classifier. Those features not used at a particular level are shaded black. For each SVM, the *Result* is calculated as the *bias term* plus the product of each feature and its weight. The "*Result* > 0" column indicates how a domain with a positive *Result* will be classified. The exception is FFx1_NArec domains, which are classified when SVM-5's *Result* is negative. Additionally, the magnitude of *Result* signifies the confidence in classification choice.

Due to the similarities some domain types share between certain features, the *order* we apply the classifiers and which features we use at each level becomes important. Using Table V as a guide, we experimented with the features used at each level of our classifier, finding an optimal order for those levels that exploits the strong differentiating features between certain domain types. For example, **SVM-1** uses F3 and F4 as strong indicators of CDN domains. This is because only CDN domains, practicing a location-aware DNS advertisement strategy, will obtain a positive *Result*. Other domains, with a large percentage of IPs from the wrong continent (F3) and similar IP distributions (F4) will generate a negative *Result*. Since CDN domains can behave similarly to FF domains in other respects (e.g., large number of IPs), removing them first will improve successive classification. Because valid non-CDN domains do not possess the recruitment and IP overlap behavior common to MAL and FF domains, **SVM-2** relies on F5, F6, and F2 on the combined (A + NA) recs to separate non-CDN domains. Notice Table VI shows that F6 is the dominating feature, i.e., if the domain demonstrates any significant recruitment period, it is unlikely to be a non-CDN domain. When reaching **SVM-3**, the test set was entirely composed of malicious domains (i.e., FF and MAL). It is logical to classify MAL domains next. Because of their faster IP recruitment rate, FF domains quickly outpace MAL domains. This allows **SVM-3** to use number of IPs (F7) and recruitment features (F5 and F6) to accurately identify MAL. After three stages of the classifier, only FF domains remained in the test set. By definition, the only thing distinguishing between FFx2, FFx1_Arec, FFx1_NArec domains is which record (A, NA or both) type demonstrates fluxiness. Thus **SVM-4** uses features related to fluxy behavior, i.e., F1, F2, F5, F6 and F7, on the individual A and NA recs to discern FFx2 from FFx1. Except F2, The same set of features are used in **SVM-5** to discriminate between FFx1_Arec and FFx1_NArec domains. F2 is ignored at this stage because the FFx1 domains experience comparable, modest-to-no IP overlap.

Applying this multi-level classifier on all ≈3.5 months of DNS data identified 17 CDNs (which we manually verified by visiting each domain) and 179 FF domains (which were also manually verified based on several heuristics, such as reverse DNS names indicating compromised computers, aggressive IP usage, etc). Additionally, the classifier identified 279 non-CDNs and 4694 MAL domains. We manually analyzed over 50 border cases from both types with *Results* closest to 0 and found them to be satisfactorily classified. The detailed analysis of these domains will be presented in the next section.

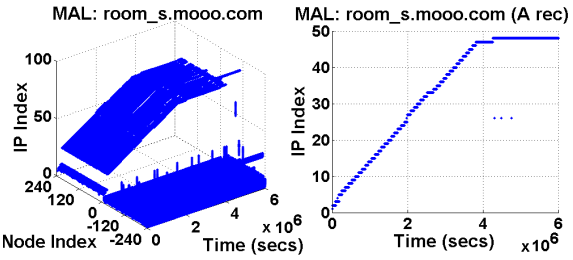

(a) 3D plot of A and NA rec (b) 2D plot of A rec

Fig. 10: *Cautious MAL domain*

C. Results

1) Cautious MAL domains: While manually validating SVM-3's results, we discovered 4 borderline MAL domains exhibiting atypical IP behavior, one of which is shown in Fig. 10. Recruiting less than 50 A rec IPs over ≈2.5 months[5], it is not fluxy enough to be considered a FFx1_Arec domain. However, its uncannily regular IP recruitment distinguishes it from other MAL domains. The domains advertise only a *single*

[5]The domain was DNS-parked afterwards. DNS parking is a technique used to block malicious domains by 'parking' the malicious domain name to an innocent IP address often controlled by the registrar.

A rec IP per query, with a max TTL of one minute. Despite this fine level of control, the domains only replace the IP about once a day, adhering to a meticulously precise schedule. Additionally, we can see from Fig. 10, that once changed, the A rec IPs are not reused. In this sense, they appear to be a type of *cautious* MAL domain, regularly and preemptively replacing their A rec IPs before they can be detected and blocked. When required, the short TTL permits rapid response. Although only 4 instances are observed, this strategy is interesting and may gain popularity among malicious domain owners trying to evade current detection technologies.

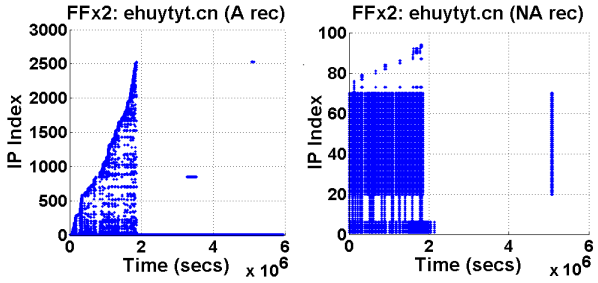

Fig. 11: *Classified FFx2 domain*

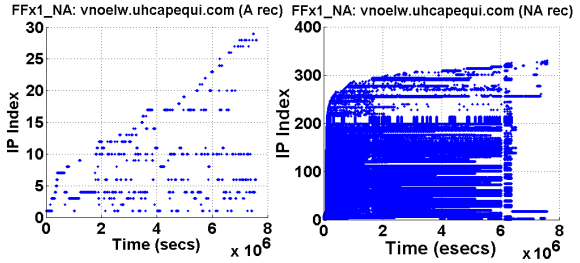

Fig. 12: *Classified FFx1_NArec domain*

2) FF domain Classification: An interesting aspect of our classifier is how it distinguishes between the FFx2, FFx1_Arec and FFx1_NA domains. Recall from Table VI that F1 is the dominant feature for SVM-4, with the NA rec being 4x as influential as the A rec. The consequence of this asymmetric weighting of fluxiness can be witnessed for the domains in Fig. 11 and Fig. 12. Both domains demonstrate definite fluxy behavior in one of their record types, i.e., A rec for Fig. 11 and NA rec for Fig. 12. With less than 30 IPs, the recruitment behavior of A rec in Fig. 12 resembles that of a MAL domain and the IP overlap is less than 4%. Thus, the classifier has performed correctly: a domain with FF behavior in its NA rec and MAL behavior in its A rec should be considered a FFx1_NArec domain. However, it isn't immediately obvious why Fig. 11 is considered fluxy in its NA rec, which appears relatively stable. Further analysis revealed that the domain has an IP overlap of ≈26%. That is, the fluxy A rec contributes over $\frac{1}{4}$ of the NA rec IPs, the less stringent fluxiness demands for the NA rec are met, and the domain is classified as FFx2.

3) Domain Type Distribution: Table VII shows the number and distribution for each domain type identified by our classifier. We find that the benign domains constitute less than 6% of the test set (i.e., 17 CDN and 279 non-CDN domains). Of the 4,873 nefarious domains, ≈96% were MAL

Domain Type	# of domains	% of ALL	% of MAL/FF	% of FF	% of FFx1
ALL	5,169				
CDN	17	0.33%			
non-CDN	279	5.40%			
MAL/FF	**4,873**	**94.27%**			
MAL	4,694	90.81%	96.33%		
FF	**179**	**3.46%**	**3.67%**		
FFx2	38	0.74%	0.78%	21.23%	
FFx1	**141**	**2.73%**	**2.89%**	**78.77%**	
FFx1_Arec	88	1.70%	1.81%	49.16%	62.41%
FFx1_NArec	53	1.03%	1.09%	29.61%	37.59%

TABLE VII: *Relative distributions of the various domain types*

domains and 179 were FF domains. Because we generated the domain list from suspicious sources, extracting this low number of benign domains and large number of MAL domains are justified. Particularly, this plethora of MAL domains results from their ease of management as the traditional and most popular mechanism employed by malicious websites (e.g., hosting nefarious content on less reputable web hosts and switching to a new server if detected). While out of this paper's scope, further research and analysis into the makeup and classification of MAL domains is warranted.

The additional level of misdirection and the nearly limitless supply of IPs provided by botnets make FF domains appealing. Thus far, it has been primarily FFx1_Arec domains observed in the wild, and their popularity is supported with our findings: ≈49% of the FF domains are FFx1_Arec. This is because FFx1_Arec domains provide the greatest return on their investment, affording botmasters an additional layer of misdirection without the hassle of maintaining volatile botnet NSes (name servers). With stable NSes, botmasters can easily replace offline IPs to avoid an interruption of service. Unfortunately for botmasters, security professionals have become aware of the FFx1_Arec botnet technique, devising clever detection strategies. While botnets provide a steady source of fresh A rec IPs, the NSes can still be blocked, crippling the botmaster's control until new NSes can be acquired. In an apparent attempt by botmasters to overcome this limitation, we witnessed a considerable presence of FFx2 domains, composing ≈21% of the FF domains. FFx2 domains provide further misdirection and protection for the botmaster, guarding the NSes against simple countermeasures such as IP blocking at the expense of a more diligent management effort.

Among FFx2 domains, we noticed a definite trend to use their more stable bots for the NA recs, often concurrently using them for the A recs as well. Interestingly, analysis of the identified FFx2 domains revealed a spectrum in the amount of NA rec fluxiness incorporated by botmasters. Obviously, we observed domains that were incredibly fluxy in both record types, as demonstrated by *old-and-girl.com* (Fig. 5). While it's interesting to observe these aggressive FFx2 domains in the wild, it was the FFx2 domains at the other end of the spectrum that proved more insightful. As an example, recall the more modest FFx2 domain *ehuytyt.cn* (Fig. 11). With over 2,500 unique A rec IPs, *ehuytyt.cn* is considerably more fluxy in its A rec than its NA rec. By using bot IPs from its A rec for roughly $\frac{1}{4}$ of its NA rec IPs, FFx2 domains

like *ehuytyt.cn* benefit from the increased control and stability provided by traditional NSes, while—for a minimal increase in management—simultaneously enhancing the domain's resilience to subversion.

Another interesting discovery is the apparent popularity of FFx1_NArec domains, accounting for ≈30% of the total FF domains observed. It seems that botmasters have become aware of security professionals analyzing domains' A recs for FF behavior. Consequently, they have migrated the fluxy behavior to the NA rec, where it is less likely to be noticed since most existing approaches only monitor A recs. Fig. 12 is a typical example of the FFx1_NArec domains identified by our classifier. It demonstrates a MAL domain strategy for its A rec IPs and a FF strategy for its NA rec IPs. This results in the domain appearing more benign when its A recs are analyzed, while providing the botmaster with a fine level of control over the NSes. Should the domain's malicious activity be detected and the A rec IPs blocked, the botmaster, having retained control over the NSes, can replace the IPs with minimal service interruption. Additionally, by controlling the NSes through the use of bots, botmasters can potentially determine when detectors perform DNS probing, allowing them to take appropriate countermeasures. The implication of this discovered behavior is straightforward: both record types must be monitored for fluxy behavior in order to quickly identify FF domains and their botnets.

IV. Related Work

The increasing popularity of FF domains has resulted in a number of proposed detection approaches [5], [9], [10] They all begin by gathering suspicious domains from various sources, actively/passively monitoring domains' DNS-query results, and extracting a set of unique features. A linear decision function [5], a naïve Bayesian classifier [9], or a decision tree [10] are then applied on these features to determine whether a domain is a FF domain. Nazario and Holz [8] later applied a similar approach to track the use of FF domains, demonstrating that continuous data mining of DNS records can yield insights into the operations of FF botnets. Caglayan *et al.*[3] analyzed several behavioral features of FF networks including their lifespan and network size. They showed that different FF networks share common life cycle characteristics. More recently, Konte *et al.*[7] studied the dynamics of FF networks from the perspective of online scam hosting infrastructures for different spam campaigns. Their results suggested that some persistent features may be useful in the detection of FF botnets.

This paper differs from previous work in the following ways. First, our goal is not to propose a deployable FF detector[6]. Instead, our classifier is developed as an analytical tool to automatically analyze the massive amount of data from a previously unattempted global perspective and provide fine-grained classification of FF domains. Second, previous work collected data from a *single* vantage point, and hence, may fail to capture useful features only discoverable with a global view. By contrast, we deployed a large number of sensors around the world, providing a global perspective for different types of FF domains. Finally, in this paper, we conduct a comparative analysis of the different IP-management schemes used by FF and CDN domains. Because of the shared goal of providing reliable content delivery despite node failure, FF and CDN domains demonstrate similar behavior (i.e., numerous, frequently changing IPs), which might cause misclassification for detectors.To solve the problem, we propose effective features based on location awareness for separating CDN domains from other domain types, which could potentially be leveraged to improve the current detectors.

[6]Incidentally, it has yielded insight into a deployable detector, and we are currently exploring the most efficient use of global vantage points and a detection window for an accurate, real-time detection system.

V. Conclusion

In this paper, we examined the global IP-usage patterns exhibited by different types of malicious and benign domains. We have deployed a lightweight probing engine on 240 PlanetLab nodes spanning 4 continents. Collecting DNS data for over 3.5 months on a plethora of domains enabled us to measure their IP-usage patterns from a global perspective and identify various behavioral traits shared among certain domains, resulting in the classification of several common domain types based on distinguishing DNS features. By applying a multi-level classifier on the data set, we have shown the relative popularity of the domain types in the malware/phishing/spam landscape. This provided valuable insight into the current state of FF domains, including the increased presence and versatile implementation range of FFx2 domains as well as an apparent trend towards using FFx1_NArec domains, which were previously unseen in the wild. We hope these new insights will foster improvements to existing defensive systems, allowing them to identify more sophisticated FF domains.

References

[1] DNS-BH - Malware Domain Blocklist. http://www.malwaredomains.com/files/domains.txt.

[2] Phishtank. http://www.phishtank.com/, 2009.

[3] A. Caglayan, M. Toothaker, D. Drapaeau, D. Burke, and G. Eaton. Behavioral analysis of fast flux service networks. In *CSIIRW '09: Proceedings of the 5th Annual Workshop on Cyber Security and Information Intelligence Research*, 2009.

[4] B. Guenter. Spam archive. http://untroubled.org/spam/.

[5] T. Holz, C. Gorecki, K. Rieck, and F. C. Freiling. Measuring and Detectin Fast-flux Service Networks. In *Proceedings of the Network and Distributed System Security Symposium*, 2008.

[6] T. Joachims. Making large-scale support vector machine learning practical. In *Advances in Kernel Methods: Support Vector Machines*. MIT Press, Cambridge, MA, 1998.

[7] M. Konte, N. Feamster, and J. Jung. Dynamics of Online Scam Hosting Infrastructure. In *Tenth Passive and Active Measurement conference*, 2009.

[8] J. Nazario and T. Holz. As the net churns: Fast-flux botnet observations. In *Malicious and Unwanted Software*, 2008.

[9] E. Passerini, R. Paleari, L. Martignoni, and D. Bruschi. FluXOR: detecting and monitoring fast-flux service networks. In *Proceedings of the 5th DIMVA*, 2008.

[10] R. Perdisci, I. Corona, D. Dagon, and W. Lee. Detecting Malicious Flux Service Networks through Passive Analysis of Recursive DNS Traces. In *Proceedings of ACSAC '09*, 2009.

[11] PlanetLab. http://www.planet-lab.org/.