

Measuring the Effectiveness of Infrastructure-Level Detection of Large-Scale Botnets

Yuanyuan Zeng[†] Guanhua Yan[‡] Stephan Eidenbenz[‡] Kang G. Shin[†]

[†] University of Michigan
{gracez, kgshin}@eecs.umich.edu

[‡] Los Alamos National Laboratory
{ghyan, eidenben}@lanl.gov

Abstract—Botnets are one of the most serious security threats to the Internet and its end users. In recent years, utilizing P2P as a Command and Control (C&C) protocol has become popular due to its decentralized nature that can help hide the botmaster’s identity. Most bot detection approaches targeting P2P botnets either rely on behavior monitoring or traffic flow and packet analysis, requiring fine-grained information collected locally. This requirement limits the scale of detection. In this paper, we consider detection of P2P botnets at a high-level—the infrastructure level—by exploiting their structural properties from a graph analysis perspective. Using three different P2P overlay structures, we measure the effectiveness of detecting each structure at various locations (the Autonomous System (AS), the Point of Presence (PoP), and the router rendezvous) in the Internet infrastructure.

I. INTRODUCTION

A botnet consists of a group of coordinated compromised computers or bots that can mount attacks, such as Distributed Denial of Service (DDoS), spamming, phishing and identity theft. Botnets have become a serious security threat to the Internet users; they can bring down the entire system and disrupt Internet services. In a botnet, its Command and Control (C&C) channel, in which the botmaster disseminates commands to, and get response from bots, is a key element. Traditional botnets utilize the IRC or HTTP protocol to implement centralized C&C. Under this design, bots have to connect to central servers and even listen on certain channels to retrieve commands. Evidently, centralized C&C is vulnerable to a single point of failure, meaning that, whenever the central servers are identified and removed, the entire botnet will be disabled. To overcome this weakness, attackers have recently devised a decentralized C&C infrastructure exploiting the P2P protocol. A few noteworthy P2P botnets include Storm, Waledac and Conficker. Their P2P implementations are either based on an existing protocol (Storm utilized Kademila [12]) or completely customized.

The decentralized nature of P2P botnets inevitably makes their detection difficult. Approaches targeting centralized C&C structures [7], [4] become ineffective under this new structure in which a botmaster can join, issue commands and leave at any time at any place. Generic detection approaches [6], [20] relying on behavior monitoring and traffic correlation analysis are mostly applicable at a small scale such as in edge networks and do not scale well because they require analysis of vast

amounts of fine-grained information. In addition, if there are only a small number of bots in an edge network, detection based on bots’ coordination may fail due to the limited number of instances seen. Given the fact that sizes of existing botnets are in the order of hundreds of thousands, effective and efficient large-scale detection needs to function at a high level without requiring fine-grained information that can only be obtained locally. As a P2P botnet has a structured overlay and connectivity patterns different from other applications from a graph analysis perspective, we naturally consider detection at the Internet infrastructure level by assessing the impact imposed by a P2P botnet at various network components and measuring the effectiveness of detection at such places.

In this paper, we evaluate the feasibility of detecting large-scale P2P botnets with different network components in the Internet infrastructure. We construct three types of P2P-botnet overlays, map them to the corresponding AS (Autonomous System)-level underlays by inferring each overlay connection’s AS-path, and thus determine the PoP (Point of Presence) path and geographical router rendezvous (co-located routers in the infrastructure) each connection goes through. We then take a close look at each individual AS, PoP and router rendezvous based on graph analysis. In particular, we calculate a few P2P traffic classification metrics to see whether the portion of botnet connections observed by a single network component can be identified as P2P traffic. We would like to answer the following three questions from our analysis: (1) which network component is the best place for detection? (2) which P2P overlay structure can help hide the botnet traffic well? (3) what are the limitations of detection at the infrastructure level? Our main contribution lies in the thorough analysis of detection potential at the three infrastructure-level network components for three different P2P overlay topologies. Our analysis has led to three key observations. First, a small number of ASes can observe almost all overlay connections, but the AS-level detection is less practical. PoPs can capture a large fraction of connections but the number of monitoring points is limited. Router rendezvous strike a balance between detection capability and feasibility. Second, a botnet has to make a tradeoff between resilience/efficiency and the ability to evade detection. Third, the infrastructure-level detection is not a panacea for all large-scale botnets: it needs to be integrated with detection

schemes in edge networks to complete a detection picture.

The remainder of the paper is structured as follows. Section II describes related work. Section III details our methodology. Section IV presents our analysis results. The paper discusses a few challenges and concludes with Section V.

II. RELATED WORK

As botnets have become a major security threat, numerous approaches have been proposed for their detection and mitigation. Most of these approaches can only be applied to specific types of botnets, requiring in-depth understanding of the C&C profiles prior to their detection. A few generic approaches can detect different types of botnets regardless of the C&C structure based on network packet and flow analysis [6] or combined host and flow analysis [20]. These approaches are effective for small-scale networks, such as in a campus or an enterprise network, but do not scale to large networks, because they need to obtain fine-grained information.

Considering the fact that P2P botnets have structured overlay topologies, our approach takes a high-level view by exploiting structural properties derived from graph analysis, and is thus not limited by the availability of fine-grained information. In this regard, our work is closely related to graph-based traffic classification and analysis. Iliofotou *et al.* [10] proposed the use of Traffic Dispersion Graphs (TDGs) to monitor and classify network traffic. TDGs focus on network-wide interactions among hosts and show that graph features, such as the average degree and directionality, can be utilized to distinguish different applications. In our analysis, we adopt some of their metrics to determine whether the portion of traffic observed by a network component is P2P. BotGrep [14] analyzes structured graphs to locate bots by extracting P2P subgraphs from a communication graph containing background traffic. This approach was evaluated on entire botnet communication graphs and graphs with some edges removed. Unlike our work, BotGrep did not assess the botnet detection potential at different network components.

We are aware of two efforts on AS-level underlays mapped from P2P overlays. Rasti *et al.* [16] examined the global impact of the load imposed by a P2P overlay on the AS-level underlay. Their focus was on the effect of overlay on the underlay, while our work is concerned with whether the effect can be utilized for detection. Jelasity *et al.* [11] constructed a modified Chord [17] topology and showed that the visibility of P2P botnet traffic at any single AS is limited and not sufficient for detection. Our work differs from theirs, as we have simulated three P2P overlay topologies and observed the traffic not only at the AS-level but also at PoPs and router rendezvous, providing a more thorough analysis.

III. METHODOLOGY

A. Overview

We would like to achieve the following two goals. First, from a defender's perspective, we want to see how much of the botnet connections can be observed at a single network component and whether the respective communication graph

has P2P properties. Second, from an attacker's perspective, we want to study which P2P overlay topology is stealthy enough so that at a single network component the graph-level information is not sufficient for detection. Our methodology consists of four main steps. In the first step, we construct a P2P overlay topology based on simulation and learn which end-device talks to which, i.e., the overlay connections. In the second step, to map the overlay to the AS-level underlay, we associate the IP addresses of the two end-devices of a connection with the corresponding ASes and calculate the AS-level path between the two ASes. Given the AS paths, we then determine PoP-level paths and geographical router rendezvous paths. Knowing the paths of all connections, in the third step, we break down the connections on a per-AS, per-PoP, and per-router-rendezvous basis. We are especially interested in the top ASes, PoPs and router rendezvous ranked by the number of connections going through. In the last step, we inspect those top network components individually. As in [11], [16], we do not consider background traffic but focus only on the traffic coming from the P2P overlay, which is the best scenario, implying that if the P2P traffic cannot be identified under this situation, it will definitely not be captured when background traffic is present. We analyze several graph properties of the communication patterns at each top network component and determine whether it has the characteristics of P2P traffic.

B. Internet Infrastructure and End-Device Modeling

Before detailing the four main steps, we briefly describe the Internet infrastructure and end-device modeling, which lays a basis for our methodology. We use multiple real-world datasets to construct a realistic model of the US Internet infrastructure. Table I lists all data sources in the model construction. In total, 73,884,296 residential computers are generated in the entire US (except Hawaii and Alaska). The distribution of Internet access routers including dial-up, DSL and Cable is based on the market share of top US broadband companies and dial-up service aggregators, and how these access routers connect to the backbone topology at Internet PoP locations is derived from AS peering relationships. We refer interested readers to [19] for details of this modeling.

C. Overlay Topology Construction

In recent years, P2P overlays have become popular in botnet construction due to their decentralized nature. Many existing P2P overlays can be utilized to facilitate botnets' C&C. We construct three types of P2P overlays: a widely-used Kademlia [12], a modified Chord [17] and a simple ring structure. We will later compare the structural properties of these three overlay topologies at each network component, the results of which will be presented in Section IV. Next, we will introduce each P2P overlay followed by the way we construct the topology.

1) *Kademlia*: Kademlia is a Distributed-Hash-Table (DHT)-based P2P overlay protocol. Under this protocol, nodes are identified by node IDs and data items are identified by keys generated from a hash function; node IDs and keys

TABLE I
DATA SOURCES USED IN OUR INTERNET INFRASTRUCTURE AND END-DEVICE MODEL

Model component	Data sources
Backbone topology	Skitter dataset: http://www.caida.org/tools/measurement/skitter/ Alias clustering data from the iPlane project: http://iplane.cs.washington.edu/data/alias_lists.txt IP geolocation dataset: http://www.ip2location.com/
Internet Point of Presence	Telegeography co-location database: http://www.telegeography.com/
Internet end-devices	US census data: census-block population in each $250 \times 250 m^2$ grid in US for a 24-hour duration [13]
Internet access routers	Dial-up service aggregators per each zip code: http://www.findanisp.com Broadband ISP market share: http://www.leichtmanresearch.com/press/081108release.html DSL central office locations: the LERG (Local Exchange Routing Guide) dataset from Telcordia Cable company service locations: Dun & Bradstreet (D&B) dataset
Internet routing	BGP routing information from the University of Oregon Route Views Project: http://www.routeviews.org/ AS prefix sets: http://www.fixedorbit.com/ AS-level path inference: Qiu and Gao's algorithm [15]

are of the same length. Data items are stored in nodes whose IDs are close to data items' keys. The distance between two IDs, X and Y , is calculated by bitwise exclusive or (XOR) operation: $X \oplus Y$. To search a data item, a node queries its neighbors for nodes whose IDs are close to this data item's key. After getting responses from its neighbors, the node continues to query those nodes that are closer to the key. This iterative process repeats until no closer nodes can be found. The benefit of Kademlia is its resilience to disruptions. Even if a few nodes are shut down or removed, the network will still be able to function. Kad network is an implementation of Kademlia. A few major P2P file sharing networks adopt the Kad implementation, such as Overnet and eMule. The Storm botnet was built upon Overnet.

An ideal way to construct the botnet overlay topology is to collect traffic traces from a real network, such as the Storm botnet. Since the Storm botnet is decentralized (i.e., there are no central venues where all communications can be observed), traces captured from the Storm botnet fall into two categories each of which has its drawbacks. In the first category, the traffic data were collected from a single or a few vantage points. They can only provide partial views of the botnet. In the second category, snapshots of the network were taken by network crawlers. The snapshots contain information such as which IPs are alive or dead, but cannot tell which IP connects to which IP. To characterize the effectiveness of detection at the underlay, a full picture capturing the entire network's connections is indispensable, so we have to construct a Kad network by simulation.

We use a high-fidelity botnet simulator BotSim [8] which integrates a popular P2P client named *aMule* [1], an implementation of Kad. Considering the fact that simulating a large-scale botnet (100,000 bots) on a single or a few machines will take a prohibitively long time, our simulator was run on a distributed platform consisting of 400 machines, each with 2 Pentium III CPUs and 4Gb RAM. The simulator is a component of MIITS [18] which is built upon PRIME SSF [3], a distributed simulation engine. To make *aMule* work seamlessly on our simulator, several modifications were made to the original *aMule* code including intercepting time-related system calls and substituting them for simulated time function calls, and replacing socket API calls with network functions developed in MIITS. The rest of the code remains intact.

In a botnet, a majority of bots are compromised residential computers and not necessarily geographically close, so we have to take locations into account. Constrained by data availability, all bots in our simulation are in the US and their locations follow the geographical distribution of 73 million residential computers by state. The simulation of 100,000 bots executed for three days in simulation time. The output files log timestamps and connections in the network. We discarded the first day in which bots bootstrapped and the entire botnet stabilized, and kept the second and the third day for analysis. With the log files keeping track of which node talks to which other node and each node's state information, we need to obtain the IP address of each end-device to completely construct the overlay topology. For this, we randomly chose an end-device address from the state a bot resides in. This way, we created two Kad overlay topologies with 100,000 nodes, one day each.

2) *Modified Chord*: Chord is a DHT-based P2P protocol under which nodes form a ring structure. Each node has a predecessor and a successor and a few long range links. For example, there are a total of N nodes in the ring. Node i connects to nodes $(i-1) \bmod N$ and $(i+1) \bmod N$. It also connects to nodes $(i+2^k) \bmod N$ for $k = 1, 2, \dots, \log_2 N - 1$ to form long-range links. In [11], modifications to Chord are proposed so that it is difficult to detect through graph analysis at any single AS. The main modification is to create clusters in the ring each of which has $\log_2 N$ consecutive nodes. This way, nodes in the same cluster can share the same set of long-range links for routing. This topology is of interest to us because we want to see whether using a more realistic AS-path calculation algorithm can make a difference in detection and whether this topology can successfully hide itself at PoPs and router rendezvous as well. Since this modified Chord's topology is relatively simple, we constructed its overlay with 100,000 nodes directly based on its protocol without simulation. Following the same practice as in Kademlia, each end-device address is a random draw from the state a bot belongs to.

3) *Simple Ring*: We also consider the simplest case: each node has only two neighbors—a predecessor and a successor—to construct a ring structure. Presumably, this structure is stealthier and harder to detect than the modified Chord due to lack of connectivity at the overlay. We will verify this

presumption in later analysis. Similar to the modified chord, this overlay has 100,000 nodes constructed directly and the bots' locations follow the same geographical distribution.

D. Overlay to Underlay Mapping

1) *AS-Path*: Given all overlay connections, the next step is to map each connection to an AS-level path. Note that each end-device IP address is associated with an AS number and determining an AS-path of a connection is actually to determine the AS-path between two ASes. We use the AS-path inference algorithm in [15] for inter-domain routing. The key idea is to infer AS paths from existing BGP routing tables.

2) *PoP-Path*: A PoP is an access point to the Internet. It is a physical location owned by an ISP or located at Internet exchange points and co-location centers. The computation of a PoP-level path is based on the respective AS-level path. Given a pair of source and destination end device IPs, the algorithm first determines the AS-level path $AS_1AS_2 \dots AS_n$, then iteratively finds the shortest IP-level path between PoPs connecting every neighboring pair of ASes and finally maps the IP-level path to the PoP-level path. We refer interested readers to [19] for details of this algorithm.

3) *Router Rendezvous Path*: Given an IP-level path of a connection, the geographical router rendezvous along this path can be determined.

E. Traffic Breakdown

Since our work focuses on structural properties of the communication graph observed by a single network component, not the entire botnet overlay per se, we need to break down the overlay connections on a per AS, per PoP and per router rendezvous basis. We then rank the three types of network components by the number of connections going through, and take a close look at the graph properties observed at each of the top 10 ASes, PoPs, and router rendezvous, respectively, in our analysis.

F. Graph Analysis

After breaking down the traffic, we know all connections that traverse a particular AS, PoP and router rendezvous. We can then generate directed graphs in which bots are represented by vertices and connections among them are represented by edges. For simplicity, all edges carry the same weight. Graph metrics to determine whether the traffic is P2P are proposed in [9] and adopted to analyze the modified chord in [11]. In our analysis, we inspect the same set of features as in [11] for consistency. The features used to characterize P2P traffic include the number of weakly-connected components, size of the largest weakly-connected component, average node degree and InO (In Out) ratio. We introduce each of them as follows.

Number of Weakly-Connected Components: A weakly-connected component is a maximal subgraph of a directed graph such that in the subgraph replacing all of its directed edges with undirected edges produces a connected undirected graph. For effective detection, we expect a small number of weakly-connected components. As one can imagine, a large

number of connected components usually means small-size components that are less likely to exhibit typical P2P patterns.

Size of the Largest Weakly-Connected Component: This metric is meaningful to us because as pointed out in [10] the graph formed by a P2P network tends to be densely connected and have a large connected component including the majority of participating nodes.

Average Node Degree: This metric counts both the incoming and outgoing edges of a node, i.e., ignoring the directionality. A graph with a high average degree tends to be tightly-connected and P2P networks normally have high average node degrees.

InO Ratio: The metric calculates the percentage of nodes in the graph that have both incoming and outgoing edges. This metric is of interest because under client-server protocols such as HTTP and SMTP, clients usually initiate connections (outgoing edges) whereas servers normally accept connections (incoming edges). But nodes in P2P networks usually serve as both clients and servers so that P2P's InO is distinctively higher than others.

IV. ANALYSIS RESULTS

This section presents our analysis results. After constructing three different P2P overlay topologies, namely, Kad, the modified Chord and the simple ring, we examine their communication graphs at three types of network components. We conduct a graph analysis first at the AS-level, then the PoP-level and finally, the router-rendezvous-level, and show the graph features at the top 10 places of each level.

A. AS-Level Analysis

We first take a look at the AS-level graphs of three different topologies. Table II shows the Kad graph properties for day1 at top 10 ASes (for brevity we omit day2's result as they are similar to day1's), ranked by the number of unique connections going through. We map the AS numbers to ISPs using the AS-name lookup list [2]. Note that the traffic percentage at a single AS is calculated by the number of unique connections observed at that particular AS divided by the total number of unique connections in the entire overlay topology. Since one connection usually can be seen at more than one AS (this is why the first column of the table adds up to more than 100%), we count each connection only once while calculating the number of connections observed at multiple ASes altogether. Following such calculations, in day1, top 10 ASes aggregated together can observe 98.95%—almost all of the Kad overlay's unique connections. In particular, the top 1 AS (3356/Level3) alone can see two thirds of the overlay connections with all nodes (100000) in the picture. Even for ASes carrying fewer connections, they have at least 99937 nodes' connections traverse through. Most importantly, at each top AS, all nodes are weakly-connected with each other, forming one giant weakly-connected component. This property can facilitate detection because one single weakly-connected graph containing a majority of connections is more likely to demonstrate P2P characteristics and easier to get caught than

TABLE II
KAD AS-LEVEL

Kad Day1 ISP	AS	Percentage	# of Nodes	# of Edges	Avg Degree	# of Conn Comp	InO
Level3	3356	65.25%	100000	38192566	763.9	1	99.02%
AT&T	7018	35.33%	100000	20679083	413.6	1	99.02%
XO	2828	23.39%	100000	13691127	273.8	1	99.02%
Sprint	1239	8.32%	99983	4872140	97.5	1	99.01%
Verizon	19262	8.30%	100000	4859686	97.2	1	100.00%
Qwest	209	8.28%	100000	4848724	97.0	1	99.02%
NTT	2914	7.78%	99993	4556302	91.1	1	99.02%
BellSouth	6389	7.78%	100000	4554972	91.1	1	99.01%
AT&T	7132	6.78%	99995	3965587	79.3	1	100.00%
UUNET	701	5.38%	99937	3148400	63.0	1	88.13%

TABLE III
MODIFIED CHORD AS-LEVEL

ISP	AS	Percentage	# of Nodes	# of Edges	Avg Degree	# of Conn Comp	InO
Level3	3356	64.25%	80620	112431	2.8	9639	66.22%
AT&T	7018	38.09%	54272	66650	2.5	10534	51.62%
XO	2828	22.73%	36234	39784	2.2	7470	47.03%
Verizon	19262	9.43%	17365	16494	1.9	3726	37.01%
NTT	2914	8.09%	15339	14151	1.8	3384	34.45%
Sprint	1239	7.64%	14908	13366	1.8	3602	31.16%
Qwest	209	7.20%	14642	12594	1.7	3757	27.99%
AT&T	7132	7.13%	13849	12482	1.8	2956	33.29%
BellSouth	6389	6.82%	13486	11934	1.8	3080	30.47%
UUNET	701	6.27%	13900	10978	1.6	4305	16.41%

TABLE IV
SIMPLE RING AS-LEVEL

ISP	AS	Percentage	# of Nodes	# of Edges	Avg Degree	# of Conn Comp	InO
Level3	3356	64.76%	79327	64755	1.6	14522	63.31%
AT&T	7018	37.51%	51316	37511	1.5	13805	46.20%
XO	2828	22.81%	32148	22805	1.4	9343	41.88%
Verizon	19262	9.30%	13632	9297	1.3	4335	36.40%
NTT	2914	8.05%	11867	8046	1.3	3821	35.60%
Sprint	1239	7.53%	11604	7532	1.3	4072	29.82%
Qwest	209	7.36%	11494	7362	1.3	4132	28.10%
AT&T	7132	7.07%	10430	7066	1.3	3364	35.49%
BellSouth	6389	6.73%	10193	6728	1.3	3465	32.01%
UUNET	701	6.17%	10831	6166	1.1	4665	13.86%




Fig. 1. Modified Chord: 10 largest components at top 5 ASes




Fig. 2. Simple ring: 10 largest components at top 5 ASes

a disconnected graph with many connected components of small sizes. As suggested in [9], two metrics can characterize P2P traffic. One is a high average degree (larger than 2.8), and the other is a high InO ratio (large than 1%). At all top ASes, the average degrees and InO values are high enough for P2P classification: the lowest value of average degree is 63.0 and that of InO is 88.13%. Thus, all top AS venues have high visibility of Kad’s overlay that demonstrates typical P2P patterns, sufficient for detection.

Table III shows graph features of the modified Chord at top 10 ASes. Compared to Kad, top 10 AS numbers remain same but their ranks change a bit. They in total observe 99.61%, an enormous fraction of overlay connections and the top 1 AS is still 3356 witnessing 64.25% of all connections. Note that the AS observing the most can see 80620 while the one observing the least can only see 13900 nodes. As for the number of connected components, to the contrary of Kad, each AS’s graph is not well connected and has thousands of connected components. Figure 1 shows in log scale the sizes of 10 largest weakly-connected components at top 5 ASes. Top 1 AS 3356’s largest component has 36532 nodes but all other components are very small containing 15 nodes or so. Top 2 AS 7018 has two large components with 8729 and 7506 nodes respectively and sizes of other components drop significantly. The component sizes remain stable at other ASes, all in the order of hundreds. Due to the relatively sparse structure of the modified Chord, unsurprisingly, the average degree at each AS is low—from 2.8 to 1.6, though the InO values are high—from 66.22% to 16.41%. Taking all metrics into account, AS 3356 is able to detect the P2P overlay since it can see a large portion of the overlay with typical P2P patterns, if not the entire one. If we relax the average degree threshold a bit, AS 7018 may also be a good venue to make detection efforts considering the two large connected components. We think it is hard for the rest of the ASes to do so due to their relatively fragmented views. Note that our observations on the modified Chord are slightly different from those in [11] which concludes that even at the most central (top) ASes the average degrees are less than 2 and connected components are mostly of size 2 and 3 with the maximal containing 29 nodes. This difference may be attributed to the way of mapping the overlay to the underlay: they make the number of overlay nodes in each AS proportional to the size of the AS whereas we consider the geographical distribution of nodes. In addition, our AS-path inference algorithm is also different from theirs: they assume shortest paths while our AS-paths are derived from real-world BGP routing tables.

When it comes to the simple ring structure (Table IV), the top AS numbers do not change, and their ranks are the same as those for the modified Chord. 99.62% of overlay connections traverse through top 10 ASes. Though the top 1 AS 3356 can see 64.76% of overlay connections, the number of nodes visible (79327) are more than the number of edges (64755), resulting in a great number of connected components (14522) and small component sizes. As seen in Figure 2, 3356’s largest component only has 34 nodes. We also verify that a majority of

3356 connected components have fewer than 10 nodes. The average degrees are all below 2, which is expected because each node only has a predecessor and a successor so that the average degree of the entire graph is only 2. Even though the InO values are high enough, detection based on scattered information at a single AS is difficult.

B. PoP-Level Analysis

At the PoP level, we also present graph features at each top PoP of three P2P structures. PoPs are represented by ID numbers and ranked by the number of unique connections going through as well. In Kad’s case (Table V), the top 10 PoPs account for 80.88% of overlay connections, a slight drop compared to that observed at top 10 ASes which can see more than 98%. This makes sense because PoPs, normally as traffic exchange points, are not able to see intra-domain traffic taking place within ASes. The top PoP 74 alone is able to observe 53.78% of all connections. Similar to the AS-level, not only almost all nodes (more than 99975) can be seen at each top PoP, but also they are weakly connected forming one single component. The average degrees and InO ratios are well above the P2P classification thresholds.

In the case of the modified Chord (Table VI), top PoPs are almost the same as those of Kad and only the ranks change, taking up 80.29% of overlay connections aggregately. 74 is still the top 1 PoP observing 54.07% of total connections containing 77488 nodes, but all other PoPs observe fewer than 20000 nodes. As for sizes of weakly connected components, shown in Figure 3 in log scale, PoP 74’s largest component is of size 23153 and others are quite small. Other PoPs’ component sizes are fewer than 300. Given all these statistics, if the average degree threshold can be relaxed a bit, PoP 74 can be a good place for detection.

In the case of the simple ring (Table VII), the PoP numbers are exactly the same as those of modified Chord. 89.25% of overlay connections reach top 10 PoPs with 54.51% traversing PoP 74. Despite the fact that half of overlay connections can be observed at PoP 74, similar to the AS-Level, the number of edges is smaller than the number of nodes. The largest component of PoP 74 is very small containing 22 nodes (Figure 4). It is the same case for all other top PoPs. Though InO values are moderate, low average degrees and a good many small connected components can prevent the P2P structure from being captured at any PoP.

C. Router-Rendezvous-Level Analysis

At the router-rendezvous level, we present results in the same way. Router rendezvous are denoted by ID numbers and ordered by the number of unique overlay connections observed. For the Kad structure, as shown in Table VIII, the top 10 router rendezvous see 89.75% of total connections. The top 1 router rendezvous number 2 is reached by 68.77% of all connections. A majority of nodes (more than 98858) appear in the graph as one giant component at each top router rendezvous. In addition, high average degrees and InO values make detection feasible. Let us take a look at the

TABLE V
KAD POP-LEVEL

Kad Day1 PoP	Percentage	# of Nodes	# of Edges	Avg Degree	# of Conn Comp	InO
74	53.78%	100000	31479094	629.6	1	100.00%
7	10.29%	100000	6024939	120.5	1	99.94%
435	8.27%	100000	4837622	96.8	1	98.50%
11	8.14%	99998	4763870	95.3	1	99.86%
128	7.77%	99981	4550316	91.0	1	99.52%
282	7.37%	99995	4315967	86.3	1	100.00%
4	7.27%	99977	4257513	85.2	1	99.73%
267	6.72%	99992	3934199	78.7	1	100.00%
291	6.26%	99975	3661420	73.2	1	100.00%
295	6.25%	99997	3658911	73.2	1	99.97%

TABLE VI
MODIFIED CHORD POP-LEVEL

PoP	Percentage	# of Nodes	# of Edges	Avg Degree	# of Conn Comp	InO
74	54.07%	77488	94629	2.4	16735	48.00%
7	9.27%	19927	16222	1.6	6095	21.91%
267	7.99%	14764	13981	1.9	3092	34.80%
11	7.98%	17225	13957	1.6	5334	18.75%
128	7.46%	17169	13058	1.5	5673	17.39%
4	7.25%	15962	12686	1.6	4834	20.36%
435	6.94%	13649	12151	1.8	3067	32.38%
282	6.81%	13677	11913	1.7	3184	31.41%
291	6.36%	12433	11137	1.8	2683	32.68%
295	5.84%	11877	10228	1.7	2803	29.32%

TABLE VII
SIMPLE RING POP-LEVEL

PoP	Percentage	# of Nodes	# of Edges	Avg Degree	# of Conn Comp	InO
74	54.51%	75999	54506	1.4	21493	43.44%
7	9.40%	16165	9400	1.2	6765	16.30%
11	7.78%	13648	7779	1.1	5869	13.99%
128	7.63%	13765	7631	1.1	6134	10.88%
267	7.52%	11079	7521	1.4	3558	35.77%
4	7.31%	12505	7305	1.2	5200	16.83%
435	7.13%	10568	7127	1.3	3441	34.88%
282	7.08%	10587	7078	1.3	3509	33.71%
291	6.37%	9392	6373	1.4	3019	35.71%
295	5.77%	8829	5774	1.3	3055	30.80%




Fig. 3. Modified Chord: 10 largest components at top 5 PoPs




Fig. 4. Simple ring: 10 largest components at top 5 PoPs




Fig. 5. Modified Chord: 10 largest components at top 5 locations




Fig. 6. Simple ring: 10 largest components at top 5 locations

TABLE VIII
KAD ROUTER-RENDEZVOUS-LEVEL

Kad Day1 Router	Percentage	# of Nodes	# of Edges	Avg Degree	# of Conn Comp	InO
2	68.77%	100000	40251799	805.0	1	100.00%
2164	14.91%	99959	8728267	174.6	1	98.96%
12	11.90%	99997	6967203	139.3	1	84.22%
98	11.75%	100000	6874621	137.5	1	100.00%
222	9.26%	100000	5419174	108.4	1	99.99%
8919	8.30%	100000	4855632	97.1	1	98.50%
745	7.82%	99997	4579803	91.6	1	99.85%
82	7.33%	99978	4288889	85.8	1	99.74%
47	6.99%	98858	4090556	82.8	1	92.32%
88	6.67%	99997	3904395	78.1	1	99.71%

TABLE IX
MODIFIED CHORD ROUTER-RENDEZVOUS-LEVEL

Router	Percentage	# of Nodes	# of Edges	Avg Degree	# of Conn Comp	InO
2	68.76%	88913	120337	2.7	13816	59.33%
2164	15.00%	29299	26245	1.8	8964	25.60%
12	11.57%	23682	20247	1.7	7629	20.07%
98	11.33%	21641	19821	1.8	5586	31.07%
222	8.73%	17779	15280	1.7	4771	27.68%
745	7.59%	16673	13275	1.6	5286	17.14%
82	7.29%	16133	12758	1.6	4926	19.98%
8919	6.94%	13649	12151	1.8	3067	32.38%
88	6.26%	12913	10962	1.7	3364	25.96%
57	6.16%	13606	10784	1.6	4029	19.42%

TABLE X
SIMPLE RING ROUTER-RENDEZVOUS-LEVEL

Router	Percentage	# of Nodes	# of Edges	Avg Degree	# of Conn Comp	InO
2	68.89%	88161	68885	1.6	19276	56.27%
2164	15.12%	25513	15122	1.2	10391	18.54%
12	11.35%	20126	11351	1.1	8775	12.80%
98	11.28%	17720	11275	1.3	6445	27.26%
222	8.93%	14243	8933	1.3	5310	25.44%
745	7.42%	13218	7419	1.1	5799	12.26%
82	7.36%	12653	7356	1.2	5297	16.27%
8919	7.13%	10568	7127	1.3	3441	34.88%
88	6.10%	9762	6102	1.3	3660	25.02%
47	6.06%	9669	6061	1.3	3608	25.37%

modified Chord at the router-rendezvous level (Table IX). There is one new router rendezvous in the top 10 list that does not appear in that of Kad’s and the ranks of the two lists are quite similar. Top 10 router rendezvous carry 89.96% of total connections and the top 1 router rendezvous is still 2 accounting for 68.76% of connections including 88913 nodes. As for the sizes of weakly connected components, the trend does not differ much from that at the AS- and PoP-level. The top 1 router rendezvous’s largest connected component is of a big size—35943 nodes (Figure 5 in log scale) and other components have small sizes (fewer than 15 nodes). With a distinctive average degree and a high InO value, this router rendezvous is a reasonable venue for capturing the modified Chord. Finally, for the simple ring structure (Table X), the top router rendezvous list is the same as that of Kad. Top 10 router rendezvous observe 80.54% of overlay connections and router 2 sees 68.89% of them. With more nodes than edges at each top router rendezvous, it is difficult to get a good view of the overlay. Similar to AS- and PoP-level, the top 1 router rendezvous’s largest component contains 33 nodes.

The average degrees are unsurprisingly low, insufficient for detection.

D. Insights from Analysis

From the above analysis, we have several key observations. First, the visibility of Kad’s overlay and structure at the top places of all levels is good enough for detection; the modified Chord’s P2P characteristics can be captured by a few top locations but not all; and the information of the hypothetical simple ring’s topology at all levels is quite fragmented and hardly useful for detection. From the attacker’s viewpoint, in terms of efficiency, Kad has the most efficient routing: contacting $O(\log N)$ nodes during a search (where N is the size of the network); the modified Chord can achieve $O(\log^2 N)$ hops; and the simple ring is the worst, requiring $O(N)$ steps. From resilience’s perspective, the Kad network is shown to be robust to a few types of mitigation strategies such as cutting off random nodes and removing peers learnt from bots’ peer lists [5]; the simple ring structure is evidently fragile—removing a couple of nodes can disconnect the overlay; and

the modified Chord structure hits the middle ground: not as resilient as Kad but better than the simple ring. We believe that, while constructing a P2P botnet, the attacker needs to strike a balance between resilience or efficiency and the ability to evade detection. Although the simple ring can hide its traffic well at various network components, to build upon this structure the botnet has to compromise resilience and C&C efficiency. The modified Chord makes a tradeoff though its structural properties cannot be concealed at some locations. Kad was successfully utilized by the Storm botnet, but given our detection strategy, to use it for a future botnet, the attacker has to come up with techniques to mask its P2P patterns.

Second, from detection's perspective, AS-level provides better overlay views than PoP- and router-rendezvous-level do, but is less practical than the other two for actual detection deployment. Since AS is only a logical concept, capturing all connections within one AS requires collaboration and synchronization among multiple physical devices at different geographical locations, which renders it highly impractical. From our analysis, we can see that at the PoP-level, detecting Kad and the modified Chord is very likely though the latter is only visible to the top 1 PoP. Compared to ASes and router rendezvous, PoPs observe less traffic due to the invisibility of traffic within ASes (intra-domain traffic). Moreover, the number of PoPs is small so that the points of monitoring are limited. Among the three, router rendezvous make a tradeoff. Their detection capabilities are comparable to PoPs' and they can observe intra-domain traffic with more monitoring points available, making detection more feasible.

V. DISCUSSIONS AND CONCLUSION

For actual implementation of botnet detection in the Internet infrastructure, there remain a few challenges to be addressed. First, since our techniques are applied at the structure-level via graph analysis, they will also identify regular P2P file-sharing topologies. To avoid misclassifying such regular P2P networks as botnets, we can perform preprocessing such as flow filtering and clustering [9] based on known patterns of regular P2P networks such as the port numbers. Also, bots identified locally in edge networks are helpful as their presence in a communication graph makes other nodes suspicious as well, so our approach may need assistance from detection mechanisms at the edge to further confirm that a graph is indeed formed by a botnet. However, if the botnet is immersed into an existing regular P2P network, detecting it solely by graph analysis at the infrastructure level would be challenging and other information is thus needed for effective detection. Second, our models on the Internet infrastructure are abstracted from real-world datasets, so the accuracy depends on how well the datasets characterize the behavior and the state of the Internet, which could be error-prone. Moreover, some datasets may be outdated and may not reflect the current state of the Internet due to its fast-evolving nature. Lastly, in the presence of a huge traffic volume, some connections could not be captured due to sampling. For densely-connected topologies such as Kad, it may not be a problem. But for the modified Chord and simple

ring's cases, it will complicate the detection. We plan to dig deeper into this issue.

To sum up, as P2P structures become a popular choice for recent botnets, especially large-scale ones, detection mechanisms have to keep up with this change and identify bots in an efficient and effective manner. In this paper, we propose detection of P2P botnets at a high-level—the infrastructure-level by analyzing their structural properties from a graph perspective. We find that detection at any of the three network components has its advantages and drawbacks. Overall, router-rendezvous-level detection is able to strike a balance between detection capability and feasibility. Also, a botnet needs to make a tradeoff between resilience and stealthiness.

ACKNOWLEDGMENTS

The work reported in this paper was supported in part by the ONR under Grant No. N000140911042 and the NSF under Grant No. CNS 0905143.

REFERENCES

- [1] "aMule," <http://www.amule.org/>.
- [2] "AS Names," <http://bgp.potaroo.net/cidr/autnums.html>.
- [3] "PRIME SSF," <https://www.primessf.net/bin/view/Public>.
- [4] J. R. Binkley and S. Singh, "An algorithm for anomaly-based botnet detection," in *Proceedings of the 2nd conference on Steps to Reducing Unwanted Traffic on the Internet*, 2006.
- [5] C. R. Davis, S. Neville, J. M. Fernandez, J.-M. Robert, and J. McHugh, "Structured peer-to-peer overlay networks: Ideal botnets command and control infrastructures?" in *Proceedings of ESORICS'08*.
- [6] G. Gu, R. Perdisci, J. Zhang, and W. Lee, "BotMiner: Clustering analysis of network traffic for protocol- and structure-independent botnet detection," in *Proceedings of the USENIX Security Symposium*, 2008.
- [7] G. Gu, J. Zhang, and W. Lee, "BotSniffer: Detecting botnet command and control channels in network traffic," in *Proc. of NDSS*, 2008.
- [8] D. Ha, G. Yan, S. Eidenbenz, and H. Ngo, "On the effectiveness of structural detection and defense against p2p-based botnets," in *Proceedings of DSN 2009*.
- [9] M. Iliofotou, H. chul Kim, P. Pappu, M. Faloutsos, M. Mitzenmacher, and G. Varghese, "Graph-based p2p traffic classification at the internet backbone," in *Proceedings of IEEE Global Internet Symposium*, 2009.
- [10] M. Iliofotou, P. Pappu, M. Faloutsos, M. Mitzenmacher, S. Singh, and G. Varghese, "Network monitoring using traffic dispersion graphs (tdgs)," in *Proceedings of IMC 2007*.
- [11] M. Jelasity and V. Bilicki, "Towards automated detection of peer-to-peer botnets: On the limits of local approaches," in *Proc. of LEET*, 2009.
- [12] P. Maymounkov and D. Mazieres, "Kademlia: A peer-to-peer information system based on the xor metric," in *Proceedings of IPTPS*, 2001.
- [13] T. N. McPherson and M. J. Brown, "Estimating daytime and night-time population distributions in u.s. cities for emergency response activities," The American Meteorological Society.
- [14] S. Nagaraja, P. Mittal, C.-Y. Hong, M. Caesar, and N. Borisov, "Botgrep: Finding p2p bots with structured graph analysis," in *Proceedings of 19th USENIX Security Symposium*, 2010.
- [15] J. Qiu and L. Gao, "As path inference by exploiting known as paths," in *Proceedings of IEEE GLOBECOM*, 2005.
- [16] A. H. Rasti, R. Rejaie, and W. Willinger, "Characterizing the global impact of p2p overlays on the as-level underlay," in *Proceedings of PAM 2010*.
- [17] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan, "Chord: A scalable peer-to-peer lookup service for internet applications," in *Proceedings of ACM SIGCOMM*, 2001.
- [18] R. Waupotitsch, S. Eidenbenz, J. Smith, and L. Kroc, "Multi-scale integrated information and telecommunications system (MIITS): First results from a large-scale end-to-end network simulator," in *Proceedings of the Winter Simulation Conference*, 2006.
- [19] G. Yan, S. Eidenbenz, S. Thulasidasan, P. Datta, and V. Ramaswamy, "Criticality analysis of internet infrastructure," *Computer Networks*, vol. 54, no. 7, 2010.
- [20] Y. Zeng, X. Hu, and K. G. Shin, "Detection of botnets using combined host- and network-level information," in *Proceedings of DSN*, 2010.