Detecting Malware with Graph-based Methods: Traffic Classification, Botnets, and Facebook Scams

Michalis Faloutsos
U. New Mexico
Albuquerque, New Mexico, USA
michalis@cs.unm.edu

ABSTRACT

In this talk, we highlight two topics on security from our lab. First, we address the problem of Internet traffic classification (e.g. web, filesharing, or botnet?). We present a fundamentally different approach to classifying traffic that studies the network wide behavior by modeling the interactions of users as a graph. By contrast, most previous approaches use statistics such as packet sizes and inter-packet delays. We show how our approach gives rise to novel and powerful ways to: (a) visualize the traffic, (b) model the behavior of applications, and (c) detect abnormalities and attacks. Extending this approach, we develop ENTELECHEIA, a botnet-detection method. Tests with real data suggests that our graph-based approach is very promising.

Second, we present, MyPageKeeper, a security Facebook app, with 13K downloads, which we deployed to: (a) quantify the presence of malware on Facebook, and (b) protect end-users. We designed MyPageKeeper in a way that strikes the balance between accuracy and scalability. Our initial results are scary and interesting: (a) malware is widespread, with 49% of our users are exposed to at least one malicious post from a friend, and (b) roughly 74% of all malicious posts contain links that point back to Facebook, and thus would evade any of the current web-based filtering approaches.

Categories and Subject Descriptors

Computer Networks [**Security**]: Miscellaneous

Keywords

Malware detection, traffic classification, botnets, socware

1. TOPICS COVERED

In this work, we address two main topics that relate to detecting malicious activities and cybercrime.

Part 1. Detecting malicious activities at the network level.

Monitoring network traffic and detecting unwanted applications has become a challenging problem, since many applications obfuscate their traffic using arbitrary port numbers or payload encryption. Apart from some notable exceptions, most traffic monitoring tools follow two types of approaches: (a) keeping traffic statistics such as packet sizes and inter-arrivals, flow counts, byte volumes, etc., or (b) analyzing packet content. In this work, we propose the use of Traffic Dispersion Graphs (TDGs) as a powerful way to monitor, analyze, and visualize network traffic [6, 4, 5, 1, 2]. TDGs model the social behavior of hosts ("who talks to whom"), while the edges can be defined to represent different interactions (e.g. the exchange of a certain number or type of packets). With the introduction of TDGs, we are able to harness the wealth of tools and graph modeling techniques from a diverse set of disciplines. First, we fully explore the abilities of TDGs as an intuitive and visually powerful tool. Second, we demonstrate their usefulness in application classification and intrusion detection solutions. Finally, we provide a hardware-aware design and implementation for TDG-based techniques. We conclude that TDGs are powerful, useful, and can be implemented efficiently in hardware. They constitute a promising new chapter for network monitoring techniques.

Specifically, we address the following questions. Is there link homophily in the application layer traffic? If so, can it be used to accurately classify traffic in network trace data without relying on payloads or properties at the flow level? Our research shows P2P that the answers to both of these questions are affirmative in real network trace data. Specifically, we define link homophily to be the tendency for flows with common IP hosts to have the same application (P2P, Web, etc.) compared to randomly selected flows. The presence of link homophily in trace data provides us with statistical dependencies between flows that share common IP hosts. We utilize these dependencies to classify application layer traffic without relying on payloads or properties at the flow level. In particular, we introduce a new statistical relational learning algorithm, called Neighboring Link Classifier with Relaxation Labeling (NLC+RL). Our algorithm has no training phase and does not require features to be constructed. All that it needs to start the classification process is traffic information on a small portion of the initial flows, which we refer to as seeds. In all our traces, NLC+RL achieves above 90% accuracy with less than 5% seed size; it is robust to errors in the seeds and various seed-selection biases; and it is able to accurately classify challenging traffic such as P2P with over 90% Precision and Recall.

As a showcase of the power of graph-based methods, we apply it on botnet detection. Detecting botnets is a critical need for securing one's network and the Internet at large. Despite significant efforts, the problem of botnet detection is still unresolved, especially, when one wants to detect: (a) decentralized or peer-to-peer botnets, (b) botnets that are in

WWW 2013 Companion, May 13–17, 2013, Rio de Janeiro, Brazil.
ACM 978-1-4503-2038-2/13/05.

a non-active period known as the â "Waiting" stage, and (c) polymorphic bots that evade signature detection. We propose a graph-based approach called ENTELECHEIA that addresses all three challenges above successfully.

Part II. Detecting social malware in Facebook.

Online social networks (OSNs) have become the new vector for cybercrime, and hackers are finding new ways to propagate spam and malware on these platforms, which we refer to as socware. As we show here, socware cannot be identified with existing security mechanisms (e.g., URL blacklists), because it exploits different weaknesses and often has different intentions. We present MyPageKeeper, a Facebook application that we have developed to protect Facebook users from socware; thus far, over 12K users have installed MyPageKeeper [7, 3].

Our work makes three main contributions. First, to enable protection of users at scale, we design an efficient socware detection method which takes advantage of the social context of posts. We find that our classifier is both accurate (97% of posts flagged by it are indeed socware and it incorrectly flags only 0.005% of benign posts) and efficient (it requires 46 ms on average to classify a post). Second, we show that socware significantly differs from traditional email spam or web-based malware. For example, website blacklists identify only 3% of the posts flagged by MyPageKeeper, while 26% of flagged posts point to malicious apps and pages hosted on Facebook (which no current antivirus or blacklist is designed to detect). Third, we quantify the prevalence of socware by analyzing roughly 40 million posts over four months; 49% of our users were exposed to at least one socware post in this period. Finally, we identify a new type of parasitic behavior, which we refer to as "Like-as-a-Service", whose goal is to artificially boost the number of "Likes" of a Facebook page.

Going one step further, we study applications that enable the spread of socware. With 20 million installs a day, third-party apps are a major reason for the popularity and addictiveness of Facebook. Unfortunately, hackers have realized the potential of using apps for spreading malware and spam. The problem is already significant, as we find that at least 13% of apps in our dataset are malicious. So far, the research community has focused on detecting malicious posts and campaigns.

We ask the question: given a Facebook application, can we determine if it is malicious? Our key contribution is in developing FRAppE (**F** acebook **R** igorous **App**lication **E**valuator)â arguably the first tool focused on detecting malicious apps on Facebook. To develop FRAppE, we use information gathered by observing the posting behavior of 111K Facebook apps seen across 2.2 million users on Facebook. First, we identify a set of features that help us distinguish malicious apps from benign ones. For example, we find that malicious apps often share names with other apps, and they typically request fewer permissions than benign apps. Second, leveraging these distinguishing features, we show that FRAppE can detect malicious apps with 99.5% accuracy, with no false positives and a low false negative rate (4.1%). Finally, we explore the ecosystem of malicious Facebook apps and identify mechanisms that these apps use to propagate. Interestingly, we find that many apps collude and support each other; in our dataset, we find 1,584 apps enabling the viral propagation of 3,723 other apps through their posts. Long-term, we see FRAppE as a step towards creating an independent watchdog for app assessment and ranking, so as to warn Facebook users before installing apps.

2. BIO

Michalis Faloutsos is a faculty member at the Computer Science Dept. at the University New Mexico. He got his bachelor's degree at the National Technical University of Athens and his M.Sc and Ph.D. at the University of Toronto. His interests include, Internet protocols and measurements, peer-to-peer networks, network security, BGP routing, and ad-hoc networks. With his two brothers, he co-authored the paper on power-laws of the Internet topology, which received the ACM SIGCOMM Test of Time award. His work has been supported by many NSF and military grants, for a cumulative total of more than $8 million. Several recent works have been widely cited in popular printed and electronic press such as slashdot, ACM Electronic News, USA Today, and Wired. Most recently he has focused on the classification of traffic and web-security, and co-founded a cyber-security company founded in 2008, offering services as www.stopthehacker.com, which received two SBIR grants from the National Science Foundation, and institutional funding in Dec 2011.

3. REFERENCES

[1] B. Gallagher, M. Iliofotou, T. Eliassi-Rad, and M. Faloutsos. Homophily in application layer and its usage in traffic classification. In *IEEE INFOCOM*, San Diego, CA, USA, March 2010.

[2] H. Hang, X. Wei, M. Faloutsos, and T. Eliassi-Rad. Entelecheia: Detecting p2p botnets in their waiting stage. In *IFIP Networking*, 2013.

[3] T.-K. Huang, M. S. Rahman, H. Madhyastha, and M. Faloutsos. An analysis of socware cascades in online social networks. In *WWW'13*, 2013.

[4] M. Iliofotou, M. Faloutsos, and M. Mitzenmacher. Exploiting dynamicity in graph-based traffic analysis: Techniques and applications. In *ACM CoNEXT*, December 2009.

[5] M. Iliofotou, B. Gallagher, T. Eliassi-Rad, G. Xie, and M. Faloutsos. Profiling-by-association: A resilient traffic profiling solution for the internet backbone. In *ACM CoNEXT*, Philadelphia, CA, USA, December 2010.

[6] M. Iliofotou, P. Pappu, M. Faloutsos, M. Mitzenmacher, S. Singh, and G. Varghese. Network monitoring using traffic dispersion graphs (tdgs). In *In ACM Internet Measurement Conference (IMC)*, pages 315–320, 2007.

[7] M. S. Rahman, T.-K. Huang, H. Madhyastha, and M. Faloutsos. Efficient and scalable socware detection in online social networks. In *USENIX Security*, 2012.