[1] E. P. Freire, A. Ziviani, and R. M. Salles. Detecting VoIP calls hidden in web traffic. *IEEE Transactions on Network and Service Management*, 5(4):204-214, Dec. 2008. [bib | DOI]

Peer-to-peer (P2P) voice over IP (VoIP) applications (e.g. Skype or Google Talk) commonly use Web TCP ports (80 or 443) as a fallback mechanism to delude restrictive firewalls. This strategy renders this kind of traffic quite difficult to be detected by network managers. To deal with this issue, we propose and evaluate a method to detect VoIP calls hidden in Web traffic. We validate our proposal considering both Skype and Google Talk generated traffic by using real-world experimental data gathered at a commercial Internet Service Provider (ISP) and an academic institution. Our experimental results demonstrate that our proposed method achieves a performance of around 90% detection rate of VoIP calls hidden in Web traffic with a false positive rate of only 2%, whereas a 100% detection rate is achieved with a false positive rate limited to only 5%. We also evaluate the feasibility of applying our proposal in real-time detection scenarios.

Keywords: ISP;Internet service provider;TCP/IP protocol;VoIP hidden calls detection;Web traffic;fallback mechanism;peer-to-peer application;voice over IP;Internet telephony;peer-to-peer computing;telecommunication traffic;transport protocols;

[2] E. P. Freire, A. Ziviani, and R. M. Salles. On metrics to distinguish Skype flows from HTTP traffic. *Journal of Network and Systems Management*, 17:53-72, 2009. [bib | DOI]

Skype is a Voice over IP (VoIP) Internet application that is gaining huge popularity in recent years. A key point to Skype popularity is its capability to dynamically adapt itself to operate behind firewalls or network proxies. A common way adopted by Skype to delude these network devices is to use port 80, normally expected to comprise HTTP traffic. In this paper, we propose metrics and investigate statistical tests intended to clearly distinguish Skype flows from HTTP traffic. We validate our study using real-world experimental datasets gathered at a commercial Internet Service Provider (ISP). Our experimental results suggest that the proposed methodology may be seen as a promising building block towards a system to detect general protocol anomalies in HTTP traffic.

[3] J. L. Garcia-Dorado, J. A. Hernandez, J. Aracil, J. E. L. de Vergara, and S. Lopez-Buedo. Characterization of the busy-hour traffic of ip networks based on their intrinsic features. *Computer Networks*, 55(9):2111-2125, 2011. [bib | DOI | http]

Internet traffic measurements collected during the busy hour constitute a key tool to evaluate the operation of networks under the heaviest-load case scenarios, and further provide a means to network dimensioning and capacity planning. In this light, this study provides a throughout analysis of the busy-hour traffic measurements of an extensive set of universities, regional networks, and Internet exchange points collected from the Spanish Research and Education Network, RedIRIS. After showing that the traffic volumes observed in the busy hour over time can be modeled by a white Gaussian process, this work takes one step further and examines the influence of the networks' intrinsic features, mainly population size and access link capacity, on the busy-hour traffic. Well-known statistical methodologies, such as ANOVA and ANCOVA, show that the network size in terms of number of users justifies most of the busy-hour traffic information. We further provide a linear-regression model that adjusts the amount of traffic that each network user contributes to the busy-hour traffic mean values, with a direct application to the problem of link capacity planning of IP networks.

Keywords: Internet traffic busy hour, Capacity planning, Bandwidth demands, University access link, ANCOVA, Network intrinsic features

[4] A. Moore, M. Crogan, A. W. Moore, Q. Mary, D. Zuev, D. Zuev, and M. L. Crogan. Discriminators for use in flow-based classification. Technical Report RR-05-13, Dept. of Computer Science, Queen Mary University of London, Aug. 2005. [bib]

Any assessment of classification techniques requires data. This document describes sets of data intended to aid in the assessment of classification work. A number of data sets are described; each data set consists a number of objects, and each object is described by a group of features (also referred to as discriminators). Leveraged by a quantity of hand-classified data, each object within each data set represents a single flow of TCP packets between client and server. The features for each

object consist of the (application-centric) classification derived elsewhere and a number of features derived as input to probabilistic classification techniques. In addition to describing the features, we also provide information allowing interested parties to retrieve these data sets for use in their own work. The data sets contain no site-identifying information; each object is only described by a set of statistics and a class that defines the causal application.

[5] D. Nechay. Controlling false alarm/discovery rates in online internet traffic classification. Master's thesis, Department of Electrical and Computer Engineering, McGill University, 2010. [bib | http]

Classifying Internet traffic flows online into applications or broader classes without in- specting the packet payloads or without relying on port numbers has become a necessity for network operators. The operators can use this information to monitor their networks and provide per-class quality of service. There has been a great deal of research done on Internet traffic classification recently and numerous techniques have been proposed. While the current techniques can obtain a high accuracy classifying Internet traffic, providing performance guarantees for particular classes of interest has never been addressed. In this thesis, we provide two novel types of online Internet traffic classifiers that can provide per- formance guarantees on the false alarm and false discovery rates, respectively. These guar- antees can be for an entire class (class-wise) or between two classes (pair-wise). Controlling false alarm rates is well-suited for application prioritization (i.e. prioritizing time-sensitive applications like VoIP over HTTP) whereas controlling false discovery rates is better suited for blocking or rate-limiting a targeted class of traffic (i.e. Peer-to-Peer). The classifier that provides false alarm rate guarantees is based on a Neyman-Pearson classification framework while the classifier that provides false discovery rate guarantees is based on the Learning to Satisfy (LSAT) framework. Both of these classifiers are implemented using a machine learning technique, namely, the 2-nu Support Vector Machine (SVM). Moreover, all pre- vious work done with these two statistical methodologies focused on binary classification only; we extend these statistical methodologies to a multi-class setting. In addition to the regular application classification problem, we also present preliminary work on a binary LSAT classifier that can detect, after the reception of only a handful of packets, whether a flow will be large, as defined by a network operator. This large flow detector can act as a preprocessor for regular application classifiers. By allowing only large flows to pass to the classifier, this allows the classifier to focus on the more resource-intensive flows. We validated our Internet traffic classifiers by testing our approaches using data provided by an ISP.

[6] L. Peterson, T. Anderson, D. Culler, and T. Roscoe. A blueprint for introducing disruptive technology into the internet. *SIGCOMM Comput. Commun. Rev.*, 33:59-64, Jan. 2003. [bib | DOI]

This paper argues that a new class of geographically distributed network services is emerging, and that the most effective way to design, evaluate, and deploy these services is by using an overlay-based testbed. Unlike conventional network testbeds, however, we advocate an approach that supports both researchers that want to develop new services, and clients that want to use them. This dual use, in turn, suggests four design principles that are not widely supported in existing testbeds: services should be able to run continuously and access a slice of the overlay's resources, control over resources should be distributed, overlay management services should be unbundled and run in their own slices, and APIs should be designed to promote application development. We believe a testbed that supports these design principles will facilitate the emergence of a new *service-oriented network architecture*. Towards this end, the paper also briefly describes PlanetLab, an overlay network being designed with these four principles in mind.

[7] C. V. Wright, F. Monrose, and G. M. Masson. On inferring application protocol behaviors in encrypted network traffic. *The Journal of Machine Learning Research*, 7:2745-2769, December 2006. [bib]

Several fundamental security mechanisms for restricting access to network resources rely on the ability of a reference monitor to inspect the contents of traffic as it traverses the network. However, with the increasing popularity of cryptographic protocols, the traditional means of inspecting packet contents to enforce security policies is no longer a viable approach as message contents are concealed by encryption. In this paper, we investigate the extent to which common application protocols can be identified using only the features that remain intact after encryption-namely packet size, timing, and direction. We first present what we believe to be the first exploratory look at protocol identification in encrypted tunnels which carry traffic from many TCP connections simultaneously, using only post-

encryption observable features. We then explore the problem of protocol identification in individual encrypted TCP connections, using much less data than in other recent approaches. The results of our evaluation show that our classifiers achieve accuracy greater than 90% for several protocols in aggregate traffic, and, for most protocols, greater than 80% when making fine-grained classifications on single connections. Moreover, perhaps most surprisingly, we show that one can even estimate the number of live connections in certain classes of encrypted tunnels to within, on average, better than 20%.

[8] E. Wustrow, S. Wolchok, I. Goldberg, and J. A. Halderman. Telex: Anticensorship in the network infrastructure. In *Proceedings of the 20th USENIX Security Symposium*, Aug. 2011. [bib]

In this paper, we present Telex, a new approach to resisting state-level Internet censorship. Rather than attempting to win the cat-and-mouse game of finding open proxies, we leverage censors' unwillingness to completely block day-to-day Internet access. In effect, Telex converts innocuous, unblocked websites into proxies, without their explicit collaboration. We envision that friendly ISPs would deploy Telex stations on paths between censors' networks and popular, uncensored Internet destinations. Telex stations would monitor seemingly innocuous flows for a special “tag” and transparently divert them to a forbidden website or service instead. We propose a new cryptographic scheme based on elliptic curves for tagging TLS handshakes such that the tag is visible to a Telex station but not to a censor. In addition, we use our tagging scheme to build a protocol that allows clients to connect to Telex stations while resisting both passive and active attacks. We also present a proof-of-concept implementation that demonstrates the feasibility of our system.
