



ELSEVIER

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Graption: A graph-based P2P traffic classification framework for the internet backbone

Marios Iliofotou^{a,*}, Hyun-chul Kim^b, Michalis Faloutsos^a, Michael Mitzenmacher^c, Prashanth Pappu^d, George Varghese^e

^a University of California, Riverside, CA, USA

^b Seoul National University, Gwanak-gu, Seoul, Korea

^c Harvard University, Cambridge, MA, USA

^d Conviva, Inc., San Mateo, CA, USA

^e University of California, San Diego, San Diego, CA, USA

ARTICLE INFO

Article history:

Received 6 February 2010

Received in revised form 18 November 2010

Accepted 29 January 2011

Available online 13 February 2011

Keywords:

Traffic classification

Behavioral-approach

Peer-to-peer

Graph mining

ABSTRACT

Monitoring network traffic and classifying applications are essential functions for network administrators. Current traffic classification methods can be grouped in three categories: (a) flow-based (e.g., packet sizing/timing features), (b) payload-based, and (c) host-based. Methods from all three categories have limitations, especially when it comes to detecting new applications, and classifying traffic at the backbone. In this paper, we propose the use of Traffic Dispersion Graphs (TDGs) to remedy these limitations. Given a set of flows, a TDG is a graph with an edge between any two IP addresses that communicate; thus TDGs capture network-wide interactions. Using TDGs, we develop an application classification framework dubbed Graption (*Graph-based classification*). Our framework provides a systematic way to classify traffic by using information from the network-wide behavior and flow-level characteristics of Internet applications. As a proof of concept, we instantiate our framework to detect P2P traffic, and show that it can identify 90% of P2P flows with 95% accuracy in backbone traces, which are particularly challenging for other methods.

© 2011 Elsevier B.V. All rights reserved.

1. Introduction

An important task when monitoring and managing large networks is classifying flows according to the application that generates them. Such information can be utilized for network planning and design, QoS and traffic shaping, and security. In particular, detecting P2P traffic is a potentially important problem for ISPs that want to manage such traffic, and for specific groups such as the entertainment industry in legal and copyright disputes. Detecting P2P traffic also has particular interest since it represents a large portion of the Internet traffic. In fact, a recent work by Labovitz et al. [27] shows that at least

20% of all Inter-domain traffic on the Internet is P2P. Even though percentage wise P2P traffic is not the dominant Internet traffic contributor today, its overall traffic volume is constantly increasing. In addition, other studies using advanced deep-packet inspection show P2P traffic to be up to 70% of the overall volume in some networks [20].

Most current application classification methods can be naturally categorized according to their level of observation: payload-based signature-matching methods [30,28], flow-level statistical approaches [10,32], or host-level methods [25,40]. Each existing approach has its own pros and cons, and no single method clearly emerges as a winner. Relevant problems that need to be considered include identifying applications that are new, and thus without a known profile; operating at backbone links [26,25]; and detecting applications that intentionally alter their

* Corresponding author. Tel.: +1 951 880 3699.

E-mail address: marios@cs.ucr.edu (M. Iliofotou).

behavior. Flow-level and payload-based classification methods require per application training and will thus not detect P2P traffic from emerging protocols. Behavioral-host-based approaches such as BLINC [25] can detect traffic from new protocols [25], but have weak performance when applied at the backbone [26]. In addition, most tools including BLINC [25] require fine-tuning and careful selection of parameters [26]. We discuss the limitations of previous methods in more detail in Section 4.

In this paper, we use the network-wide behavior of an application to assist in classifying its traffic. To model this behavior, we use graphs where each node is an IP address, and each edge represents a type of interaction between two nodes. We use the term *Traffic Dispersion Graph* or *TDG* to refer to such a graph [19]. Intuitively, with TDGs we enable the detection of network-wide behavior (e.g., highly connected graphs) that is common among P2P applications and different from other traffic (e.g., Web). While we recognize that some previous efforts [6,9] have used graphs to detect worm activity, they have not explored the full capabilities of TDGs for application classification. This paper is an extension of a workshop paper [18] and the differences will be clarified in the related work section (Section 4).

We propose a classification framework, dubbed **Graption** (*Graph-based classification*), as a systematic way to combine network-wide behavior and flow-level characteristics of network applications. Graption first *groups* flows using flow-level features, in an unsupervised and agnostic way, i.e., without using application-specific knowledge. It then uses TDGs to *classify* each group of flows. As a proof of concept, we instantiate our framework and develop a P2P detection method, which we call **Graption-P2P**. Compared to other methods (e.g., BLINC [25]), Graption-P2P is easier to configure and requires fewer parameters.

The highlights of our work can be summarized in the following points:

- *Distinguishing between P2P and client-server TDGs.* We use real-world backbone traces and derive graph theoretic metrics that can distinguish between the TDGs formed by client-server (e.g., Web) and P2P (e.g., eDonkey) applications (Section 2.2).
- *Practical considerations for TDGs.* We show that even a single backbone link contains enough information to generate TDGs that can be used to classify traffic. In addition, TDGs of the same application seem fairly consistent across time (Section 2.3).
- *High P2P classification accuracy.* Our framework instantiation (Graption-P2P) classifies 90% of P2P traffic with 95% accuracy when applied at the backbone. Such traces are particularly challenging for other methods (Section 3.2.2).
- *Comparison with a behavioral-host-based method.* Graption-P2P performs better than BLINC [25] in P2P identification at the backbone. For example, Graption-P2P identifies 95% of BitTorrent traffic while BLINC identifies only 25% (Section 3.3).
- *Identifying the unknown.* Using Graption, we identified a P2P overlay of the Slapper worm. The TDG of Slapper was never used to train our classifier. This is a promis-

ing result showing that our approach can be used to detect both known and unknown P2P applications (Section 3.4).

The rest of the paper is organized as follows. In Section 2 we define TDGs, and identify TDG-based metrics that differentiate between applications. In Section 3 we present the Graption framework and our instantiation, Graption-P2P. In Section 5 we discuss various practical issues. In Section 4 we discuss related work. Finally, in Section 6 we conclude the paper.

2. Studying the TDGs of P2P applications

2.1. Traffic dispersion graphs (TDGs)

Definition. Throughout this paper, we assume that packets can be grouped into flows using the standard 5-tuple $\{\text{srcIP}, \text{srcPort}, \text{dstIP}, \text{dstPort}, \text{protocol}\}$. Given a group of flows S , collected over a fixed-length time interval, we define the corresponding TDG to be a directed graph $G(V, E)$, where the set of nodes V corresponds to the set of IP addresses in S , and there is a link $(u, v) \in E$ from u to v if there is a flow $f \in S$ between them.


In this paper, we consider bidirectional flows. We define a TCP flow to start on the first packet with the SYN-flag set and the ACK-flag not set, so that the initiator and the recipient of the flow are defined for the purposes of direction. For UDP flows, direction is decided upon the first packet of the flow.

Visualization examples. In Fig. 1, we show TDG examples from two different applications. In order to motivate the discussion in the rest of the paper, we show the contrast between a P2P and a client-server TDG. From the figure we see that P2P traffic forms more connected and more dense graphs compared to client-server TDGs. In Section 2.2, we show how we can translate the visual intuition of Fig. 1 into quantitative measures that can be used to classify TDGs that correspond to different applications.

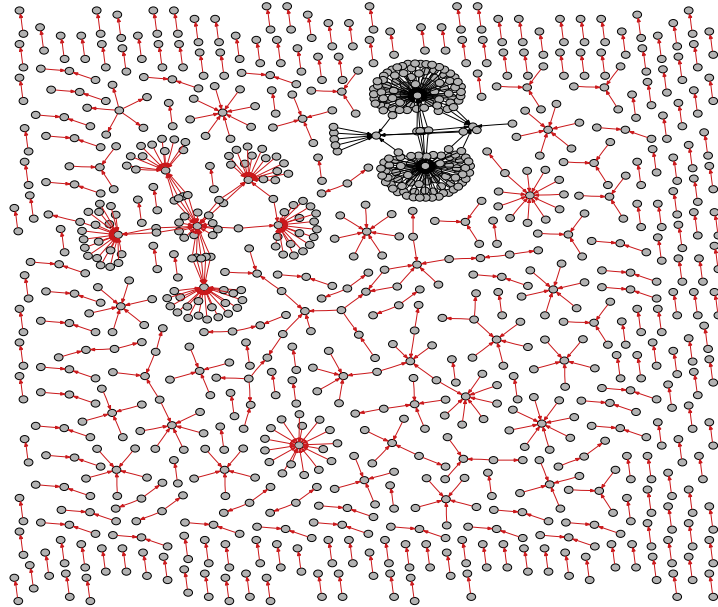
Data set. To study TDGs, we use three backbone traces from a Tier-1 ISP and the Abilene (Internet2) network. These traces are summarized in Table 1. All data are IP anonymized and contain traffic from both directions of the link. The TR-PAY1 and TR-PAY2¹ traces were collected from an OC48 link of a commercial US Tier-1 ISP at the Palo Alto Internet eXchange (PAIX). To the best of our knowledge, these are the most recent backbone traces with payload that are available to researchers by CAIDA [5]. The TR-ABIL trace is a publicly available data set collected from the Abilene (Internet2) academic network connecting Indianapolis with Kansas City. The Abilene trace consists of five randomly selected five-minute samples taken every day for one month, and covers both day and night hours as well as weekdays and weekends.

Extracting ground truth. We used a Payload-based Classifier (PC) to establish the ground truth of flows for the

¹ The authors thank CAIDA for providing this set of traffic traces. Additional information for these traces can be found in the DatCat, Internet Measurement Data Catalog [8], indexed under the label "PAIX".



(a) The Gnutella TDG (P2P application).



(b) The HTTPS TDG (client-server application).

Fig. 1. Two TDG visualization contrasting a P2P (top) with a client-server application (bottom). Largest component is with bold (black) edges – hence the graphs are best viewed on a computer screen or a colored print-out. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 1

Set of backbone traces from the Cooperative Association for Internet Data Analysis (CAIDA). Statistics for the TR-ABIL trace, are reported only for the first five-minute interval since IPs were anonymized differently at each five-minute sample.

Name	Date/time	Duration	Flows
TR-PAY1	2004-04-21/17:59	1 h	38,808,604
TR-PAY2	2004-04-21/19:00	1 h	37,612,752
TR-ABIL	2002-09/(N/A)	1 month	2,057,729

TR-PAY1 and TR-PAY2 traces. Both traces contain up to 16 bytes of payload in each packet, thereby allowing the labeling of flows using the signature matching techniques described in [26,25]. Running the PC over the TR-PAY1 and TR-PAY2 traces we find 14% of the traffic to be P2P, 27% Web, 7% DNS, and the rest to belong to other applications. A detailed application breakdown is summarized in Table 2. In the table we further report the traffic information for the top six P2P applications with the highest

Table 2

Application breakdown for TR-PAY1 and TR-PAY2 traces. Values in parenthesis show the percentage of each P2P application over the entire identified P2P traffic (All-P2P).

Name	% in flows	% in bytes	% in packets
Gnutella	0.95(6.78)	0.17(1.59)	0.81(6.32)
eDonkey	2.96(21.16)	2.22(21.17)	2.84(22.21)
FastTrack	0.55(3.92)	0.74(7.10)	0.97(7.61)
Soribada	7.76(55.44)	0.07(0.63)	0.97(7.63)
MP2P	0.41(2.93)	0.01(0.14)	0.07(0.53)
BitTorrent	0.60(4.26)	4.59(43.81)	4.37(34.24)
All-P2P	13.85	9.19	12.10
Web (80/443)	27.45	41.99	35.51
SMTP	1.65	1.10	2.431
DNS	6.65	0.32	1.586
Games	0.71	0.54	2.84
Unknown	0.97	1.98	3.22
No-Payload	28.06	0.46	6.04
Rest	20.66	44.42	36.27

number of flows in our data set. These six applications contribute ~95% of the flows and ~75% of the overall P2P bytes. In our traces, we also have flows by P2P applications that are not longer in use, such as WinMX, Direct Connect, Goboogy, Ares, SoulSeek, and EarthStation. We observed that some of these protocols were used for transferring very large files and contributed many bytes, but their overall number of flows was significantly lower. In our evaluation we include all P2P flows.

For our study, we remove the 1% of traffic that remained unknown and the 28% that contained no payload. The small number of unknown flows is strongly correlated with the small number of encrypted flows in our traces.

2.2. Identifying P2P TDGs

Identifying the right metrics to compare graph structures is a challenging question that arises in many disciplines [31]. Our approach is to consider several graph metrics, each capturing a potentially useful characteristic, until a set of metrics is found that distinguishes the target graphs.

To select an appropriate set of metrics, we generate a large number of TDGs using all our traces (Table 1), thus observing TDGs over two different locations at the backbone. For the TR-PAY1 and TR-PAY2 traces, we use the payload-based classifier (PC) in order to select which flows belong to each TDG. Since the TR-ABIL trace does not have any payload information, we use port numbers [26] to assign flows to applications. We can use port numbers for the TR-ABIL trace since it was collected in 2002 where most P2P applications used their default port numbers [13,23]. We only use the TR-ABIL trace to verify our TDG observations over a second location in the backbone and we do not use it in the final evaluation of our classifier. By using the month-long TR-ABIL trace, we can study the consistency of TDGs over different times of the day and over weekdays and weekends.

We observe TDGs over five-minute intervals. This interval length gives good classification results and stability of TDG metrics over time as we show later in this section. For each TDG we generate a diverse set of metrics. Our

metrics capture various aspects of TDGs including the degree distribution, degree correlations, connected components, and distance distribution. For additional details about these metrics we refer the reader to [19,31].

To select the right set of metrics we use various graph visualizations and trial and error. Finding a less ad hoc approach is beyond the scope of this work. Two TDG visualization examples are shown in Fig. 1. We see that FastTrack (P2P) has a denser graph than HTTPS, or a higher **average degree**, where the average node degree \bar{k} is given by $\bar{k} = 2|E|/|V|$.

We utilize two other metrics that capture the directionality of the edges in the graph and the distances between nodes. The directionality is useful since we know that pure clients only initiate traffic, pure servers should never initiate traffic, and that some P2P nodes play both roles. To capture this quantitatively, we define **InO** to be the percentage of nodes in the graph that have both incoming and outgoing edges.

The distance between two nodes is defined as the length of their shortest path in the graph. The diameter of a graph is defined as the maximum distance between all pairs of nodes, which is sensitive as a metric [31]. For a more robust metric, we use the *effective diameter* (EDiam), which we define as the 90-th percentile of all pairwise distances in the graph.

We show the high graph diameter of the BitTorrent P2P TDGs in Fig. 2. The graph represents the BitTorrent interactions between 3000 nodes from the TR-PAY1 trace. The visualization is suggestive of a spider's web with large paths between nodes at opposite edges. This intuition is precisely captured by the large value of the effective graph diameter. In general TDG visualizations are suggestive, but the metrics make the intuitions precise and allow classification.

From our measurements, we empirically derive the following two rules for detecting P2P activity. **Rule 1:** $\bar{k} > 2.8$ and $InO > 1\%$; **Rule 2:** $InO > 1\%$ and $EDiam > 11$. With these simple rules, we can correctly identify all P2P TDGs from both backbone locations (Abilene backbone and Tier-1 ISP). Intuitively, P2P hosts need to be connected with a large set of peers in order to perform tasks such as answering content queries and sharing files, which can explain the higher average degree compared to client-server applications. An additional characteristic of P2P applications is the duality of roles, with many hosts acting both as client and server. The duality of roles is in turn captured by the high InO value. We further speculate that the decentralized architecture of some P2P applications (such as BitTorrent), can explain the high diameters in some P2P TDGs.

We do not claim that our thresholds are universal, but our measurements suggest that small adjustments to these simple parameters allow our methodology to work on different backbone links.

Distinguishing collaborative applications from P2P: Some well-known applications other than P2P exhibit collaborative behavior, such as DNS and SMTP. This is not surprising since in these applications servers communicate with each other and with other clients (high k), and servers act both as clients and servers (high InO). This is exactly what our metrics are set out to detect. Next, we outline two




Fig. 2. Graphical representation of the BitTorrent TDG with 3000 IPs, showing the formation of long paths connecting P2P hosts. The largest connected component is highlighted with darker color edges.

approaches that can help to separate between collaborative applications (e.g., DNS) and P2P: (1) One solution is to rely on port numbers or payload information for this task. In fact, it has been reported recently [26] that port numbers are fairly accurate in identifying such legacy applications, although they fail to identify P2P and other applications with dynamic use of port numbers. Therefore, one could use legacy ports to pinpoint and isolate such collaborative applications and then use graph metrics on the remaining traffic. In addition to port inspection, we can also examine the payload of a flow in order to verify that it follows the expected application-layer interactions. (2) In a recent work on dynamic graphs [16], we show that measuring how much TDGs change over time gives additional information about the nature of the applications. In particular, we observed that P2P TDGs are more dynamic (i.e., how frequently nodes and edges change over time) than legacy applications, such as DNS. If the classification problem at hand requires this further separation, we can incorporate such dynamic techniques in Graption.

2.3. Practical considerations for TDGs

2.3.1. Stability of metrics over time

If thresholds derived from TDGs changed significantly with time, then the classifier must be trained constantly and that would detract from its value. Fortunately this does not appear to be the case. We show the stability of TDGs in time and space using traces from different points in time at two different backbone locations for Abilene and a commercial Tier 1 ISP (Section 2.2).

To test stability over time, we measured P2P TDGs from our longest trace (TR-ABIL) that span over one month and include samples taken over day and night hours, and both weekdays and weekends. To the best of our knowledge such month-long traces with payload are not available. Stability was also observed in all the one-hour long traces (TR-PAY1, TR-PAY2).

Fig. 3 summarizes the stability study for TR-ABIL showing the average degree metric. As we see, average degree takes values in a small range over the entire month. Soribada tends to have very small graphs ranging from 30 to 1000 nodes and hence the higher variability. Stability was also observed for the diameter and InO metric with the figures not included for brevity. For this study, we used TDGs representing five-minutes of traffic. In Section 2.3.2




Fig. 3. The average degree for various P2P protocols over one month in trace TR-ABIL. Candle sticks show the maximum and minimum recorded values together with the average (horizontal line) and \pm standard deviation. For visualization purposes, for the MP2P application we show only the minimum value which is very high (23.73).

we show experiments with TDGs where we vary the observation interval from five seconds to 15 min. Note that BitTorrent is not shown in Fig. 3 since it was not well-known in 2002 and resulted only in a handful of flows in the TR-ABIL trace.

2.3.2. Selecting the interval of observation

The effect of the observation interval on the graph metrics. We observe that by increasing the observation interval, graph metrics are more stable across successive TDGs of an application. To show this, we measure the variation of all our graph metrics over time as a function of the interval of observation. We vary the interval from one second up to 15 min. For each interval length, we generate consecutive graph snapshots with non-overlapping time intervals. For each interval, we extract the graph and calculate the average degree, InO, and effective diameter metrics.

We examine the variability of each metric over all intervals of observation. We use the commonly-used *range* metric, which is the difference between the maximum and minimum observed values over all intervals for each graph metric. For each metric, we then normalize the range by the average value calculated over the time series. This normalization makes the range for different metrics somewhat comparable. For each application, we report the average range over all three graph metrics, and we plot this average range versus the duration of the observation interval in Fig. 4. Our experiments showed very high range values for intervals smaller than 60 s. The range shows significant decrease for intervals larger than a minute. Increasing the interval from five minutes to 10 and 15 does not significantly effect the range. This trend is shown in Fig. 4, where we report the average range over all metrics for a set of our applications. Even though we have variability in our results, there was a clear trend in all our measurements. Similar trends we observed in all our traces (Table 1). For the rest of our study, we use five minute intervals.

3. The Graption framework

The Graption traffic classification framework consists of the following three steps.

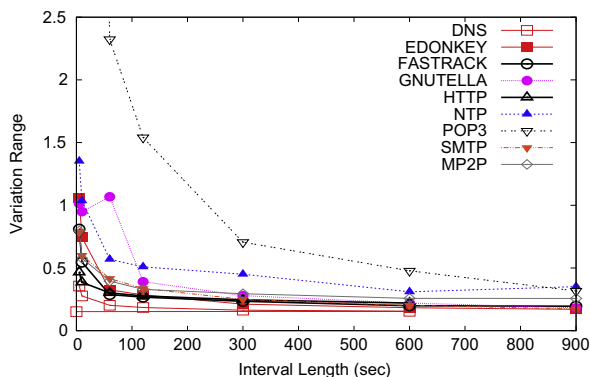


Fig. 4. The effect of changing the interval of observation for TDGs ranging from five seconds up to 15 min over a large set of protocols. To reduce variability we can choose to use longer intervals. After 300 s the reduction in variability is not significant.

Step 1: Flow isolation. The input is network traffic in the form of flows as defined in Section 2. The goal of this first optional step is to utilize external information to isolate any flows that can already be classified. This knowledge could be based on payload signatures, port numbers, or IP address (e.g., exclude flows from a particular domain such as `google.com`).

Step 2: Flow grouping. We use similarity at the flow and packet level to group flows. The definition of similarity is flexible in our proposed methodology. We can use flow statistics (duration, packet sizes, etc.) or payload if this is available. Eventually, the output of this step is a set of groups with each group ideally containing flows from a single application (e.g., Gnutella, NTP, etc.). However, at this step, the exact application of each group is not known.

Step 3: Group classifier. For each group of flows, we construct a TDG. Next, we quantify each TDG using various metrics. The classifier uses these metrics to identify the application for each group of flows. For the classification decision, we use a set of rules which in general depend on the focus of the study.

Although this paper focuses on P2P detection, we believe Graption can be used for general application classification by choosing metrics and parameters appropriately. We next describe how we specialized Graption to detect P2P traffic (Graption-P2P).

3.1. Implementation details of Graption-P2P

Step 1: This is an optional step in our methodology. Experiments without this step are discussed later in the section. Recent work [26] suggests that port-based classification works very well for legacy applications, as legacy applications use their default ports and tunneling of P2P at such ports is not very common. Thus, in this study, we isolated flows with port 80 for Web, port 53 for DNS, and port 25 for SMTP. These applications turn out to be about 65% of the total number of flows. In our traces, the proportion of P2P actually using one of these ports is as low as 0.1%.

Step 2: To implement flow grouping we use the fact that application-level headers are likely to recur across flows from the same application. Therefore, payload similarity can be used to group flows. In Graption-P2P, we only use the first sixteen bytes from each flow. As we show, sixteen bytes are sufficient to give very good classification results. This observation agrees with findings in [30,28]. Even though we use the payload bytes, our grouping is agnostic to application semantics, as each byte is considered as a single independent categorical feature. We consider each byte as a single categorical feature in the range $\{0, 1, \dots, 255\}$.

The flow grouping step comprises two sub-steps: cluster formation and cluster merging.

(a) *Forming clusters.* Given the set of discriminating features, the next step is to cluster “similar” flows together. We use the term **cluster** to describe the outcome of an initial grouping using the selected features. Clusters may be merged in the next function of this step to form **groups**, which produces the final output.

Feature-based clustering is a well-defined statistical data mining problem. For this task we used the popular **K-means** algorithm [38]. This algorithm has been commonly used for unsupervised clustering of network flows [10,30], with very good results and low computational cost. K-means operates with a single parameter that selects the number of final clusters (k). As we show later in our evaluation, our classifier gives very good results over a large range of k .

The similarity between two flows is measured by the *Hamming* [38] distance calculated over the 16 categorical features (i.e., the payload bytes). Even though more involved similarity measures such as edit-distance (also known as Levenshtein distance) exist, Hamming distance has been used successfully before [15] and performs very well in our application.

(b) *Cluster merging.* During clustering, it is likely that the same application generates multiple clusters. For example, many P2P protocols exhibit a variety of interaction patterns, such as queries (UDP flows) and file transfers (TCP flows), each with significantly different flow and packet characteristics [23]. This motivates mergings clusters that we expect to belong to the same application into groups. This grouping provides a more complete view of the application and aids in understanding the structure of the P2P protocol, as we show in Section 3.2.

Cluster merging cannot be based on the chosen set of flow-level features that were already used to create the clusters originally. Instead, in the case of a P2P protocol, it is natural to assume that the TDGs corresponding to each cluster of the same protocol would share a large number of common nodes (IP addresses).

Based on these observations, we use an Agglomerative (Hierarchical) Clustering Algorithm that recursively merges clusters with significant similarity in IP addresses. We used the following metric to calculate similarity between clusters: $Sim(C_1, C_2) = (\text{Number of flows having their source or destination IPs present in both clusters}) / (\text{The number of flows of the smaller cluster})$. The cluster merging process starts by hierarchically merging clusters with high similarity and stops when the similarity between all new cluster pairs is below a **similarity threshold (ST)**. As we show later in our evaluation, our classifier gives very good results over a large range of similarity thresholds.

Step 3: The outcome of the previous step is a set of groups of flows, with each group consisting of flows that we hope stem from a single application. In order to classify each group, we generate a TDG on the group in the same way as described in Section 2. Each group yields a TDG that can be summarized using graph metrics. To identify P2P TDGs, we used the rules extracted from Section 2.2. When a group is labeled as P2P then all the flows of that group are classified as P2P flows.

3.2. Evaluating Graption-P2P

To evaluate Graption-P2P, we use traces TR-PAY1 and TR-PAY2, where we have the ground truth using the payload classifier (Section 2). We compute the True Positives, False Positives, and False Negatives. The True Positives (TP) measures how many instances of a given class are correctly classified; the False Positives (FP) measures how many instances of other classes are confused with a given class; and the False Negatives (FN) measures the number of misclassified instances of a class. In our comparisons, we used the following standard metrics: **Precision** (P), defined as $P = TP / (TP + FP)$; **Recall** (R), defined as $R = TP / (TP + FN)$; and the **F-Measure** [38], defined as $F = 2P \cdot R / (P + R)$, combining P and R. In all our results, the metrics above report the classification performance in terms of flows.

3.2.1. Forming clusters

The K-means algorithm generates k clusters and assigns each flow to a different cluster. We first test the effectiveness of K-means to form **pure** clusters. A pure cluster will ideally contain flows from a single application. To label the clusters, we use the *dominant heuristic* [3]. In this heuristic, using the ground truth of flows we label each cluster as belonging to the application with the majority of flows in the cluster. All the flows of a cluster are then classified to belong to this dominant application.

Using the dominant heuristic we assign clusters to applications and calculate the precision (P) and recall (R) for different values of k . The clustering results as we increase k for both traces are shown in Fig. 5. We observe that with sufficiently large k (>120) we achieve very good results with P and R above 90%. By increasing k , on one hand we have more “pure” clusters but on the other we make the cluster merging step harder. In the extreme case, each cluster will contain a single flow giving 100% P and R, but making cluster merging challenging. We will return to this topic later in Section 3.5.

The Graption-P2P classifier is not sensitive to the exact value for k , since we do not require that each application maps to exactly one cluster. Instead, we only need enough clusters so that flows from different applications do not share a cluster. Hence, it is reasonable to slightly overesti-

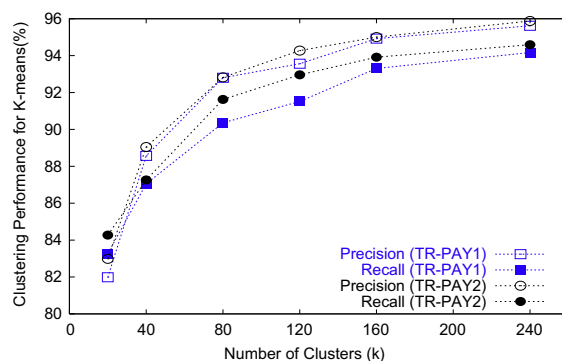


Fig. 5. Evaluating K-means: With sufficiently large k (>120) K-means can efficiently separate the flows of different application. All classification metrics are in terms of flows.

mate the k . As we show next, with k in the range 120–240 Graption-P2P has high classification performance.

3.2.2. Identifying P2P traffic

Using Graption-P2P, we achieve high flow P2P F-Measure over a range of values of k (K-means) and similarity thresholds (ST). We show this in Fig. 6, where we vary the ST from 0.01 to 1 and use a sufficiently large k (see Fig. 5). All experiments are averaged over each disjoint five minute interval of both traces. Intuitively, by using a very large ST, the clusters of an application are not grouped together, which results in TDGs that are harder to classify as P2P. On the other hand, with a very small ST, clusters belonging to different applications are merged together leading to poorer classification performance. The results in Fig. 6 show that we achieve good classification performance (>90% F-Measure), over a large range of similarity thresholds and number of clusters (k).

In Fig. 7, we compare our approach with labeling each cluster using the ground truth (i.e. without merging any clusters and labeling each cluster using the dominant heuristic). Intuitively, for a given clustering of flows, the ground truth shows the best that our cluster labeling mechanism can achieve. For merging, we use a ST of 0.5. From Fig. 7, we see that Graption-P2P deviates only slightly from labeling clusters using the ground truth. In the same plot, we also compare Graption-P2P without the cluster




Fig. 6. Graption-P2P achieves >90% flow F-Measure over a large range of similarity thresholds and number of clusters (k).

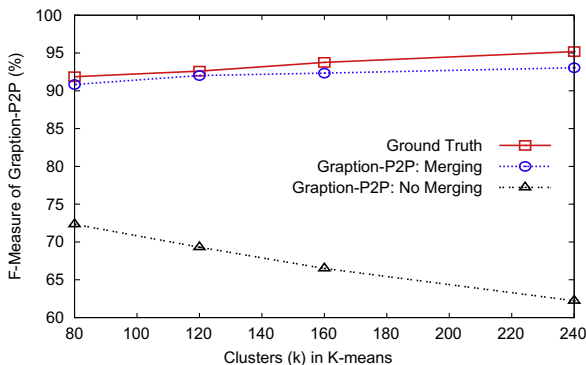


Fig. 7. Graption-P2P compared to cluster labeling based on ground truth. Results are also compared with and without cluster merging.

merging step, highlighting the benefit of merging clusters of the same application together.

Using a ST of 0.5 and $k = 160$, Graption-P2P achieves above 90% Recall and above 95% Precision over all disjoint five minute intervals for both traces. To apply Graption-P2P to other backbone links, the same selection process can be repeated to adjust the values of ST and k . Our experiments show that the classification performance can degrade with a bad choice of parameters. However, as shown in Figs. 5–7, for reasonable choices for k and ST, our method provides very good results. The Graption framework is flexible enough to incorporate state-of-the-art clustering methods.

3.3. Comparison with BLINC [25]

We used BLINC to classify traffic for both TR-PAY1 and TR-PAY2 traces. BLINC was optimized after several trial and error efforts to achieve its best accuracy for classifying P2P traffic over these traces. A detailed description of the selection process for the parameters for BLINC is provided in [26]. As reported by Kim et al. [26], tuning the 28 parameters of BLINC is a time consuming process. This is because small changes to the thresholds have a significant effect on the classification results.

The Recall and Precision for BLINC are 84% and 89% respectively compared to 90% and 95% of our method. The classification performance per P2P protocol for both BLINC and Graption-P2P are shown in Fig. 8. Given that both methods do not distinguish between particular P2P protocols (e.g., Gnutella vs BitTorrent), it only make sense to report the portion of identified traffic (recall) for each protocol. Graption-P2P can identify more traffic for all P2P applications tested with the exception of MP2P. We can see from Fig. 3 that MP2P has a much higher average node degree than the other P2P applications which can explain the high performance of BLINC over this protocol. Going back to Fig. 8, we see that BLINC has significantly lower performance for some P2P applications. For example, Graption-P2P detects 95% of BitTorrent traffic, while BLINC detects only 25%! In addition, Graption-P2P detects 91% of the flows from the three popular P2P applications {BitTorrent, Gnutella, and eDonkey} while BLINC detects only 73%.

Our experiments suggest that BLINC and possible other behavioral-host-based approaches work well when applied

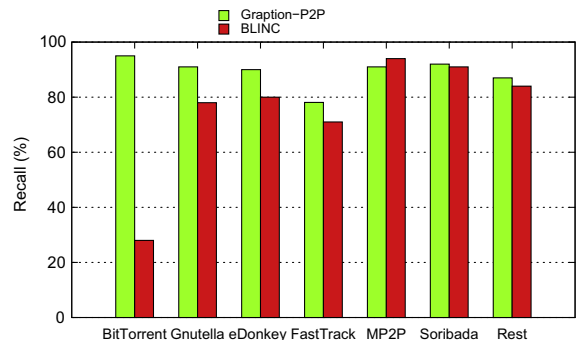


Fig. 8. The percentage of P2P flows detected by Graption-P2P and BLINC. Flow precision for Graption-P2P is 95% and for BLINC is 89%.

at the edge, where a large fraction of host flows are observed and hence enough evidence is collected to profile each node. However, this is not always true for backbone monitoring points which can explain BLINC's lower performance. These observations are also supported by findings in [26].

We acknowledge that BLINC provides a methodology for detecting traffic from multiple classes of applications (e.g., HTTP, FTP, Chat, etc.) and not only for P2P. Our results only show that Graption-P2P does better in P2P detection.

3.4. Identifying the unknown

Our goal here is to answer the following question: *Can Graption identify ports used by unknown P2P applications in our traces?* To achieve this, we configure Graption as follows. For step 1, we isolate legacy flows using the same approach as in Graption-P2P. In addition, we isolate flows from ports that belong to known P2P applications (e.g., port 1214 for FastTrack). At step 2, we cluster flows using their port number. This produces one cluster for each distinct port in the trace. For merging clusters, we use the same similarity threshold as in the Graption-P2P configuration. Finally, for classifying each group (step 3) we use the same rules as in Graption-P2P. We will refer to this configuration as Graption-P2P-Ports.

By manually inspecting the TDGs that Graption-P2P-Ports reports as P2P, we discovered interesting behaviors in the Abilene Trace (TR-ABIL). Our simple method was able to detect three ports (2002, 4156, 1978) that we initially thought were false positives. On further research, we found that these three ports were used by the Slapper² worm. Slapper has been reported to use these ports in order to form a P2P botnet overlay for launching DDoS attacks. Even though Graption-P2P-Ports was never explicitly trained to detect the Slapper overlay, it used the TDG profile of known P2P applications (e.g., eDonkey, Gnutella, etc.) to identify it. This is an advantage of using network-wide behavior to classify traffic.

Graption-P2P-Ports also captured an unusual behavior of IPs executing `traceroute` to each other forming a near clique. Such highly connected group of hosts was not observed in any of the P2P applications present in our traces. We speculate these interactions are by host performing active network measurements inside the Internet2 (Abilene) backbone. Given this unique connection pattern, we can always isolate such unusual behavior and distinguish it from a potentially new P2P application. Besides the active measurement traffic and the flows from the Slapper worm, Graption-P2P-Ports reported only 0.3% other flows as P2P. With no payload information, we were not able to manually match these flows to a known exploit or P2P application.

3.5. Other Graption configurations

For completeness, we show here preliminary results with other system configurations assuming that access to payload is limited or isolation cannot be used. We proceed by describing these configurations.

Graption-P2P-NI: without isolation. This is the same configuration as the Graption-P2P, but without using isolation. By increasing k in K-means above 300 we achieve good results; >92% precision and >85% recall. A larger k was to be expected, since by including all the Web, SMTP, and DNS flows the total number of signatures increased as well. Our preliminary experiments showed that by increasing the number of clusters, it makes the cluster merging step more difficult. This observation highlights the advantage of using isolation.

Graption-P2P-NP: using only flow-level features. This configuration is the same as the Graption-P2P, but assuming that payload and port numbers cannot be used. We achieved >90% precision and >88% recall over a range of k and similarity thresholds. For each flow, we extracted more than 40 flow features ranging from packet size information (size of first 10 packets, max/min packet size, etc.), timing information (flow duration, min and max inter-packet gap, etc.), TCP flags, total volumes in bytes, number of packets, etc. In order to test the relevance of each feature we applied the Information Gain Ranking Filter [38] over various time intervals. For this configuration, we used the most prominent features: packet size information (i.e. min, max, and the size of the first five packets) and protocol (UDP or TCP). A recent work [11] also supports our observation that packet sizes are good features to cluster/classify network flows.

4. Related work

This journal paper extends our six-page workshop version [18]. In addition to the work presented in the workshop version, we here include: (a) A study on practical considerations for using TDGs (see Section 2.3, Figs. 3 and 4). (b) We provide additional details regarding our data sets and our classification results (see Sections 2.2 and 3.5, Fig. 2, and Table 2). (c) We provide a more detailed comparison with BLINC (see Section 3.3 and Fig. 8). (d) We include additional experiments showing how Graption can be used to detect previously unknown traffic (see Section 3.4).

4.1. Traffic classification

As an alternative to port-based methods, some work use payload [30,28]. Other approaches use machine learning (ML) algorithms to classify traffic using flow features (e.g., packet sizes). For an exhaustive list and comparison of ML algorithms we refer the reader to [32,26]. The problem of training a ML algorithm on one trace and applying it on another was recently addressed in [37]. Even though supervised solutions to traffic classification are continuously improving, all supervised methods require per application training and will thus not detect traffic from new applications. Our work has more in common with unsupervised methods which group similar flows together. All previous methods [3,29,10] require manual labeling of clusters. Our work bridges this gap by providing a method to automatically label clusters of flows based on their network-wide behavior.

In BLINC [25], the authors characterize the connection patterns (e.g., if it behaves like using P2P) of a single host at the transport layer and use these patterns to label the

² <http://www.cert.org/advisories/CA-2002-27.html>

flows of each host. BLINC uses graph models called graphlets to model a host's connection patterns using port and IP cardinalities. Unlike TDGs, graphlets do not represent *network-wide* host interaction. In some sense, TDGs represent a further level of aggregation, by aggregating across hosts as well. Thus it is perhaps fair to say that while BLINC hints at the benefit of analyzing the node's interaction at the "social" level, it ultimately follows a different path that focuses on the behavior of individual nodes. As we show, our approach performs better than BLINC in our backbone traces (see Section 3.3). Similar to BLINC, other host-based methods [2,22] target the identification of P2P users inside a university campus (i.e., network edge). The connection patterns of neighboring host (e.g., their degree distribution) were also used as features to profile network hosts [1]. Unlike Graption-P2P, in [2,22,25,1] they do not use *network-wide* host interaction. In [7], the authors use a port-based method to identify P2P users, using their temporal appearance and connection patterns in a trace.

A different approach to host profiling was introduced by Trestian et al. [36]. They used readily available information from the Web to classify traffic using the Google search engine. They show very good results for classifying flows for legacy application, but their results are not promising for P2P detection because of the dynamic nature of P2P IP hosts. Our method can thus be used to complement the work in [36].

The problem of classifying traffic from applications that use sophisticated obfuscation techniques is the topic in [17]. The solution proposed uses advanced community detection algorithms to group related hosts together. It then requires external seed information about the application used by few hosts in each group in order to classify the entire group. The solution in [17] is very different from this work since it does not classify applications using network-wide behavioral information. We believe that Graption-P2P can provide a way to identifying P2P communities without relying on external seeding information.

4.2. Worm detection

Graphs have been used for detecting worm activities within enterprise networks [9]. Their main goal was to detect the tree-like communication structure of worm propagation. This characteristic of worms was also used for post-mortem trace analysis (for the identification of the source of a worm outbreak, the so-called patient zero) using backbone traces [39]. More recent studies use graph techniques to detect hit-list worms within an enterprise network, based on the observation that an attacker will alter the connected components in the network [6].

4.3. Measurements on network-wide interactions

Statistical methods are used in [40] for automating the profiling of network hosts and ports numbers. The connectivity behavior and habits of users within enterprise networks is the focus of many papers, including [35]. In [34], the authors study P2P overlays using passive measurements, but target mainly the profiling of P2P hosts. The most recent work on network-wide interactions is by Jin et al. [21]. In their work [21], they use graph-partitioning methods to ex-

tract and study the evolution of smaller communities within a TDG. None of the above papers targets P2P detection.

5. Discussion

Enhancing the isolation step. To improve isolation we can enforce payload inspection in addition to port-based filtering. For example, we can test all DNS flows at port 53 to see if they also have a DNS payload signature or if another protocol is tunneling its traffic under the DNS port. If payload is encrypted, then we can choose to use flow-level feature such as packet sizes [26] or white-listed IP addresses [36].

Applying our framework to other backbone locations. In our study we show that is feasible to correctly identify P2P TDGs from all our backbone traces using a single set of thresholds. In other words, we do not use different classification threshold for each individual trace. Even though TDGs from different locations may have differences between them, in our study, we always find TDGs from the same family (e.g., P2P) to have similar characteristics (e.g., high average degree). In Section 2.2, we explicitly state that our thresholds are not universal. Therefore, it might not always be possible to use exactly the same thresholds to identify particular TDGs in all backbone locations. Nevertheless, our measurements suggest that small adjustments to our simple parameters (see Section 2.2) allow our methodology to work on different backbone links.

TDG formation and packet sampling. Packet sampling is used in order to reduce the overhead of processing every single packet by a monitoring system. Sampling effectively reduces the overall number of observed flows by the monitor. Less overall flows translates to fewer observed edges over time when forming TDGs. In our experiments in Section 2.3, we observe that the longer we collect edges for a TDG (i.e., applying a longer observation interval) the number of edges and the stability of metrics increases. With sampling, we believe the curves in Fig. 4 may change. In particular, in order to reduce TDG variability in the presence of sampling, the monitoring interval will need to be increase.

Encryption in P2P applications. If an application encrypts all traffic classification methods using payload [33,28,15,30] will be out-dated. Note, however, that it has been shown before [4] that even if an application encrypts its traffic there are ways to distinguish it from other flows. To show that even in the case where P2P encryption becomes prevalent our framework will still be useful, we also experimented with Graption-P2P-NP using only flow level features and observed good results with recall and precision of above 88% in both traces.

Obfuscation of P2P applications at multiple levels. In the future, obfuscation might become more common. In the case where IPsec (encryption up to Layer-4 headers) becomes prevalent, transport layer [24], and host-based [25] methodologies will most likely not work. If P2P protocols start using *polymorphic blending methods* [12,14] to masquerade their traffic at the packet (e.g., packet sizes) and payload level to resemble a legacy application (e.g., Web), then all current machine learning [32] methods will not work. In such a scenario, TDGs (IP-to-IP graphs) might still prove to be useful because the interaction patterns seem hard to hide.

6. Summary and conclusions

The underlying theme of our work is to go beyond just monitoring individual packets, flows, or hosts, to *also* monitoring the interactions of a group of hosts. Towards this end, we introduce TDGs and show their potential to generate novel tools for traffic classification. Based on TDGs, we developed Graption, a graph-based framework, which we then specialize (Graption-P2P) for the detection of P2P applications. We show that Graption-P2P classifies more than 90% of P2P traffic with above 95% precision when tested on real-world backbone traces. While our best results are obtained assuming access to some initial payload bytes, Graption still performs well (90% precision with 88% recall) even when payload is not used. We believe the network-wide characteristics of P2P traffic (e.g., InO, average degree, and diameter) are not as easy to hide without significantly altering P2P functionality.

Acknowledgments

This work was supported partly by NSF NECO 0832069, a CISCO URP grant, ARL CTA W911NF-09-2-0053, by NAP of Korea Research Council of Fundamental Science and Technology, and the ITRC support program [NIPA-2010-C1090-1011-0004] of MKE/NIPA. Michael Mitzenmacher was supported in part by NSF grant CNS-0721491 and research grant from CISCO. George Varghese was supported by grants from NSF and CISCO. The SDSC's TeraGrid and compute resources are supported by the NSF grant CONMI CRI-0551542. Support for CAIDA's Internet traces is provided by the National Science Foundation, the US Department of Homeland Security, and CAIDA Members.

References

- [1] E.G. Allan, W.H. Turkett, E.W. Fulp, Using network motifs to identify application protocols, in: IEEE GLOBECOM, 2009.
- [2] G. Barlett, J. Heidemann, C. Papadopoulos, Inherent behaviors of on-line detection of peer-to-peer file sharing, in: IEEE GI, 2007.
- [3] L. Bernaille, R. Teixeira, K. Salamatian, Early application identification, in: ACM CoNEXT, 2006.
- [4] D. Bonfiglio, M. Mellia, M. Meo, D. Rossi, P. Tofanelli, Revealing Skype Traffic: when randomness plays with you, in: ACM SIGCOMM, 2007.
- [5] CAIDA. <http://www.caida.org>.
- [6] M.P. Collins, M.K. Reiter, Hit-List worm detection and bot identification in large networks using protocol graphs, in: RAID, 2007.
- [7] F. Constantinou, P. Mavrommatis, Identifying known and unknown peer-to-peer traffic, in: IEEE NCA, 2006.
- [8] DatCat, Internet Measurement Data Catalog. <http://www.datcat.org>.
- [9] D. Ellis, J. Aiken, K. Atwood, S. Tenarglia, A behavioral approach to worm detection, in: ACM CCS WORM, 2004.
- [10] J. Erman, M. Arlitt, A. Mahanti, Traffic classification using clustering algorithms, in: ACM SIGCOMM MineNet, 2006.
- [11] A. Este, F. Gringoli, L. Slagerelli, On the stability of the information carried by traffic flow features at the packet level, ACM SIGCOMM CCR 39 (3) (2009) 13–18.
- [12] P. Fogla, M. Sharif, R. Perdisci, O. Kolesnikov, W. Lee, Polymorphic blending attacks, in: USENIX Security Symposium, 2006.
- [13] A. Gerber, J. Houle, H. Nguyen, M. Roughan, S. Sen, P2P, the Gorilla in the cable, in: National Cable and Telecommunications Association (NCTA), 2003.
- [14] G. Gu, R. Perdisci, J. Zhang, W. Lee, BotMiner: clustering analysis of network traffic from protocol- and structure-independent botnet detection, in: USENIX Security Symposium, 2008.
- [15] P. Haffner, S. Sen, O. Spatscheck, D. Wang, ACAS: automated construction of application signatures, in: ACM SIGCOMM MineNet, 2005.
- [16] M. Iliofotou, M. Faloutsos, M. Mitzenmacher, Exploiting dynamicity in graph-based traffic analysis: techniques and applications, in: ACM CoNEXT, 2009.
- [17] M. Iliofotou, B. Gallagher, G. Xie, T. Eliassi-Rad, M. Faloutsos, Profiling-by-association: a resilient traffic profiling solution for the internet backbone, in: ACM CoNEXT, 2010.
- [18] M. Iliofotou, H. Kim, P. Pappu, M. Faloutsos, M. Mitzenmacher, G. Varghese, Graph-based P2P traffic classification at the internet backbone, in: IEEE GI, 2009.
- [19] M. Iliofotou, P. Pappu, M. Faloutsos, M. Mitzenmacher, S. Singh, G. Varghese, Network monitoring using traffic dispersion graphs (TDGs), in: ACM IMC, 2007.
- [20] IPOQUE – Bandwidth Management with Deep Packet Inspection, Internet Study: The Impact of P2P File Sharing, Voice over IP, Instant Messaging, One-Click Hosting and Media Streaming on the Internet, 2009.
- [21] Y. Jin, S. Esam, Z.L. Zhang, Unveiling core network-wide communication patterns through application traffic activity graph decomposition, in: ACM SIGMETRICS, 2009.
- [22] W. John, S. Tafvelin, Heuristics to classify internet backbone traffic based on connection patterns, in: IEEE ICOIN, 2008.
- [23] T. Karagiannis, A. Broido, N. Brownlee, K. Claffy, M. Faloutsos, Is P2P dying or just hiding? in: IEEE GLOBECOM, 2004.
- [24] T. Karagiannis, A. Broido, M. Faloutsos, K. Claffy, Transport layer identification of P2P traffic, in: ACM IMC, 2004.
- [25] T. Karagiannis, K. Papagiannaki, M. Faloutsos, BLINC: multi-level traffic classification in the dark, in: ACM SIGCOMM, 2005.
- [26] H. Kim, K. Claffy, M. Fomenkov, D. Barman, M. Faloutsos, K. Lee, Internet traffic classification demystified: myths, caveats, and the best practices, in: ACM CoNEXT, 2008.
- [27] C. Labovitz, S. Lelkel-johnson, J. Oberheide, F. Jahanian, Internet inter-domain traffic, in: ACM SIGCOMM, 2010.
- [28] J. Ma, K. Levchenko, C. Kreibich, S. Savage, G.M. Voelker, Unexpected means of protocol inference, in: ACM IMC, 2006.
- [29] A. McGregor, M. Hall, P. Lorier, J. Brunskill, Flow clustering using machine learning techniques, in: PAM, 2004.
- [30] A. Moore, K. Papagiannaki, Toward the accurate identification of network applications, in: PAM, 2005.
- [31] M.E.J. Newman, The structure and function of complex networks, SIAM Review 45 (2003) 167.
- [32] T.T.T. Nguyen, G. Armitage, A survey of techniques for internet traffic classification using machine learning, IEEE Communications Surveys and Tutorials, 4th ed., March, 2008.
- [33] S. Sen, O. Spatscheck, D. Wang, Accurate, scalable in-network identification of p2p traffic using application signatures, in: WWW, 2004.
- [34] S. Sen, J. Wang, Analyzing peer-to-peer traffic across large networks, IEEE/ACM Transaction on Networking 12 (2) (2004) 219–232.
- [35] G. Tan, M. Poletto, J. Gutttag, F. Kaashoek, Role classification of hosts within enterprise networks based on connection patterns, in: USENIX Annual Technical Conference, 2003.
- [36] I. Trestian, S. Ranjan, A. Kuzmanovic, A. Nucci, Unconstrained endpoint profiling (Googling the Internet), in: ACM SIGCOMM, 2008.
- [37] G. Urvoy-Keller, T. En-Najjary, Challenging statistical classification for operational usage: the ADSL case, in: ACM IMC, 2009.
- [38] I.H. Witten, E. Frank, Data Mining: Practical Machine Learning Tools and Techniques, 2nd ed., Morgan Kaufman, 2005.
- [39] Y. Xie, V. Sekar, D. Maltz, M. Reiter, H. Zhan, Forensic analysis of epidemic attacks in federated networks, in: IEEE ICNP, 2006.
- [40] K. Xu, Z. Zhang, S. Bhattacharyya, Profiling internet backbone traffic: behavior models and applications, in: ACM SIGCOMM, 2005.



Marios Iliofotou received the B.S. degree in Electrical and Computer Engineering from the Technical University of Crete, in Greece, and the M.S. and Ph.D. degree in Computer Science from the University of California, Riverside, in 2007 and 2011, respectively. His current research investigates the effectiveness of graph-based algorithms for managing and protecting large IP networks. During the course of his Ph.D., Marios has worked as an intern at Microsoft Research in Cambridge, Telefonica Research in Barcelona, the Lawrence Livermore National Laboratory, Narus, and Cisco Systems.



Hyun-chul Kim is an assistant professor in the School of Computer Science and Engineering at Seoul National University. His research interests include content-centric networking, measurement, analysis and modeling of Internet traffic and topology, and user behavior patterns in a variety of large-scale applications such as IPTV, VoD, VoIP, P2P file sharing, massively multi-player online role-playing games (MMORPGs), and the Web. Kim received his Ph.D. in computer science from the Korean Advanced Institute of

Science and Technology (KAIST).



Michalis Faloutsos is a faculty member at the Computer Science Department in University of California, Riverside. 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. He is actively involved in the community as a reviewer and a TPC member in many conferences and journals. With his two brothers, he co-authored the

paper on powerlaws of the Internet topology SIGCOMM'99, which is one of the top ten most cited papers of 1999. His most recent work on peer-to-peer measurements 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 identification of abnormal network behavior. He also works in the area of Internet routing (BGP), and ad hoc networks routing, and network security, with emphasis on routing.



Michael Mitzenmacher is a Professor of Computer Science in the School of Engineering and Applied Sciences at Harvard University. He graduated summa cum laude with a B.A. in mathematics and computer science from Harvard in 1991. After studying mathematics for a year in Cambridge, England, on the Churchill Scholarship, he obtained his Ph. D. in computer science at U.C. Berkeley in 1996.

His work on low-density parity-check codes shared the 2002 IEEE Information Theory Society Best Paper Award and won the 2009 ACM SIGCOMM Test of Time Award. His textbook on randomized algorithms and probabilistic tech-

niques in computer science was published in 2005 by Cambridge University Press.



Prashanth Pappu received his Doctor of Science degree from Washington University in St. Louis where his research was focused on design of scalable routers and included QoS architectures, packet lookup and scheduling algorithms. Later, he was a Technical Leader at Cisco Systems where he led the design of Cisco's data center, storage and Ethernet switches. He's currently the Director of Product Management at Conviva Inc. that is creating a real-time, viewer-centric platform for delivering the highest resolution, broadcast

quality video on the Internet.



George Varghese worked at DEC for several years designing DECNET protocols and products (bridge architecture, Gigaswitch) before obtaining his Ph.D. in 1992 from MIT. He worked from 1993–1999 at Washington University. He joined UCSD in 1999, where he currently is a professor of computer science. He won the ONR Young Investigator Award in 1996, and was elected to be a Fellow of the Association for Computing Machinery (ACM) in 2002. Together with colleagues, he has 16 patents awarded in the general field of

Network Algorithmics. Several of the algorithms he has helped develop have found their way into commercial systems including Linux (timing wheels), the Cisco GSR (DRR), and Microsoft Windows (IP lookups). He also helped design the lookup engine for Procket's 40 Gbps forwarding engine. He has written a book on building fast router and endnode implementations called "Network Algorithmics", which was published in December 2004 by Morgan-Kaufman. In May 2004, he co-founded NetSift Inc., where he was the President and CTO. NetSift was acquired by Cisco Systems in 2005. From Aug 2005 to Aug 2006, he worked at Cisco Systems to help equip future routers and switches to detect traffic patterns for measurement and security.