Self-Learning Peer-to-Peer Traffic Classifier

Ram Keralapura[1], Antonio Nucci[1], and Chen-Nee Chuah[2]
[1] Narus Inc.   [2] University of California, Davis

Abstract—

The popularity of a new generation of smart peer-to-peer applications has resulted in several new challenges for accurately classifying network traffic. In this paper, we propose a novel 2-stage p2p traffic classifier, called *Self Learning Traffic Classifier* (SLTC), that can accurately identify p2p traffic in high speed networks. The first stage classifies p2p traffic from the rest of the network traffic, and the second stage automatically extracts application payload signatures to accurately identify the p2p application that generated the p2p flow. For the first stage, we propose a fast, light-weight algorithm called *Time Correlation Metric* (TCM), that exploits the temporal correlation of flows to clearly separate peer-to-peer (p2p) traffic from the rest of the traffic. Using real network traces from tier-1 ISPs that are located in different continents, we show that the detection rate of TCM is consistently above 95% while always keeping the false positives at 0%. For the second stage, we use the LASER signature extraction algorithm [17] to accurately identify signatures of several known and unknown p2p protocols with very small false positive rate ($< 1\%$). Using our prototype on tier-1 ISP traces, we demonstrate that SLTC automatically learns signatures for more than 95% of both known and unknown traffic within 3 minutes.

I. INTRODUCTION

Managing large networks involves several critical aspects like traffic engineering, network planning and provisioning, security, billing, and fault management. The ability of a network operator to accurately classify traffic into different applications (both known and unknown) directly determines the success of many of the above network management tasks. For example, identifying non-profitable peer-to-peer (p2p) traffic could help an Internet Service Provider (ISP) in providing better quality of service to other revenue-generating delay/loss sensitive applications. Hence it is imperative to develop traffic classification techniques that are fast, accurate, robust, and scalable in order to meet current and future needs of ISPs.

Over the past few years, peer-to-peer (p2p) networks have become extremely popular for many different applications like file (audio, video, and data) sharing, live video streaming, IP-TV, and VoIP services, among several others. In fact, many studies (like [5][9]) show that over 60% of the Internet traffic today comprises of p2p traffic. Hence, accuarately identifying p2p flows is an important task for network operators.

Traditionally, ISPs have used *port numbers* to identify and classify network traffic. For example, TCP port 80 is HTTP traffic, TCP port 1214 is Kazaa p2p traffic, and so on. This approach is easy to implement and introduces very little overhead on the traffic classifier. However, in order to circumvent detection, p2p networks have started using non-standard ports for communication [9][13][18][16]. In other words, p2p networks can choose random ports or standard ports used by other applications to send their traffic. These strategies at the application-level have made port number based traffic classification inaccurate and hence ineffective [9][16].

To address the above problems, techniques that rely on application payload signatures (i.e., comparing stored signatures to the packets from applications) were developed [12][18][16]. Although this approach is fast, accurate, robust, and reusable in different contexts (firewalls, routers, NATs, etc.), it faces the problem of *scalability*: (i) Keeping up with the number of applications that come up everyday is impractical. (ii) Reverse engineering these applications to find accurate signatures is not trivial, and hence keeping an up-to-date list of signatures becomes an herculean task for engineers.

Given the shortcomings of port- and signature-based approaches, the focus shifted to developing techniques that are less dependent on individual applications, and more dependent on capturing commonalities in the behavior of p2p applications based on layer-3/layer-4 information. We refer to these as *pattern classification* techniques. Some approaches like [21][22][15][20][11] examine the connection patterns, and classify traffic into different p2p applications using machine learning and/or clustering algorithms; others like [14] look at specific attributes of flows to group them into different applications. Although pattern classification techniques seem to be very promising (and thus deserves a strong attention from the research community), we believe that there are several open questions about their applicability in the real world: (i) Due to their dependence on statistical techniques that need multiple flows and multiple packets from each flow, the time required to detect and report the discovery of an application, is much longer when compared to signature matching techniques. (ii) Most of them are incapable of differentiating individual applications behaving in a similar fashion at the macroscopic level. For instance, they can detect p2p traffic, but cannot identify individual protocols like eDonkey, BitTorrent, or Gnutella. (iii) They are not as accurate and reliable as signature-based techniques since they are heavily dependent on the point of observation and network conditions (e.g., traffic asymmetry). (iv) Although pattern classification appears to be less resource consuming compared to the signature matching approaches (since it requires to monitor only layer-3/layer-4 data), it is in fact not true. Pattern classification requires maintaining considerably larger number of states in memory for processing, and thus severely limits their effectiveness in operating at very high speeds.

In this work, we address the problem of identifying traffic originating from known and unknown p2p networks. In particular, we focus on: (i) *real-time* identification of p2p

Name	Network	Duration	Flows	E/G/B/S/K Flows $(\times 10^3)$
Trace-1	ISP-A	591.209s	7.26×10^6	84.2 / 31.4 / 26.4 / 7.5 / 0.6
Trace-2	ISP-B	1067.927s	10.21×10^6	6.3 / 1.3 / 70.6 / 2.2 / 0.3
Trace-3	ISP-B	1083.837s	11.61×10^6	5.6 / 1.9 / 68.5 / 3.1 / 0.4
Trace-4	ISP-B	1183.637s	11.35×10^6	6.1 / 1.4 / 75.4 / 1.7 / 0.3

TABLE I

DETAILS OF THE FOUR DATA TRACES USED IN THIS WORK (E-EDONKEY; G-GNUTELLA; B-BITTORRENT; S-SKYPE;K-KAZAA).

traffic in *large* networks (for example, tier-1 and tier-2 ISPs) by monitoring the traffic at the *network edge*, and (ii) p2p networks that use superpeer technology (edonkey, gnutella, etc). In this paper, we propose a novel two-stage p2p traffic classifier called *Self-Learning Traffic Classifier* (SLTC) that brings together the benefits of signature matching (speed, accuracy, and reusability) and pattern classification (scalability) techniques. In the first stage, SLTC separates p2p traffic from the rest of the traffic by exploiting the general behavior of superpeer-based p2p protocols, and in the second stage, it automatically extracts payload signatures to identify specific p2p protocols. Note the second stage is required to separate p2p *protocols* (like edonkey, gnutella, kazaa, etc.) from each other. SLTC populates the extracted signatures into a signature database that will be used to classify all future flows. Thus all flows for which SLTC has a signature in the database bypass the expensive pattern classification step. SLTC can automatically *learn* (i.e., classify and extract signatures) both known and unknown p2p applications in a matter of minutes.

To the best of our knowledge, this is the first work to propose a multi-stage, self-learning, real-time p2p traffic classification system that can be used in high speed networks with minimum manual intervention. Our main contributions are:
• We propose a 2-stage SLTC system that can quickly learn known and unknown p2p applications and classify them in real-time. For the first stage, we propose a new pattern classification algorithm, called *Time Correlation Metric* (TCM) that explores the *temporal* correlation of incoming and outgoing p2p flows. TCM first identifies p2p nodes, and subsequently classifies flows to/from these nodes as p2p flows. We show how this new concept clearly outperforms previous metrics in (i) discovering p2p nodes with an accuracy well above 95% with 0% false positive, and (ii) distinguishing p2p nodes as either peers or super-peers. (Section V).
• We comprehensively explore the feasibility of SLTC with many tier-1 ISP packet traces. Our experiments show that SLTC can learn over 95% of all p2p traffic in less than 3 minutes. Once SLTC learns about an application, future flows that belong to the application are directly classified using application payload signatures and hence do not go through the two stages in SLTC. Furthermore, the 90-percentile detection lag (i.e., the total time from detecting the first packet of a p2p flow to extracting a signature for the application corresponding to the flow) is less than 60 seconds (Section VII).

II. DATA DESCRIPTION

We collect and analyze 4 packet traces captured from two tier-1 ISP networks (say ISP-A and ISP-B) that are located in different continents (Table I). For both the ISPs, we captured

all packets (with no sampling) between the ISP and one of its customers (a tier-2 ISP).

Validating our classification algorithm needs *ground truth*. In other words, we should be able to classify the traffic in these traces using an alternative method (other than the algorithms in SLTC), so that we can compare the results from SLTC algorithms. To accomplish this, we built a *L7 protocol analyzer* (L7PA) based on the application payload signatures publicy available at L7-Filter [6] and used this to verify the accuracy of SLTC algorithms. L7PA has signatures for 25 different applications that include both p2p (BitTorrent, Gnutella, EDonkey, Skype, and KaZaa) and non-p2p (HTTP, SMTP, POP3, DNS, etc.) protocols. Table I also shows the number of flows for different p2p protocols identified by L7PA in each of the four traces. We use these flows as the ground truth in the rest of this paper.

III. RELATED WORK AND CHALLENGES

A. Peer-to-Peer Traffic Classification

Given the shortcomings of the basic approaches, there has been a lot of effort in developing p2p traffic classification techniques that rely just on the layer-3/layer-4 information. Some approaches (like [21][22][15][20][11]) examine the connection patterns at layer-3 and classify traffic into different applications using machine learning techniques and/or clustering algorithms, while others (like [13], [14]) look at specific attributes of flows to group them into different applications. The authors in [13] propose many flow-based heuristics to identify p2p nodes: (i) p2p nodes use both TCP and UDP protocols as their transport layer protocol, (ii) p2p node are characterized by both incoming and outgoing connections, and (iii) the ratio of the number of source IP to source port for all incoming flows into a p2p node approaches 1. However, as we show later in this section, these heuristics lead to false positives[1] and false negatives.

Although the above approaches result in high detection rates, their main limitation is that they are infeasible for *real-time* classification in high speed ISP networks for three reasons. (i) They rely on time consuming algorithms (like statistical clustering, machine learning, etc.) that have to be applied on every flow seen by the classifier; thus keeping up with the traffic rate becomes extremely difficult. (ii) They are designed to identify high-level application classes, but cannot identify individual applications, an important requirement for network operators to prioritize traffic. (iii) None of them can effectively identify new (i.e., currently unknown) applications that come up in the future.

Our approach addresses the above issues and is geared towards real-time identification of known and unknown p2p applications in high speed networks where asymmetric routing is commonly used. Our algorithm is simple, fast, accurate, and resistant to data obscured by asymmetric routing.

---

[1]The authors in [13] acknowledge that the heuristics can lead to false-positives and provide refinements to the basic heuristics to minimize false-positives. However, we find that in the case of ISP networks there could also be a lot of false negatives. Please see [7] for more details.

B. Peer-to-Peer Networks

In this work, p2p traffic refers to the traffic originating from: (*i*) *unstructured* p2p networks where different peers join and leave the network as and when they please, (*ii*) dynamic p2p networks that are used to exchange files, music, video, and other forms of data, and (*iii*) p2p networks that use superpeer technology to manage their network. Examples include eDonkey [2], Gnutella [4], BitTorrent [1], etc.

Unstructured p2p networks are distributed in nature providing an infrastructure to exchange files, music, and video with each other without relying on any centralized servers. Many popular p2p networks have several million users at anytime, and hence a *completely distributed* approach to finding and exchanging information leads to network meltdown. Most of the successful p2p networks that exist today adopt the strategy of constructing *hybrid* networks, where the p2p network elects a few nodes as *leaders* for a group of nodes based on the nodes' computing/network resources. These leaders are usually referred to as *superpeers* or *ultrapeers*.

Superpeers are typically connected to several other superpeers and the main objective is to ensure that these superpeers (and hence the peers connected to them) are connected to the rest of the network. We can think of this architecture of p2p networks as a two-level hierarchy. The first level contains all the superpeers connected to several other superpeers in the same level. The second level contains peers connected to one or more superpeers in the first level. Note that these peers at the second level may or may not be connected to other peers in the same level. This architecture ensures that when peers join or leave a network, the impact on the network (in terms of connectivity of other peers) is minimal. However the impact is higher when superpeers leave the network. Hence nodes that have significantly higher uptimes are chosen to be superpeers.

Although the actual functionality of a superpeer varies depending on the particular p2p application, in general, a superpeer acts as a gateway to the rest of the network for the group of peers that are connected to it.

C. Challenges to P2P Traffic Detection

Although p2p networks are application layer networks built on top of the IP layer, traffic from these networks behave very similar to the rest of the Internet traffic and is virtually indistinguishable. Hence, most strategies proposed in the past for classifying p2p traffic based on only layer-3/layer-4 information rely on first detecting *nodes* that are running p2p applications, and then identifying p2p traffic based on these p2p nodes. In this subsection, we present the most obvious metrics (that are feasible to be used for p2p traffic classification)[2] for identifying p2p nodes, and show why these metrics fail to accomplish their objective in the context of our problem definition (i.e., real-time superpeer-based p2p traffic classification at network edges).

A common strategy adopted by most p2p networks to get around the connectivity problem introduced by firewalls is to use both TCP and UDP protocols on any of the open ports. Furthermore, to optimize their performance, p2p nodes typically use both TCP and UDP protocols for control, signaling, and/or data flows. For example, a Skype peer connects to its superpeer usig both TCP and UDP [8]. Another characteristic that distinguishes a p2p node from a node that does not run any p2p applications is the p2p node's ability to act as both a client and a server. Several heuristics have been proposed to take advantage of these properties of p2p nodes [13].

However, there are several problems while using the above heuristics to detect p2p nodes: (*i*) **False Positives:** Several other protocols in the Internet, like DNS, gaming, streaming, IRC, etc., also exhibit these properties. In other words, these non-p2p applications also use both TCP and UDP protocols to communicate between node pairs. As an example consider Figure 1(a), where several nodes running DNS protocol in Trace-1 also have both TCP and UDP connections. Also, nodes running these non-p2p applications can both accept and open connections to other nodes. Figure 1(b) once again uses DNS nodes as an example to show that nodes running non-p2p applications can have both incoming and outgoing connections. Note that we saw similar results for a number of other protocols like SMTP, gaming, etc. Hence the above heuristics could lead to a lot of false positives [13], [10]. (*ii*) **False Negatives:** All p2p nodes (or p2p node pairs) do not always satisfy the above heuristics. For example, not all p2p node pairs use both TCP and UDP protocols to talk to each other. Several p2p protocols use TCP port 80 (a port most likely to be open in almost every firewall) as a way to bypass firewalls and hence may not use both TCP and UDP protocols. Also, several p2p nodes may not be observed (from the perspective of the monitoring point) to act as both server and client[3] (see Figure 1(c)). Thus there could be a lot of false negatives while using these heuristics as well.

We propose a novel lightweight approach to easily detect p2p nodes for the problem defined in Section I. We present the intuition and algorithm for our approach in Section V.

IV. SLTC ARCHITECTURE

SLTC is a 2-tiered system comprised of a distributed collection tier and a centralized processing tier. Data is collected directly off-the-wire using high speed *passive* listeners, called *High Speed Monitors* (HSM). These monitors passively observe network traffic on different links, and try to classify traffic using application payload signatures. If a monitor can successfully identify a transiting flow using known signatures, then the flow is marked as "known", and no further analysis is necessary to classify it. However, if a flow cannot

---

[2]Several other techniques have been proposed (like [21], [22], [15], [20], [11], [14]) in the literature. However we believe that these techniques are infeasible for real-time P2P traffic classification in high speed networks.

[3]Even if all p2p nodes act as both clients and servers, observing all these connections depends on the location/s of the monitoring equipment. In several cases, the location of the monitoring equipment could force us to believe that p2p nodes are acting only as servers or only as clients. Hence we believe it is reasonable to assume that not all the connections can be observed to make accurate conclusions here.

(a) Nodes in Trace-1 running DNS use both TCP and UDP.

(b) Number of incoming and outgoing connections in nodes running DNS (Trace-1).

(c) Number of incoming/outgoing connections in nodes running p2p applications (Trace-1).
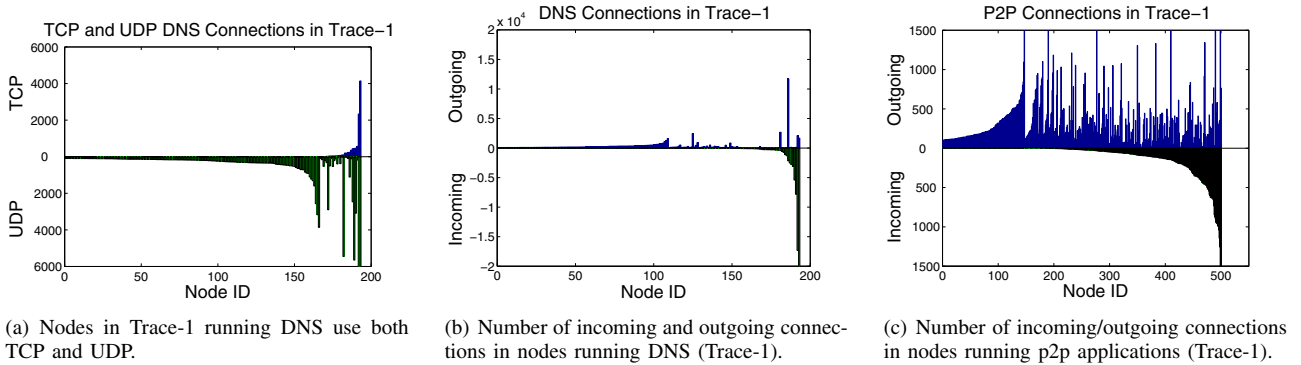
Fig. 1. False positives and false negatives when using basic metrics for p2p node detection.

be successfully classified, then HSM forwards layer-3/layer-4 information of the flow to the centralized server, called *Logic Server* (LS). This centralized server runs our p2p traffic classification algorithm, TCM, to identify p2p superpeers. LS conveys the findings back to the HSMs, and the HSMs in turn starts forwarding all the packets (including application payload) of flows destined to the identified p2p superpeer to the LS. The LS forwards this to the SE algorithm that will analyze several flows destined to the same p2p superpeer and extracts a L7 signature. These signatures are then sent back to the HSMs, and the HSMs use them to identify any future flow that belongs to the application. Note that the signature from one p2p superpeer can be used to identify all flows that belong to the particular p2p application that the superpeer is running. Please see [7] for more details about the architecture.

V. PEER-TO-PEER PATTERN CLASSIFIER

A. P2P Traffic Classification: Intuition and Approach

Our approach to identifying p2p traffic class relies on the following observations in *hybrid* p2p networks (i.e., p2p networks that use superpeer technology). When a peer (or host) joins a p2p network, it typically connects to one or more servers and/or superpeers. If the peer connects to a server (as in the case of a few hybrid p2p networks) at the start, then the server provides the peer with the superpeer contact information. The peer will eventually contact the superpeer to let the superpeer know of its arrival. Figure 2 depicts the above process. When a new peer, **Peer A**, joins the p2p network, it talks to a superpeer, **Superpeer S**, in its host cache (i.e., a table containing all the neighboring peers). The information about the superpeer could already be in **Peer A**'s host cache due to the past activity of the peer in the network, or it could be obtained from a centralized database by first connecting to a central server. Either ways, **Peer A** ultimately connects to **Superpeer S**, and sends the information that can be used by other peers to contact **Peer A**. As soon as **Superpeer S** receives this information, it forwards the information to other peers (like **Peer B**), and superpeers that are connected to it. This process of disseminating the peer contact information is critical in p2p networks for two reasons: (i) Fault tolerance: A typical p2p network experiences a lot of churn (i.e., peers joining and leaving the network). In Figure 2, if **Superpeer S**

decides to leave the network, then **Peer A** loses connectivity to the rest of the network. If other superpeers know about **Peer A**, then they can take over the responsibility from **Superpeer S**, thus providing all the required services to **Peer A**. (ii) File download/upload: Peers in the network need **Peer A**'s contact information to upload/download files.

Based on the above observations, when a new node arrives into the p2p network, a superpeer accepts a connection from the node, and subsequently opens a connection to one or more other nodes in the network. From the perspective of a superpeer, *an incoming connection is closely followed in time by one or more outgoing connections*. Hence our hypothesis is that by observing connections coming into and leaving a node in close succession we can accurately identify superpeers in p2p networks. We call this as the *time correlation metric* (TCM), i.e., a metric that captures the temporal correlation between the incoming and outgoing connections in p2p protocols.

There are several reasons why we believe that TCM is ideal in the context of our problem (i.e., for identifying superpeer-based p2p traffic in high speed networks in real-time by monitoring peering links):

• **Churn in P2P Networks.** P2P networks experience a lot of churn [19] and support constant searches. Since TCM aims to exploit the p2p network behavior during these commonly occurring events, we believe that TCM can be very successful.

• **Location of HSM.** HSM monitors all bidirectional traffic on a peering link. Given that peers in p2p networks typically connect to random superpeers, there is a very high probability that several incoming and outgoing connections from the *same* superpeer crosses a peering link multiple times. Notice that,
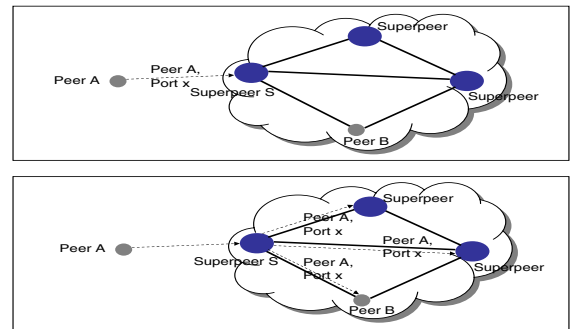

Fig. 2. A peer joining (or searching for information) a p2p network.

TCM(CURTIME, NEWFLOW(SIP, DIP, SPORT, DPORT, L4PROT))

1: $OutgoingFlows(sIP, sPort).add(curTime, newFlow)$
2: $IncomingFlows(dIP, dPort).add(curTime, newFlow)$
3: **for all** (F such that $F \in IncomingFlows(sIP, sPort)$) **do**
4:    **if** (($consideredDstIP$ does not contain dIP) and ($curTime - F.time$) $< T_{th}$) **then**
5:       $consideredDstIP.add(dIP)$
6:       $pattern(sIP, sPort) \leftarrow pattern(sIP, sPort) + 1$
7:       **if** ($pattern(sIP, sPort) > P_{th}$) **then**
8:          $repetition(sIP, sPort) \leftarrow repetition(sIP, sPort) + 1$
9:          **if** ($repetition(sIP, sPort) > R_{th}$) **then**
10:             $P2PNodePortPair.add(sIP, sPort)$

Fig. 3. TCM Algorithm

unlike other approaches, in TCM, *we do not need to observe all the flows, but only require to monitor a few flows for every superpeer*. Once again, this results in an advantage for TCM.

Another important point that we wish to draw the reader's attention to is that monitoring network traffic on peering links also eliminates *false positives* to a large extent. For instance, DNS nodes could also exhibit the TCM property as described earlier. Consider a hierarchical DNS system where higher level DNS servers are located outside a network. A recursive DNS query to a DNS server in this system could result in the server opening new connections to other DNS servers. However, the natural association of network borders and the DNS servers in a hierarchial system, ensures that the HSM (sitting on a peering link at the network edge) *either captures the incoming or the outgoing connections but not both*. The same is true for other applications like smtp, pop3, etc.

The TCM algorithm based on the above approach is shown in Fig 3. The algorithm is characterized by 3 parameters:
- **TCM Time Threshold** (T_{th})**.** This represents the maximum time difference between incoming and outgoing connections in superpeers. A large value of T_{th} implies that we will group together unrelated incoming and outgoing flows. A very small value implies that we do not group together even the flows that are correlated. Hence choosing an optimal value of this metric is critical to the effectiveness of TCM.
- **TCM Pattern Threshold** (P_{th})**.** As we explained earlier (in Figure 2), *every* incoming connection to the superpeer from a new peer (or a search query from the existing peer) results in *several* outgoing connections from the superpeer. P_{th} represents the number of outgoing connections that should be temporally correlated with an incoming connection to assume that a TCM pattern has occurred. A very high value of P_{th} could lead to a lot of false negatives (i.e., p2p superpeers not identified as superpeers), where as a small value could lead to false positives (non-p2p nodes identified as superpeers).
- **TCM Repetition Threshold** (R_{th})**.** If a superpeer is observed for a long period of time, the TCM pattern (i.e., one incoming connection resulting in several outgoing connections) should occur several times. We use R_{th} as a parameter to specify the number of times the TCM pattern should be observed before declaring a node to be a superpeer. Note that a very small value could imply that non-p2p nodes could be included in the superpeer set by pure coincidence. However, a large value could once again lead to false negatives.

B. Reducing False Negatives

An assumption that we make in the TCM heuristic is that when a new peer establishes a connection with the superpeer, the superpeer opens *new* connections to other existing superpeers and/or peers to convey the information about the new peer. However, in reality this might not always be true. That is, the superpeer may convey the information about the new peer using *existing* connections to other superpeers/peers.

Figure 4 shows the total data rate of p2p connections that lasted for more than 10 minutes in Trace-4 (the graph was generated using our ground truth and zooms in to show the region of interest). We can see that there are several connections whose overall data rate is very small (1-2 bytes/sec). These connections last more than 10 minutes, but only exchange 1000-2000 bytes of data. This suggests that these are long lasting *control* connections that carry small control data. Thus, ignoring communications on long lasting connections results in several false negatives in TCM. Hence, in our TCM algorithm, instead of always looking for *new* outgoing connections, in addition, we also look at existing connections carrying small control packets. Furthermore, we consider small outgoing control packets on existing connections only if: (i) the connection lasts for a long time, and (ii) the average packet size of the connection is also small (< 150 bytes). This heuristic eliminates the possibility of considering small packets from non-control flows. In our experiments, we incorporate these changes to the algorithm in Figure 3.
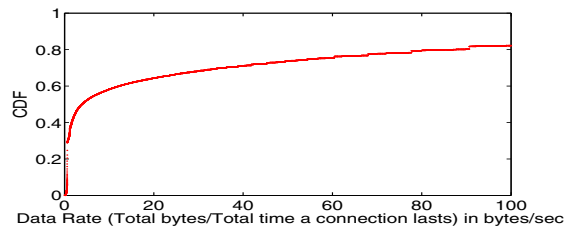

Fig. 4. Overall data rate of p2p flows in Trace-4 that last more than 10 mins.

VI. SIGNATURE EXTRACTOR (SE)

The signature extraction component resides in the logic server (LS) and its goal is to *automatically* extract signatures for p2p applications. Note that SE is critical for self-learning mode of SLTC for 2 reasons: (i) It helps in differentiating p2p applications from each other, and (ii) It helps to bypass the expensive pattern classification step for known flows. We use the LASER algorithm to extract signatures from packet payloads [17]. LASER uses the least common subsequence (LCS) algorithm that is popularly used in DNA sequencing. LASER is very efficient and accurate in extracting signatures for unknown applications [17]. Due to lack of space we omit details about the SE component. See [7] for details.

VII. SYSTEM EVALUATION

To evaluate the proposed SLTC architecture, we built a prototype of SLTC (both HSM and LS components). We replayed all the four traces (described in Section II) and used them as input to the system. In Section VII-A, we show how

we select the various TCM parameters. In Sections VII-B and VII-C, we use the *ground truth* generated by the L7 protocol analyzer (L7PA) to evaluate the *accuracy* of TCM and SE algorithms. In Section VII-D, we evaluate the performance of the end-to-end SLTC system by exploring all (known and unknown) flows in the four traces.

A. TCM Parameter Tuning

As mentioned in Section V, the TCM algorithm depends on 3 parameters - T_{th}, P_{th}, and R_{th}. The choice of values for these parameters directly influence the accuracy and efficiency of TCM. In this section, we tune the values of these parameters using our L7PA such that the output of TCM results in high detection rate and low false positive rate.

We use Trace-1 for all the parameter tuning experiments, however, the results from the other three traces were very similar. To eliminate the dependence of our parameter values on the length of the trace (i.e., the trace time interval), we first split the trace into multiple segments each of which is 180 seconds long. We compute the detection and false positive rates for each of these segments. The results in Figures 5, 6, and 7 show the average detection and false positive rates for all the 180-second segments in Trace-1.

Figure 5 shows the detection and false positive rates as a function of the time threshold (T_{th}) for different values of P_{th} and $R_{th} = 2$. We can see that large values of T_{th} results in higher detection rates, but also results in higher false positive rates. Also, the maximum detection rate decreases as the value of P_{th} increases. In other words, for P_{th} values between 2 and 4, the maximum detection rate reaches 100%. However, for $P_{th} = 5$, the maximum detection rate falls below 100%, showing that we will be unable to detect all the superpeers using TCM if we set a large value for P_{th}. Finally, from Figure 5, we can clearly see that there is no region in the graph where the detection rate is 100% and the false positive rate is 0%. In other words, we cannot find any parameter values that result in optimal detection and false positive rates.

Figures 6 and 7 are similar to Figure 7, but for $R_{th} = 3$ and $R_{th} = 4$ respectively. From Figure 7, we can see that for $R_{th} = 4$, the maximum detection rate is always less than 100% irrespective of the values of T_{th} and P_{th}. Hence increasing the value of R_{th} beyond 3 will not result in optimal detection rate. Finally, from Figure 6, we can see that the optimal detection and false positive rates can be obtained when $T_{th} \in (2, 3]$, $P_{th} = 4$, and $R_{th} = 3$. Hence, in the experiments in the rest of this paper we use $T_{th} = 2.5s$, $P_{th} = 4$, and $R_{th} = 3$.[4]

B. Peer-to-Peer Traffic Classifier

Although our L7PA has signatures for several p2p and non-p2p protocols, the list is surely not exhaustive. Hence the output of L7PA contains several flows that are just marked as "unclassified" TCP/UDP traffic. We represent the set of nodes identified by L7PA that belong to p2p networks (i.e., Gnutella,

[4]The ideal parameter values vary with the deployment scenario. In our case, we are monitoring peering links between a tier-1 and tier-2 ISP. The ideal parameter values will change with the location of the monitoring point.
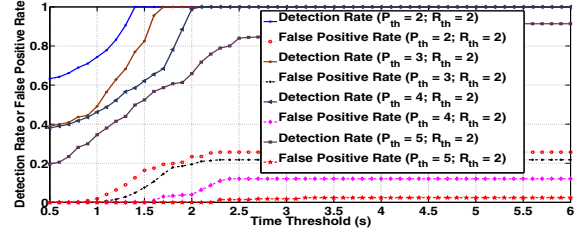

Fig. 5.   Detection rates and false positive rates for $R_{th} = 2$.

eDonkey, BitTorrent, Skype, and Kazaa) by P, the set of nodes belonging to non-p2p applications (like HTTP, DNS, SMTP, IMAP, POP3, etc.) by N, and the unclassified nodes by O. In order to eliminate ambiguity of a node belonging to multiple sets, we ensure that if a node belongs to P then it does not belong to the other two sets. Similarly, if a node belongs to O, then it is not included in N. Hence all the three sets are *mutually exclusive* of each other.

Let T represent the total time of a data trace. We represent T as a sum of several *time intervals* of fixed length z. We represent the total number of such intervals in any trace file by n. Now, let P_i denote the set of p2p nodes identified by L7PA in the i^{th} time interval where, $i \in 1 \ldots n$. Similarly, N_i and O_i represent the set of nodes in the non-p2p and unclassified set identified in the i^{th} interval. Note that the sets P_i, N_i, and O_i constitute the ground truth in the i^{th} interval. Similarly, let D_i represent the set of p2p nodes detected by our TCM algorithm in the i^{th} interval.

We use 3 metrics to evaluate the accuracy of TCM: *detection rate* (DR), *false positive rate* (FP), and *false negative rate* (FN). We define the detection rate in any interval, DR_i, and overall detection rate until (and including) the interval i, DR_i^{full}, as:

$$DR_i = \frac{|D_i \bigcap P_i|}{|P_i|}; DR_i^{full} = \frac{|\{\bigcup_{k=1..i} D_k\} \bigcap \{\bigcup_{k=1..i} P_k\}|}{|\bigcup_{k=1..i} P_k|} \quad (1)$$

Similarly, we define false positive and false negative rates as:

$$FP_i = \frac{|D_i \bigcap N_i|}{|N_i|}; FP_i^{full} = \frac{|\{\bigcup_{k=1..i} D_k\} \bigcap \{\bigcup_{k=1..i} N_k\}|}{|\bigcup_{k=1..i} N_k|} \quad (2)$$

$$FN_i = \frac{|P_i - D_i|}{|P_i|}; FN_i^{full} = \frac{|\{\bigcup_{k=1..i} P_k\} - \{\bigcup_{k=1..i} D_k\}|}{|\bigcup_{k=1..i} P_k|} \quad (3)$$

We evaluate the accuracy of our TCM algorithm by directly replaying all flows in the four traces into the TCM component in the LS. Once the TCM algorithm identifies the
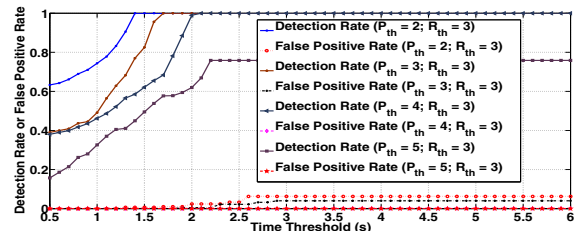

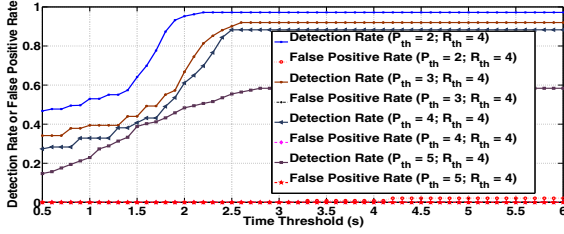Fig. 6.   Detection rates and false positive rates for $R_{th} = 3$.

Fig. 7. Detection rates and false positive rates for $R_{th} = 4$.

Fig. 9. Overall TCM detection rate for different P2P applications.

$< superpeer, port >$ pairs in the trace, we flag all the other nodes that connect to $< superpeer, port >$ pairs as a p2p node. We use this information in Equations (1)-(3).

The top graph in Fig 8 shows the overall detection rate (DR_i^{full}) as a function of time for all the traces. The detection rate is computed every 5 seconds (i.e., $z = 5$). In other words, we accumulate all our findings in every 5-second interval and use them to update the detection rate at the end of the interval. We can see that the overall detection rate is between $40-60\%$ after the first time interval, but it increases beyond 90% within 3 mins in all of the traces. The overall detection rate reaches close to 100% in most of the traces within 10 mins.

The bottom graph in Fig 8 shows the detection rate (DR_i) in every interval. We can once again notice that the detection rate in any interval reaches 95% within 3 mins and remains close to 100% after 7 mins. The main take-away point here is that after the first few minutes, the detection rate in any interval remains very close to 100%, showing that TCM can identify all p2p nodes in the network very effectively.

Fig 9 shows the overall detection rate of 5 popular p2p applications in the Internet for which we have the ground truth. We can see that the overall detection rate for most of them in all the four traces is over 95%. Note that Fig 9 shows the overall detection rate after full trace replay. However, the detection rate in every interval for all the applications is much higher (over 99%) after the first few intervals.

Based on our four traces, we find that some of the applications (like Kazaa) are not very popular. We found only a few Kazaa flows in all our traces (See Table I). Given that the TCM parameters are tuned to capture superpeers that are significantly active, we miss a few nodes in these applications that are not popular; hence the detection rates are smaller
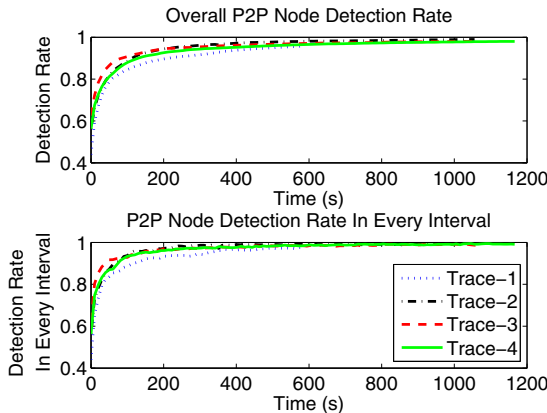

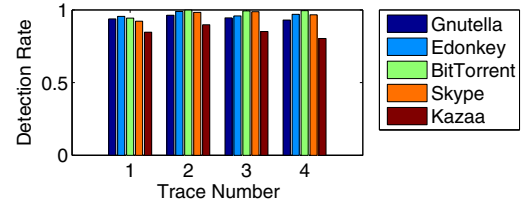Fig. 8. Overall and per-interval detection rates of TCM vs. time

compared to the other applications. However, from Figure 8, we can clearly see that the p2p detection rate is still very high.

Although our detection rates are very high, the accuracy of TCM depends on the false positive and false negative rates as well. **The false positive rates in all our above experiments were consistently *zero* for all the four traces.** The TCM algorithm did not identify any node that belonged to the non-p2p set as a p2p node. The false negative rate, by definition in Eqn (3), is the complement of the detection rate, i.e., $DR_i = 1 - FN_i$, and $DR_i^{full} = 1 - FN_i^{full}$. Hence, we do not explicitly plot the results for the false negative rate.

C. Signature Extractor

Due to lack of space we do not present the results from the SE component here. However, in [7], we show that signatures extracted are very robust with high recall rate.

D. SLTC System

There are 2 main objectives for SLTC as a system: (i) Learn signatures for *all* known and unknown p2p applications seen by the HSM, and use them to classify future incoming flows, and (ii) Learn the signatures as quickly as possible so that the number of "unclassified" flows can be minimized. For experiments in this subsection we replay all flows in the trace files into the HSM, and any flow that cannot be directly classified in the HSM is sent to the TCM component in the LS. The output (i.e., $< ip, port >$ pairs of superpeers) from the TCM is used by the HSM to send packets to the SE component. The SE algorithm finds the signatures of p2p applications and populates the signature database in the HSMs.

We demonstrate how well SLTC meets both its objectives using two metrics: (i) *Fraction of Flows* that are "unclassified" in every interval. This represents the fraction of p2p flows that SLTC has not learnt about, i.e., the p2p flows sent to LS for classification and signature extraction, and (ii) *Time Lag*, i.e., the total time taken to extract a signature and populate the database after seeing the first packet of a flow.

The top graph in Fig 10 shows the fraction of "unclassified" flows in SLTC, i.e., the p2p flows that are sent to the logic server (LS) for classification, as a function of time. Note that at time 0 there are no signatures in the database and hence all flows are sent to the logic server. However, as time goes by, LS extracts more and more signatures, thus ensuring that most of the traffic is classified by the HSM. In fact, more than 90% of the p2p traffic is classified within 1 minute. We can clearly see that the fraction of traffic entering LS decreases over time. Note that the fraction of traffic sent to LS steadies at about 5% in all the four traces. This means that our SE component is
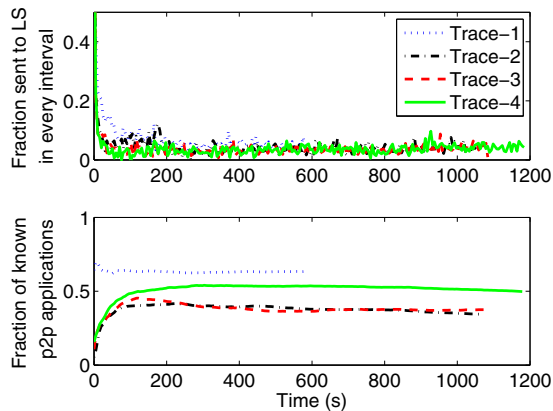
Fig. 10.   (a) Classified flows; (b) Known applications in classified flows.

Protocol	String
Gigaget	0x29000000 0x00000007000000434f4e4e454354

TABLE II
NEW P2P APPLICATION SIGNATURE - GIGAGET.

unable to extract signatures for these flows even after repeated attempts. These flows could be either encrypted or simply do not have a signature. We defer further exploration of this issue as a part of our future work.

The bottom graph in Fig 10 shows the fraction of the "classified" flows that are "known". That is, among all the classified flows (i.e., flows for which SLTC learns a signature) from the top graph in Fig 10, there are some flows that belong to "known" applications and others belong to "unknown" or "new" p2p applications. We use BitTorrent, Gnutella, eDonkey, Skype, and Kazaa as known p2p applications and classify the rest as unknown. From the bottom graph in Fig 10, we can see that more than 40% of the classified p2p applications are unknown for Trace-1 whereas the percentage increases to about 60% for the other traces. Learning unknown or new p2p applications is one of the key features of SLTC and from the bottom graph of Fig 10 we can clearly see that SLTC has learned several new p2p applications. In fact, we manually explored the signature of an unknown application and found that it belongs to Gigaget p2p network (Table II) [3].

Figure 11 shows the CDF of the total time lag. Total time lag is the sum of the times taken TCM and SE algorithms, and both these times depend on the parameters set in those components. For example, the time delay in TCM is characterized by three parameters T_{th}, P_{th}, and R_{th}. Changing the values of these parameters could result in a tradeoff between accuracy and time lag. Using the default parameter values for TCM and the SE algorithm, in Figure 11, we can see that classification and signature extraction of over 90% of the flows in all the traces takes less than one minute, making it feasible to use this architecture in real-time with strict time constraints.

VIII. DISCUSSION AND CONCLUSIONS

In this paper, we presented SLTC, a self-learning p2p traffic classifier that can learn known and unknown applications with minimum manual intervention. We showed that SLTC can learn over 95% of p2p applications in a few minutes. Although
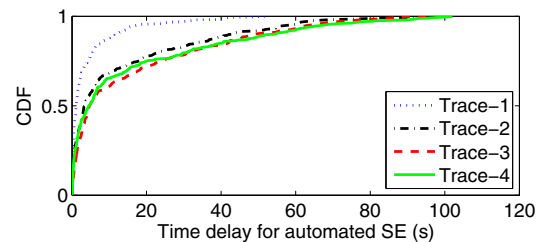

Fig. 11.   CDF of the total time lag.

our focus in this paper was on p2p applications, we strongly believe that this framework can be used for classifying other classes of Internet traffic as well.

We presented a simple, light-weight, and effective algorithm, called TCM, for superpeer-based p2p traffic identification using temporal correlation of flows. We showed that the detection rate is very high with no false positives. One of the main reasons for having no false positives is the fact that the monitoring points are at the network edge. Unlike p2p-protocols, most of the non-p2p protocols in the Internet (like smtp, ftp, pop3, imap, http, etc.) are typically configured such that they are aware of network borders.

In this work we used packet traces with no sampling. However, we envision TCM to work even with sampling. We plan to explore this further as a part of our future work.

REFERENCES

[1] Bittorrent. http://www.bittorrent.com/.
[2] Emule project. http://www.emule-project.net/.
[3] Gigaget. http://www.gigaget.com/.
[4] Gnutella. http://www.gnutella.com.
[5] Ipoque survey. http://www.ipoque.com/.
[6] L7 filter. http://l7-filter.sourceforge.net/.
[7] A novel architecture for self-learning traffic classifier. Narus Technical Report.
[8] S. Baset and H. Schulzrinne. An Analysis of the Skype Peer-To-Peer Internet Telephony Protocol. In *Tech Report, Columbia Univ.*, 2004.
[9] CacheLogic. The True Picture of Filesharing. http://www.cachelogic.com/home/pages/research/p2p2004.php.
[10] F. Constantinou and P. Mavrommatis. Identifying Known and Unknown Peer-to-Peer Traffic. In *IEEE NCA*, 2006.
[11] J. Erman, M. Arlitt, and A. Mahanti. Traffic Classification using Clustering Algorithms. In *ACM Sigcomm Minenet*, 2006.
[12] P. Haffner, S. Sen, O. Spatscheck, and D. Wang. ACAS: Automated Construction of Application Signatures. In *ACM Minenet*, 2005.
[13] T. Karagiannis, A. Broido, M. Faloutsos, and K. Claffy. Transport Layer Identification of P2P Traffic. In *ACM IMC*, 2004.
[14] T. Karagiannis, K. Papagiannaki, and M. Faloutsos. BLINC: Multilevel Traffic Classification in the Dark. In *ACM Sigcomm*, 2005.
[15] A. McGregor, M. Hall, P. Lorier, and J. Brunskill. Flow Clustering Using Machine Learning Techniques. In *PAM*, April 2004.
[16] A. Moore and K. Papagiannaki. Toward the Accurate Identification of Network Applications. In *PAM*, March 2005.
[17] B-C. Park, Y.J. Won, M-S. Kim, and J.W. Hong. Towards Automated Application Signature Generation for Traffic Identification. In *IEEE/IFIP NOMS*, 2008.
[18] S. Sen, O. Spatscheck, and D. Wang. Accurate, Scalable In-Network Identification of P2P Traffic Using Application Signatures. In *WWW2004*, May 2004.
[19] D. Stutzbach and R. Rejaie. Understanding Churn in Peer-to-Peer Networks. In *IMC*, 2006.
[20] N. Williams, S. Zander, and G. Armitrage. A Preliminary Performance Comparison of Five Machine Learning Algorithms for Practical IP Traffic Flow Classification. In *CCR*, October 2006.
[21] S. Zander, T. Nguyen, and G. Armitage. Automated Traffic Classification and Application Identification using Machine Learning. In *LCN*, 2005.
[22] S. Zander, T. Nguyen, and G. Armitage. Self-Learning IP Traffic Classification Based on Statistical Flow Characteristics. In *PAM*, 2005.