Traffic classification for application specific peering

Balachander Krishnamurthy and Jia Wang

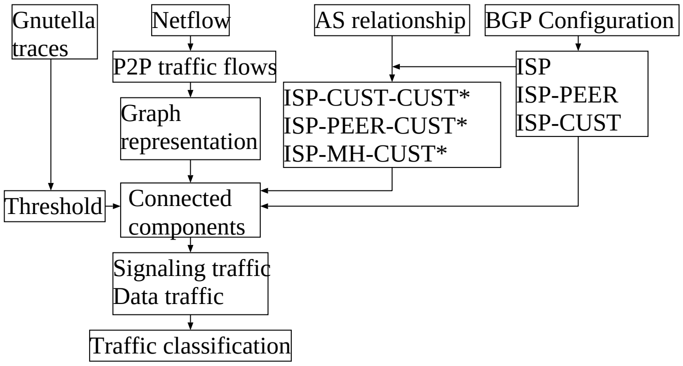

Fig. 1. The process of automatic traffic classification.

I. Introduction

Peer-to-Peer (P2P) applications [4], [6], [2], [3] have increased traffic significantly. We propose *application specific peering* that allows searches for a resource (using existing or modified P2P protocols) to be directed to a copy of resource on an ISP's network, using a novel technique that automatically classifies traffic (shown in Fig 1) as belonging to an ISP, its customers, peers, or their customers. We examined a large amount of traffic belonging to three popular P2P protocols and used a graph transformation to detect connected components. We partition traffic in each component into signaling and data sections based on request/response sizes of P2P protocols and extract high traffic volume (by bytes) entities by examining large connected components in the graph. We discuss the important components of our traffic classification system here.

We analyzed router-level data from multiple Internet Gateway Routers (IGRs) across AT&T IP backbone using Cisco's *NetFlow* services. Note that statistics are from a subset of IGRs at some AT&T's peering links. They are not representative of overall traffic pattern on AT&T IP backbone. We examined DirectConnect[2], Gnutella[3], and FastTrack[4], extracting records that matched the default application ports (source or destination, 411/412, 6346/6347, and 1214 (FastTrack) respectively), involving TCP traffic. Data was gathered over three separate months in 2001, each lasting 5-7 days. The number of netflow records varies from a low of 0.5M (DirectConnect) to a high of 340M (FastTrack), while the daily traffic volume varies from 773 GB to 1776 GB and unique IPs per day varied from 1M to 2M. To separate P2P signaling traffic and data traffic we set a threshold t on the flow sizes classifying all the flows of size $\leq t$ bytes as *signaling traffic* and rest as *data traffic*. The threshold was chosen according to P2P content size distribution observed in the Gnutella traces collected in earlier work [5] as 4K bytes. We used the graph library discussed in [1] to store all the IP addresses, prefixes, and AS numbers from the various P2P protocols along with the traffic volume (bytes) exchanged between them. Each entity (IP address, prefix, or AS) is represented by a node and each flow is represented by a direct edge between a pair of nodes in the graph. The bytes exchanged in the flow is represented as the weight of the corresponding edge. The graph software quickly separates out strongly connected components allowing us to zoom in on the heaviest connected component in terms of bytes exchanged between its members. Most of the bytes in all our traces are in the heaviest connected component. Smaller components likely represent the different communities of interest. The weighted indegree and outdegree information was used to separate the nodes in the heaviest connected component into *data* and *signaling* sections. The heavy entities (IP addresses, prefixes, AS numbers) in the heaviest connected component were classified using our classification mechanism.

II. Internet hierarchy classification

AS relationship can be classifed as *provider-customer* and *peer-peer* relationship [7]. In the former, the customer pays the provider for accessing the Internet. In the latter, the two AS peers find it mutually advantageous to exchange traffic between their respective customers. The relationship between a pair of ASes is inferred from the AS paths seen from multiple vantage points in the Internet. We assume that every AS adheres to the BGP export rules: (i) An AS can export its routes and routes of its customers to its providers and peers, but cannot export routes learned from other providers or peers. (ii) An AS can export its routes, routes of its customers, and routes learned from other providers and peers to its customer. Based on this, all the paths can be classified into an AS path either traverse zero or more customer-provider edges and then zero or more provider-customer edges; or it first traverses zero or more customer-provider edges, then traverse one peer-peer edge, and finally zero or more provider-customer edges. If AS Y can be reached from AS X via m provider-customer edges, we call AS Y a customer of AS X and AS X a provider of AS Y.

ASes are characterized via their relationships: (1) *ISP*: ASes owned by the ISP including its siblings, (2) *ISP-CUST*: ASes are direct BGP customers (m = 1) of the ISP and reachable via one provider-customer edge, (3) *ISP-PEER*: ASes are peers of the ISP, (4) *ISP-CUST-CUST**: ASes are indirect customers ($m >$ 1) of the ISP but not customers of any AS in *ISP-PEER* (i.e., they are only reachable from one or more ASes in *ISP-CUST*

Authors are with AT&T Labs–Research, 180 Park Ave., Florham Park, NJ 07932, USA. email:{bala,jiawang}@research.att.com.

Type	Signaling		Data	
Metrics	Traffic (%)	IP (%)	Traffic (%)	IP (%)
ATT	2.73	2.20	2.07	2.26
ATT-CUST	37.63	43.27	37.51	42.98
ATT-PEER	12.17	8.59	13.74	8.83
ATT-CUST-CUST*	0.00	0.00	0.00	0.00
ATT-PEER-CUST*	31.20	32.61	30.51	32.41
ATT-MH-CUST*	4.22	5.83	5.42	5.42
UNKNOWN	12.05	7.49	10.75	8.11

TABLE I
HEAVIEST CONNECTED COMPONENT IN GNUTELLA 09/01.

Type	Signaling		Data	
Direction	Src (%)	Dst (%)	Src (%)	Dst (%)
ATT	1.00	3.77	0.72	3.08
ATT-CUST	18.46	62.05	20.85	60.75
ATT-PEER	10.58	6.54	9.29	6.86
ATT-CUST-CUST*	0	0	0	0
ATT-PEER-CUST*	59.38	12.05	59.18	14.73
ATT-MH-CUST*	4.99	8.22	5.65	7.13

TABLE II
TRAFFIC DIRECTION IN GNUTELLA (09/01).

via multiple provider-customer edges and not reachable from any AS in *ISP-PEER*), (5) *ISP-PEER-CUST**: ASes are indirect customers of ASes in *ISP-PEER* but not customers of the ISP (i.e., only reachable from one or more ASes in *ISP-PEER* via multiple provider-customer edges and not from any AS in *ISP-CUST*), (6) *ISP-MH-CUST**: ASes are indirect customers of the ISP and of ASes in *ISP-PEER*. ISP-MH-CUST* are multihomed in the sense that they are reachable from one or more ASes in *ISP-PEER* via multiple provider-customer edges, and reachable from one or more ASes in *ISP-CUST* via multiple provider-customer edges. Unclassified ASes are labeled *UNKNOWN*.

We gathered list of ASes that belong to the AT&T IP backbone (*ATT*), *ATT-CUST* and *ATT-PEER* directly from the BGP configuration of access routers. Lists of ASes belonging to *ATT-CUST-CUST**, *ATT-PEER-CUST**, and *ATT-MH-CUST** are computed by traversing the provider-customer edges obtained in [7] from ASes belonging to *ATT-CUST* and *ATT-PEER*. The experiments were conducted at IP, prefix, and AS levels separately. The heaviest connected component had 99% of the IP addresses and over 99% of the traffic volume in all three protocols at all layers (DirectConnect at IP level ranged between 87% to 92%) during the three weeks. Over 90% of the total IP addresses contribute to the signaling traffic (only 0.4% of the total traffic volume). About 50% of total IP addresses contributed to data traffic (99.6% of total traffic volume). Similar results were observed for prefixes and ASes.

We partitioned signaling and data traffic flows of the heaviest connected component into 7 categories. At all levels, *ATT-CUST* and *ATT-PEER-CUST** are big contributors of both signaling and data traffic. Table I shows a summary of the classification of the heaviest connected component in Gnutella traffic at the IP level (DirectConnect and FastTrack are similar). *ATT-CUST* consists of 43% of the IP addresses that have signaling traffic and contributes 38% of the signaling traffic in Gnutella. Though *ATT* and *ATT-CUST* currently contribute ∼40% of the total traffic, the traffic that are exchanged within AT&T (i.e., both source and destination belong to either *ATT* or *ATT-CUST*) counts for 5% of the total Gnutella traffic.

Examining the heaviest connected component showed data traffic contributed by each of the individual IP addresses is highly skewed. In Gnutella, the top 0.1% and 8% of the IP addresses that have data traffic, contribute 30% and 60% of the data traffic, respectively. The signaling traffic contributed by individual IP addresses is much less skewed than that of the data traffic. The top 1% and 5% of IP addresses, that have signaling traffic, contribute 25% and 50% of the signaling traffic. The observation at the prefix and AS levels are similar. This indicates that popular contents are stored at very few places on the P2P networks, while the signaling traffic distribution depends on the application-level topology of the P2P network. To further classify how P2P traffic belonging to the heaviest connected component, we take the direction of the traffic flows into account. Table II shows the classification of the source and destination of Gnutella traffic. Most of the control and data traffic originate at *ATT-PEER-CUST** and are destined for *ATT-CUST*, which count for over 60% of the total traffic. The observation also holds on the heaviest connected components in DirectConnect and in FastTrack.

REFERENCES

[1] Charles D. Cranor, Emden Gansner, Balachander Krishnamurthy, and Oliver Spatscheck. Characterizing Large DNS Traces Using Graphs. In *Proceedings of ACM Sigcomm Internet Measurement Workshop*, November 2001.

[2] Direct Connect. http://www.neo-modus.com.

[3] Gnutella hosts. http://www.gnutellahosts.com.

[4] KaZaA. http://www.kazaa.com.

[5] B. Krishnamurthy, J. Wang, and Y. Xie. Early Measurements of a Cluster-based Architecture for P2P Systems. In *Proceedings of ACM Sigcomm Internet Measurement Workshop*, November 2001.

[6] Morpheus. http://www.musiccity.com.

[7] L. Subrammanian, S. Agarwal, J. Rexford, and R. H. Katz. Characterizing the Internet Hierarchy from Multiple Vantage Points. In *Proceedings of IEEE Infocom*, June 2002.