A Measurement Study of SOPCast

Florin Coras[1]*, Thomas Silverston[2], Jordi Domingo-Pascual[1], and Albert Cabellos-Aparicio[1]

[1] Universitat Politècnica de Catalunya, Barcelona, Spain
{fcoras,acabello,jordi.domingo}@ac.upc.edu
[2] The University of Tokyo, Tokyo, Japan
thomas@hongo.wide.ad.jp

Abstract. The sustained increase of network access speeds Internet users have experienced in recent years has lead to a proliferation of multimedia services that are now reportedly among the largest and the fastest growing global bandwidth consumers. It is the advent of these technologies and their flexibility in interacting with content that has enabled and encouraged users to start transitioning from traditional broadcasting systems to Internet based IPTV or similar streaming services. In order to transmit such multimedia content the industry has researched and developed scalable P2P live systems that are used daily by millions of users in the Internet. Given the relevance of such systems from the standpoint of both users and their bandwidth demands, it is important to understand their characteristics. In this paper we present and analyze a set of experimental datasets that reveal the properties of users and their traffic exchanges in the SOPCast live P2P streaming system.

Keywords: p2p live streaming; peer locality; SOPCast

1 Introduction

The Internet is gradually becoming the preferred infrastructure for delivering live content such as sports events or news to large user sets. According to recent reports, video streaming is among the largest and the fastest growing bandwidth consumers [15] and IPTV driven revenues are to rise [6] from less less than USD 6B in 2010 to USD 17B in 2016.

For such scenarios, where one-to-many content delivery is required, IP-multicast [11] is the most efficient solution in terms of bandwidth. However, inter-domain multicast has failed to be widely adopted by most ISPs. Reasons, among others, have to do with its management complexity and the lack of a clear commercial service. While the former leads to high operational expenditure the latter leads to situations when multicast implementation over links with unicast economical

* This work has been partially supported by the Spanish Ministry of Education under scholarship number AP2009-3790, research project TEC2011-29700-C02, Catalan Government under project 2009SGR-1140 and a Cisco Unsolicited Research Proposal Grant.

Fig. 1: Vantage Points

agreements leads to loss of revenue. Additionally, for end-customers, multicast does not provide any benefit since it has the same cost as unicast in terms of bandwidth. For ISPs, multicast is profitable when the bandwidth savings are higher than deployment and management costs, but typically, bandwidth is cheaper than management. In short, multicast never found its *sweet spot* [12] for large-scale deployment.

In light of this deployment predicament, both the industry and the academia have switched efforts to Application-Layer Multicast (ALM) where innovation was possible. Consequently, many academic [8, 10, 7, 18, 13, 9] and commercial [3, 2, 5, 4, 21] streaming products are now available. Most notably, PPLive [2] is reported to be used daily by millions of users [14].

The resulting popularity of the commercial ALM solutions has lead to an intense scrutiny [14, 20, 17] of their performance. In this paper we present and analyze a set of experimental datasets that show some of the characteristics of SOPCast, a P2P live content streaming application.

2 Dataset

In this section we describe the setup of the testbed we used to collect the dataset analyzed in this paper.

The aim, when building our testbed, was to create an infrastructure capable of performing a world-wide distributed passive capture of large P2P live content streaming overlays. The choice of the streaming content was driven by the subjectively perceived importance of the ongoing events at the time the experiment took place. As a result, all traces in our dataset consist of traffic pertaining to popular sports events. Reasons for our choice were threefold. First, users are generally interested in consuming such traffic live (as it gets produced). Secondly, interest for such events tends to be high world-wide. Finally, the usefulness of such content has been proven by previous [16] works that also aimed to characterize P2P live streaming overlays. In particular, we captured the content streamed by several SOPCast channels during the closing stages of the 2011 UEFA Champions League. Although we obtained datasets from multiple encounters, in this paper we focus solely on a 2011 UEFA Champions League semifinal match. In spite of the fact that interest for such football matches is highest in Europe, the teams involved are both highly appreciated worldwide and amount players spanning many nationalities. Additionally, interest was increased as this was the penultimate phase of the prestigious competition.

For the capturing process we used 2 vantage points in USA, 3 in Europe and 2 in Asia, spanning a total of 6 countries. A map depicting their geographic position is presented in Fig.1. In order to better capture the client interest for the streamed content, on each machine involved in the experiment we joined a number of SOPCast channels streaming the same event. Statistical properties of the traces are given in Table 1. All machines involved ran Ubuntu Linux and had a 100Mbps Ethernet connectivity to the Internet. The packet capturing was done with *tcpdump*. Because multiple SOPCast instances ran in parallel on each PC, filtering of packets per channel was done based on UDP port number. The capture duration was always higher than the time span of the match in order to observe transitory peer behavior.

Peer IPs were mapped to their autonomous systems (ASes) with the help of an origin AS database obtained from RouteViews [19].

For brevity, in what follows we shall be referring to the traces by associating a vantage point with one of the four captured channels (e.g., *ch1-barcelona*).

2.1 Observations

Some of the statistical properties of the traces captured are presented in Table 1. Among them, we have the count of peering IPs encountered in each trace, which may be used as an estimate for the number of connected end-hosts. However, because we did not identify hosts behind NAT boxes, this value should be held as a lower bound estimate. From the peer IPs, the number of Autonomous Systems (AS) exchanging traffic with our nodes is inferred. A similarity metric (described lower) is also computed for both IPs and ASes and the ratio of uploaded/downloaded traffic from each vantage point.

The balance between the upload and download ratios provides an insight into the peers altruistic nature. Content to be consumed is downloaded only once however depending on the peer's upload capacity it may be replicated several times. Channels with higher bitrate require an even larger effort from the peers. The second column in Table 1 shows the volume of content downloaded by each of our vantage points. The variations between the peers serving content in the same channel can be accounted for by congestion and differences in signalling traffic. In fact, an altruistic peer offering to replicate traffic receives in return a large volume of signalling traffic.

As in [16], we defined a similarity metric in order to evaluate the breadth of the peer and AS population that we measured. For a vantage point in a chan-

Table 1: Table with trace properties

	Download (GB)	**Number of IPs**	**Number of ASes**	**Up** (%)	**Down** (%)	**Similarity IP** (%)	**Similarity AS** (%)
ch1@850Kbps	9.63	64586	2839	89	11	12	68
california	1.45	19250	1469	76	24	93	98
cluj	1.50	32229	1980	90	10	92	96
ireland	0.99	13522	1294	66	34	92	98
barcelona	1.31	34320	1940	91	9	78	94
singapore	1.39	37164	2039	89	11	89	96
tokyo	1.18	37822	2028	95	5	91	96
virginia	1.33	21864	1745	88	12	92	96
ch2@345Kbps	2.48	19987	1539	82	18	24	87
california	0.42	9213	1060	74	26	92	98
cluj	0.47	11432	1212	91	9	85	93
ireland	0.19	6164	878	69	31	92	97
barcelona	0.38	7475	962	78	22	88	96
singapore	0.30	6025	709	73	27	82	96
tokyo	0.29	7014	865	77	23	87	96
virginia	0.43	7044	937	78	22	90	97
ch3@525Kbps	6.77	4815	820	93	7	26	90
california	0.93	3439	656	92	8	97	99
cluj	1.24	3515	659	94	6	98	98
ireland	0.87	3269	645	95	5	96	98
barcelona	1.06	3333	640	91	9	94	97
singapore	0.76	2715	568	79	21	98	98
tokyo	0.80	2853	590	69	31	98	99
virginia	1.11	3756	702	94	6	95	97
ch4@800Kbps	7.37	49814	2303	89	11	19	82
california	1.05	23307	1592	89	11	91	97
cluj	1.12	16734	1403	86	14	92	97
ireland	0.86	23983	1593	91	9	92	96
barcelona	1.36	23040	1614	90	10	92	96
singapore	0.79	16240	1185	85	15	88	98
tokyo	1.00	22926	1539	89	11	92	96
virginia	1.19	25254	1636	91	9	91	96

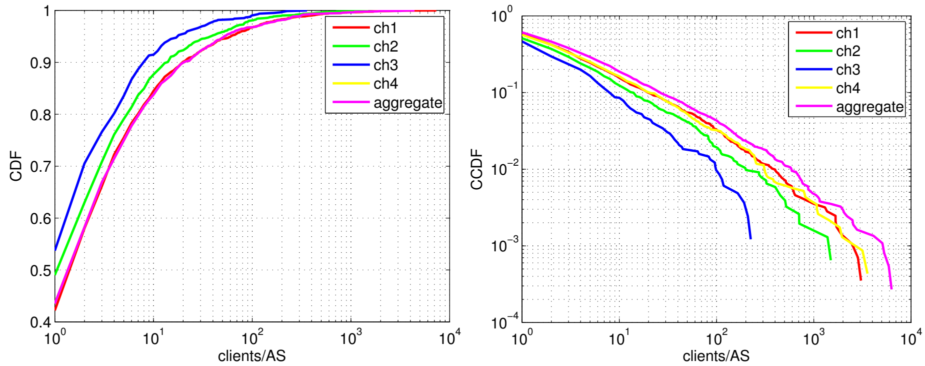

(a) CDF (b) CCDF

Fig. 2: Clients per AS distribution

nel, the metric was defined as the ratio of IPs/ASNs that overlap with those encountered in traces from other vantage points. For instance, in the case of the trace *ch1-california*, we observed a total of around 19k IPs and 1.4k ASNs which possess a similarity of 93% and respectively 98%. Overall, the similarity for IPs seldom drops under 85% and for ASNs never drops under 93%. The high IP similarity values indicate that in each channel our vantage points exchanged traffic with a large fraction of the peer population, leading to an accurate aggregate view of the whole overlay. Furthermore, the AS similarity values suggest that we have a precise estimate of the ASes exchanging traffic in all the channel overlays. This is also confirmed by the similarity of the curves describing the distribution of the clients in ASes (see Fig. 2).

If we aggregate the traces pertaining to each channel and perform the same analysis we observe that there is little overlap between the peer IPs. However, we do observe high values for the ASN overlap with the outlier being $ch1$ due to its much larger client population. Overall, we can infer that channels generally have non-overlapping clients (as expected) however, their clients pertain to autonomous systems that have a larger overlap.

Differences between channel client population sizes are explained by differences in popularity between the channels. From our dataset we deduced that the streaming bitrate and the language are two important factors to influence a channel's popularity. For instance, $ch1$ and $ch4$ are fairly popular due to their better streaming quality whereas $ch2$ and $ch3$ raise a lower interest. Moreover, the fact that Romanian was the language used in $ch3$ explains the lower number of clients.

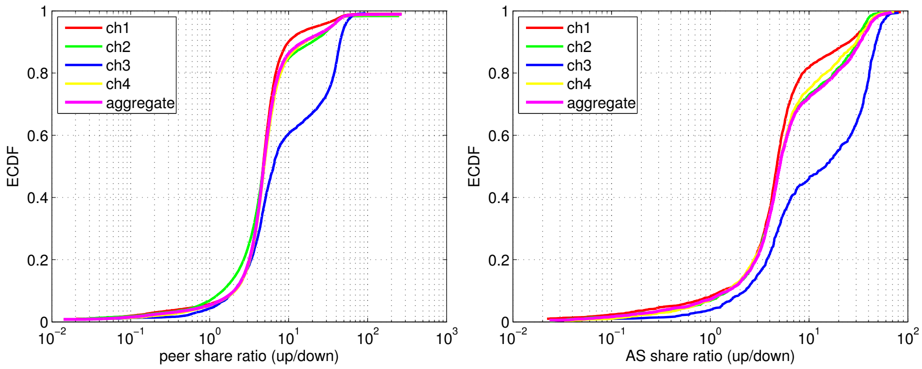

(a) Peer share ratio (b) AS share ratio

Fig. 3: Peer and AS share ratios

3 Distribution of clients in ASes

The distribution of clients in ASes is depicted in Fig. 2 as both cumulative distribution function (CDF) and complementary cumulative distribution function (CCDF). Additionally, we have added a curve depicting the distribution of clients in ASes for the aggregated clients sets of all channels. It can be seen that the plots have a similar shape and the differences between them are only due to the inter-channel client variations. This also holds for the aggregate distribution which is slightly different from the curves pertaining to $ch1$ and $ch4$ and which account for the largest client populations.

As we have seen in the previous section, channels tend to have non-overlapping client populations. Therefore the reasons behind the similarity of the curves have to do with a more subtle phenomenon probably related to user behavior and localized user interest. Figure 2a depicts the CCDFs of the distributions and from them we can deduce that clients roughly distribute in autonomous systems according to a power law. The reasoning being that power laws have as CCDF representation a straight line.

4 Collaboration between peers

Within P2P systems, it is the responsability of the peers to replicate content to other members. This, in fact, being the fundamental requisite for the scaling of P2P overlays. In this section we study the amount of traffic exchanged by our nodes with their peers. To evaluate their level of collaboration, we define and compute a for each overlay member a *sharing ratio*. Specifically, for each peer, the sharing ratio is computed by dividing its volume of uploaded traffic by the download one.

For all channels, Fig. 3 shows the CDF of the sharing ratios (upload/download) observed at our nodes at peer and AS level. The results are aggregated for channels and then for the whole experiment. A ratio larger than 1 means that one of our nodes has acted altruistically with a peer (it provided more data that it requested) whereas one lower means that the peer has provided content to one of our nodes. Of course, a value of 1 means that the peer is in traffic balance with one of our nodes.

Save for *ch*3 our nodes in all channels behave similarly and this holds also for the global (aggregate) traffic exchange. We speculate that the curve for *ch*3 has a different shape in the zone of larger share ratios due to the lower number of peers in the overlay and a higher need for nodes with high upload capacities.

In absolute terms, the plots show that our nodes acted as seeds for the overlay as they typically obtained the content from under 5% of their peers and uploaded content for 95%. Similarly, at AS level content is downloaded from around 7.5% of the peer ASes and replicated towards the rest of 92.5%. The fact that for all vantage vantage points their AS curve grows slower than the IP share ratio indicates that our nodes exchange large amounts of traffic with several nodes within the same AS. This is suggestive of an inefficient intra-domain replication of traffic, due to a lack of peer collaboration.

5 Peer and content locality

In order to understand the traffic requirements for ISPs involved in SOPCast overlays, we investigated the inter-AS traffic exchanges. Figure 4 depicts the distribution of both upload and download traffic for vantage points in *ch*1 and *ch*2.

Although the number of ASes with contacted peers is above $1k$ for *ch*1 and above 700 in the case of *ch*2, both for upload and download, the greater part of traffic is exchanged with a restricted AS subset. Nevertheless, its size is not small. In the case of *ch*1, on average, more than 50 AS were contacted to download 90% of the streamed content. Similarly, the vantage points receiving *ch*2, to obtain 90% of the content, had to connect to more than 20 ASes. For both channels, the vantage points replicate the content to a slightly larger set of ASes than the one from which they had obtained it. However, it is interesting to note here that two different kinds of seeding behaviours were observed. They are easily distinguishable in Fig. 4c for *ch*1 where we see that a first group of the vantage points replicate content towards a limited number of ASes, actually comparable with the one from which they received their content. The second group replicates content to a larger set of ASes however still smaller than 100 if we consider the destinations of 90% of the traffic. This behavior is not that easily distinguishable in the case of *ch*2 except as the number of ASes increases.

Overall, we could conclude that the number of ASes towards which a large part (90%) of the content gets replicated is only slightly larger than the one from which content is obtained. Then, if we consider the large replication factors of each vantage point seen in Table 1 we can infer that our nodes generally replicate

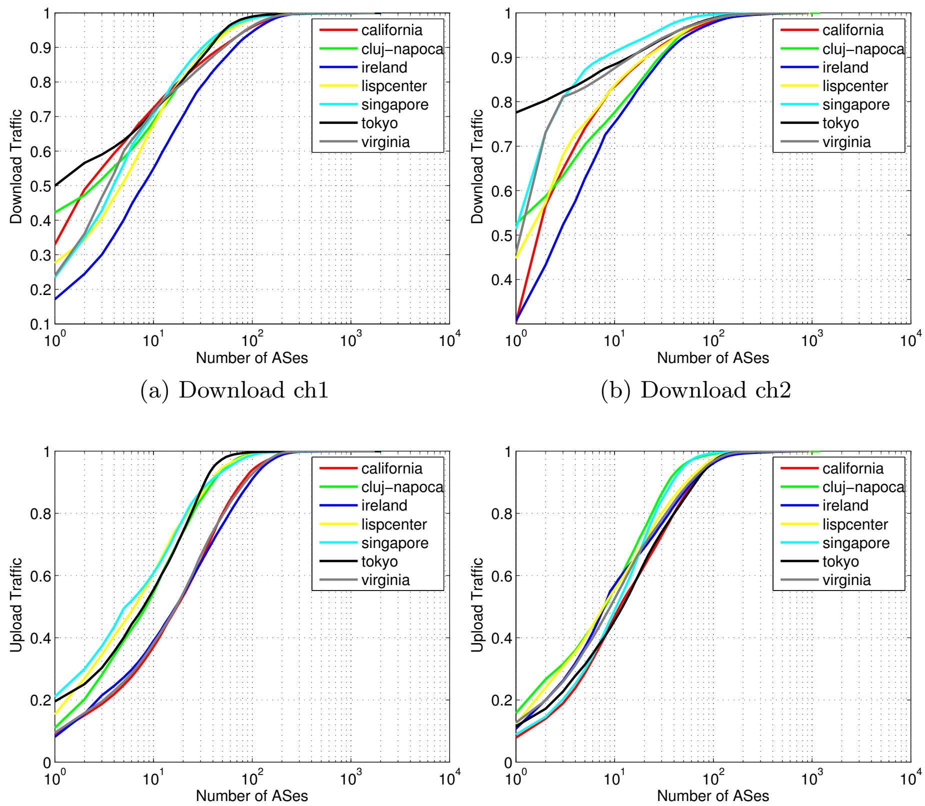

(a) Download ch1 (b) Download ch2

(c) Upload ch1 (d) Upload ch2

Fig. 4: Inter-AS traffic

content to a large number of peers in each AS. A rather inefficient use of inter-AS bandwidth that is probably caused by the inability of nodes to identify closeby (same AS) peers.

5.1 Geographic location of peers

To better understand the criteria (if any) based on which traffic is replicated between peers and ASes, we studied the geographic distribution of participating hosts and the volume of traffic exchanged at country level. ASes were mapped to their corresponding country with the help of the cymru [1] whois service. Although an AS can span many countries, for our intended purposes this mapping was accurate enough.

Figures 5 and 6 show our results for $ch1$ and $ch2$. We did not distinguish between peers used for uploading and those for downloading as these two groups tend to have a high percentage of overlap. Due to readability constraints, the

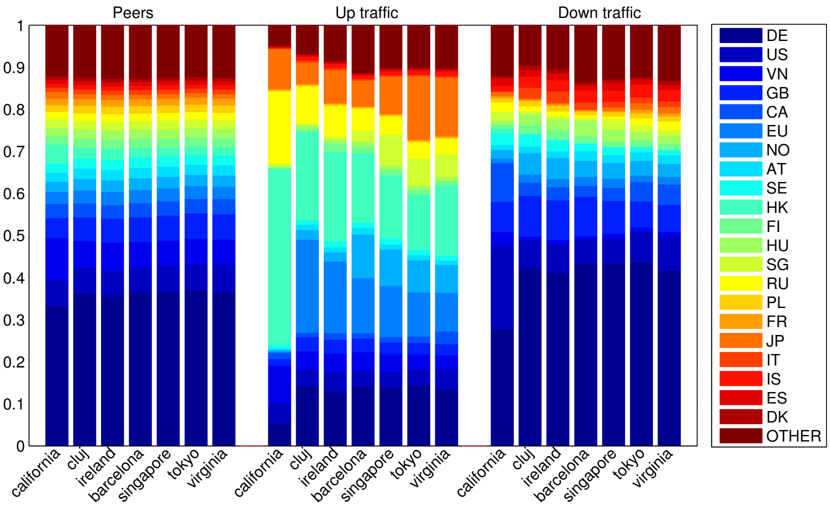

Fig. 5: Geographic location of peers in ch1

stacked histograms present just the first 20 countries, when ranked by number of clients. The legend entry *OTHER* stands as an aggregate for the rest of the countries that could not be shown.

The country with the largest population in both channels is Germany. This is somewhat unexpected as both football teams involved were Spanish. In fact, Spanish peers make up just a small part of the *ch*1 overlay and are not present in the top 20 of *ch*2. Additionally, at no vantage point did we observe any sort of biasing of the peer population towards the country where the measurement was performed. In effect, the *Peer* bars, in both figures, indicate that all vantage points observe almost the same distribution of peers in countries. This corroborates the conclusion that our nodes had an accurate view of the channel population.

Surprisingly, the *Download traffic* bar plots are also very similar although, with more important differences than the *Peer* ones. This suggests that the vantage points are completely unbiased (from geographic perspective) in their choice of upstream peers. For *ch*1, traffic is mainly obtained from a handful of countries in Northern Europe and America: Germany, USA, GB, Canada and Iceland. Similarly, a large part of the content for *ch*2 is obtained from a limited set of countries: Germany, GB, Hungary, Egypt, France. As it may be seen, again the majority are placed in Europe in spite of the fact that the channel is of Chinese origin.

Similarities between the upload patterns of the vantage points are not as obvious as the one for download but they are still discernible. For *ch*1, save for the fact that the vantage point in California uploaded little content to EU (ASes registered with the European Union country code) and Norway customers, the typical node behavior is to upload mainly to clients in Germany, Hong Kong, EU,

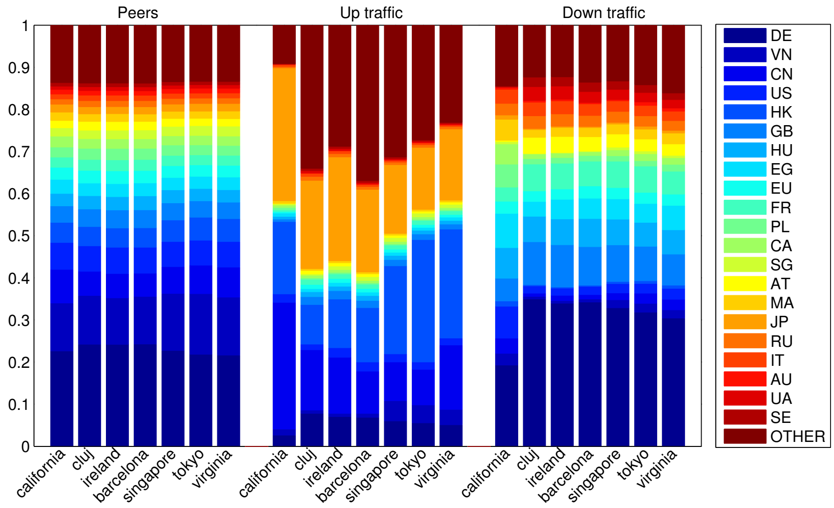

Fig. 6: Geographic location of peers in ch2

Norway and Italy. The fact that traffic is resent to Germany after it was mainly received from there, speaks again against the efficiency of the routing which does not seem to be aware of the peer locality. The bouncing between the countries is even more detrimental if one considers that the content is delay sensitive. Same observations hold for $ch2$ just that the destination countries are different. The predilect destinations for the uploaded traffic were China, Hong Kong and Japan. A considerable amount of content has been uploaded to destinations with few peers and from which a low volume of content has been downloaded.

In summary, the analysis has shown that the routing in the overlay is not aware of peer locality. As a result, nodes interact in a similar fashion, in this sense, our nodes observe the same distribution of peers in countries, obtain their content from the same sources and roughly upload to the same destinations.

6 Conclusions

In this paper, we presented our large-scale measurements of SOPCast P2P live streaming overlays.

Throughout the study we quantified several relevant parameters for live streaming overlays. From traffic volume perspective we characterize both global and peer level exchanges and provide information about upload/download ratios and respectively peer share ratios. Furthermore, we also analyze and provide information about traffic locality at AS level. Finally, we study the geographic location of peers and the traffic exchange at country level.

Our results indicate that SOPCast does not consider peer geographic location and generally exchanges traffic arbitrarily between a large set of ASes. As a result the same content may be exchanged multiple times between two ASes and intra-domain traffic exchanges are not prioritized. We did not observe a correlation between the streamed traffic and the geographical location of clients and the application seemed to be used world-wide indiscriminately of continent.

References

1. Cymru, http://www.team-cymru.org
2. PPLive, http://www.pplive.com/
3. Sopcast P2P Internet TV, http://www.sopcast.com
4. TVAnts, http://www.tvants.com/
5. UUSee, http://www.uusee.com/
6. Digital TV World Revenue Forecasts (Jun 2011), http://www.digitaltvresearch.com/products/product?id=21
7. Banerjee, S., Kommareddy, C., Kar, K., Bhattacharjee, B., Khuller, S.: Construction of an efficient overlay multicast infrastructure for real-time applications. In: Proc. IEEE INFOCOM. pp. 1521–1531 (2003)
8. Banerjee, S., Bhattacharjee, B., Kommareddy, C.: Scalable application layer multicast. In: SIGCOMM, Proceedings ACM. pp. 205–217 (2002)
9. Castro, M., Druschel, P., Kermarrec, A.M., Rowstron, A.: Scribe: a large-scale and decentralized application-level multicast infrastructure. Selected Areas in Communications, IEEE Journal on 20(8), 1489 – 1499 (oct 2002)
10. hua Chu, Y., Rao, S., Seshan, S., Zhang, H.: A case for end system multicast. Selected Areas in Communications, IEEE Journal on 20(8), 1456 – 1471 (oct 2002)
11. Deering, S.: Host extensions for IP multicasting. RFC 1112 (Standard) (Aug 1989), http://www.ietf.org/rfc/rfc1112.txt, updated by RFC 2236
12. Diot, C., Levine, B.N., Lyles, B., Kassem, H., Balensiefen, D.: Deployment Issues for the IP Multicast Service and Architecture. IEEE Network 14(1), 78–88 (Jan 2000)
13. Francis, P.: Yoid: Extending the internet multicast architecture. Unpublished paper, available at http://www.aciri.org/yoid/docs/index.html (2000)
14. Hei, X., Lian, C., Lian, J., Liu, Y., Ross, K.W.: A Measurement Study of a Large-Scale P2P IPTV System. TOM 9(8), 1672–1687 (Dec 2007)
15. Labovitz, C., Iekel-Johnson, S., McPherson, D., Oberheide, J., Jahanian, F.: Internet inter-domain traffic. ACM SIGCOMM Computer Communication Review 40(4), 75–86 (2010)
16. Silverston, T., Jakab, L., Cabellos-Aparicio, A., Fourmaux, O., Salamatian, K., Cho, K.: Large-scale measurement experiments of p2p-tv systems insights on fairness and locality. Signal Processing: Image Communication (2011)
17. Silverston, T., Fourmaux, O.: Measuring P2P IPTV Systems. NOSSDAV (2007)
18. Tran, D., Hua, K., Do, T.: A peer-to-peer architecture for media streaming. Selected Areas in Communications, IEEE Journal on 22(1), 121 – 133 (jan 2004)
19. University of Oregon: RouteViews Project, http://www.routeviews.org
20. Vu, L., Gupta, I., Liang, J., Nahrstedt, K.: Measurement and modeling of a large-scale overlay for multimedia streaming. QSHINE (1), 1 (2007)
21. Zhang, X., Liu, J., Li, B., Yum, Y.S.: Coolstreaming/donet: a data-driven overlay network for peer-to-peer live media streaming. In: INFOCOM, Proceedings IEEE. vol. 3, pp. 2102 – 2111 (march 2005)