1

Characterizing Dynamic Properties of the SopCast Overlay Network

Kênia Carolina Gonçalves[1], Alex Borges Vieira[2], Jussara M. Almeida[1],
Ana Paula C. da Silva[2], Humberto Marques-Neto[3], Srgio V. A. Campos[1]

[1]Computer Science Dept.
Universidade Federal de Minas Gerais, Brazil
{keniacarolina,jussara,scampos}@dcc.ufmg.br

[2]Computer Science Dept.
Universidade Federal de Juiz de Fora, Brazil
{alex.borges,anapaula.silva}@ufjf.edu.br

[3]Computer Science Dept.
PUC Minas, Brazil
{humberto}@pucminas.br

Abstract

Peer-to-Peer live video streaming systems are becoming increasingly popular. Nevertheless, in spite of various studies of client behavior aspects and system optimizations, the current knowledge about the dynamic properties of the system, particularly how the P2P overlay network changes over time during a live transmission, is still superficial. In this paper, we provide a characterization of the dynamic properties of a popular P2P live streaming media application, namely SopCast. We use complex network metrics to analyze how the structure of the network evolves over time from the perspective of individual nodes (local view) and of the whole network (global view). We find that SopCast peers may be clustered into three profiles based on their centrality properties in the network. Moreover, inspire of peers changing their partners over time, they tend to remain with the same centrality profile. Also, the global network structure tends to remain roughly stable over time, except for a decaying clustering coefficient. Our findings can be used to generate more realistic synthetic P2P workloads and to drive future system designs and simulations.

1. Introduction

Video applications popularity increases each day. Recent studies report that the number of users of this type of application may reach 83 million in 2012 [9]. In particular, Peer-to-Peer (P2P) live video streaming applications, such as SopCast[1], PPLive[2] and PPStream[3], are also currently very popular.

[1]www.sopcast.org
[2]www.pptv.com
[3]www.ppstream.com

Previous studies analyzed different aspects of P2P live streaming applications such as peer behavior[4, 10] and overlay properties[12, 5]. Most of these studies focus on a aggregated view of the system during a transmission. However, P2P overlay network evolves over time[16]. Such dynamic properties may impact system maintenance, resource allocation, and, ultimately, the quality of the live video delivered to clients. To our knowledge, previous analysis of such properties focused only on how the number of partners of a peer change over time[13, 4].

We provide a broader characterization P2P overlay network dynamics during live transmissions. Our characterization is performed from two perspectives: (a) the local view from individual nodes, and (b) the global view of the network as a whole. We characterize the dynamics of individual nodes within the network structure using complex network metrics. We use centrality metrics, namely degree, betweenness and closeness, to capture how the node is localized within the network. We analyze the dynamics of the network structure as a whole characterizing how the diameter, clustering coefficient, maximum degree and average shortest path changes over time.

Our characterization relies in SopCast data, a widely popular application (Google Trends). We ran hundreds of SopCast clients on PlanetLab[4] [8] to monitor a popular Chinese channel. Our results show that peers can be grouped into three profiles with distinct centrality properties: *High Centrality* (HC), *Intermediate Centrality* (IC) and *Low Centrality* (LC). HC peers are more central nodes in the network structure with large values of degree, betweenness and closeness. LC peers, in turn, are more peripheral, whereas IC peers have more intermediate values of the centrality metrics.

We also show that peers have high probability of remaining with the same centrality profile over time, that is, a non negligible portion changes their profile. Analyzing the partnerships established by a peer during a transmission, we

[4]www.planet-lab.org

note that even though a peer tends to keep the same partners during short periods of time (i.e., on average 70% of the partners are maintained during consecutive snapshots), the fraction of partners in common tends to decrease for longer time window. Our results also show that network structure tends to remain roughly stable in terms of diameter, average shortest path and maximum degree, whereas the clustering coefficient tends to decrease over time, possibly due to the changes in partnerships. **We believe these findings can be exploited to build more realistic P2P workloads and simulation environments and to drive future designs and evaluations protocols.We believe if the protocols could detect peer centrality profile in time, the broadcast server can suggest better initial peers to new clients and knowing how the High Centrality (HC) peers are could help to broadcast the content.**

The remainder of this paper is organized as follows. Section 2 discusses related work, whereas SopCast data collection methodology is described in Section 3. Section 4 introduces the complex network metrics we use to characterize peer and network dynamics. Our characterization results are discussed in Section 5. Section 6 offers conclusions and future work.

2 Related Work

Previous studies on P2P live streaming focus on understanding client behavior[4, 10] and proposing new algorithms to build and maintain the P2P overlay[5, 3]. For instance, Hei *et. al.*[4] shows that PPLive reaches a good performance using the Internet infrastructure. Further, they characterized the user behavior and the information that they exchanged. In [5], authors present a methodology for analyzing the user behavior. This analysis is based on the system performance and the quality perceived by users, without considering the partnerships among peers on the overlay structure.

Silverston *et. al.*[10, 11] analyzed traffic patterns from four P2P live streaming applications. They characterized metrics such as upload/download rates and average packet sizes. Other studies have analyzed the number of partnerships established by a peer and peer lifetime in the system[4, 13]. Bermudez *et. al.*[1] analyzed the SopCast traffic monitored in IPSs focused in bandwidth, peer lifetime(arrive and depart) and peer localization in the world. However, unlike our work, these previous studies did not analyze how the topological structure of the overlay network evolves.

Tang *et. al.*[13] studied the SopCast overlay. In particular, authors characterized control and data packets exchanged between partners. They show that out-degree is correlated with average upload volume. Indeed, Wu *et. al.*[16] analyzed network metrics over time as degree, clus-

tering coefficient, reciprocity and average shortest path. They found that the clustering coefficient has a variation while the average shortest path not. They do not analyze from two viewpoints and other metrics: global(diameter and maximum degree) and local(betweenness and closeness), and how the individual peers evolve over time in their centrality.

Stutzbach *et. al.*[12] proposed a methodology to analyze the evolution of the overlay network in a P2P file sharing application. Our work is fundamentally different as the two types of applications have inherently different characteristics: live video applications have much stricter delay constraints which may impact in the network dynamics.

The temporal evolution of networks, particularly the "small-world" phenomenon, was previously studied by Watts and Strogatz [15]. They analyzed the dynamics of several networks, using metrics such as diameter, degree and clustering coefficient. Easley and Kleinberg's [2] also presents analysis of the growth and reduction of various types of networks.

3 Data Collection Methodology

The data characterized in this work is describe as following, we first briefly describe the SopCast application (Section 3.1) and then present our data collection methodology (Section 3.2).

3.1 The SopCast Application

SopCast is one of the most popular P2P live streaming applications. We collected data from CCTV-1, one of the most popular open channels in China, which, in turn, is one of the countries with the largest number of SopCast clients.

Each SopCast channel has its own overlay. Each overlay has three components: a live streaming server, a bootstrap server and a set of clients(peers). Clients exchange data among them to watch the live streaming video. The server is a special client which encodes the media. The bootstrap server maintains a centralized record of all clients. When a new client connects to a particular channel, it contacts the bootstrap, which returns a list of clients already receiving the live content. The joining client may establish partnerships with a subset of these clients.

Figure 1 shows two consecutive snapshots of a hypothetical overlay network formed for the transmission of the CCTV-1 SopCast channel. As illustrated, Client 2 has partnerships established with Clients 1 and 3 in snapshot i. In the next snapshot, the partnership with Client 1 is undone, and Client 2 starts exchanging data with Client 4.
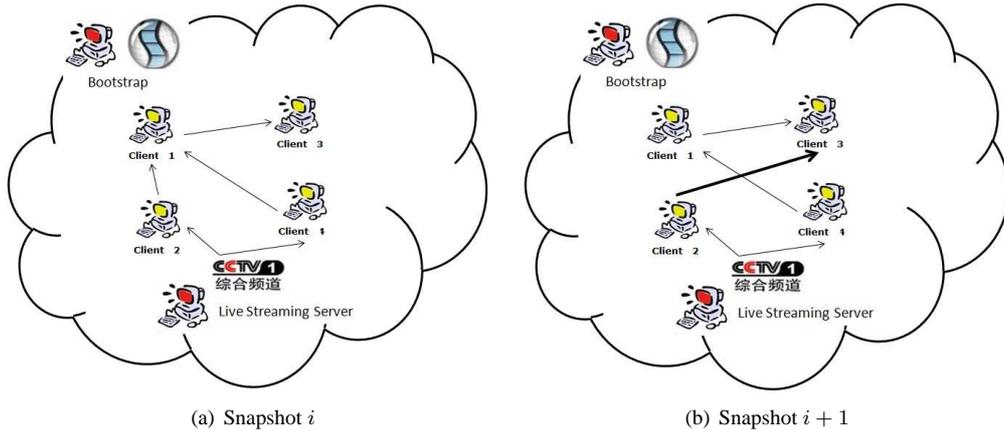
(a) Snapshot i (b) Snapshot $i+1$

Figure 1. Hypothetical example illustrating the dynamic nature of the P2P overlay network.

3.2 Data Crawling

We ran a series of 7 experiments with SopCast using between 200 and 465 PlanetLab computers between October and November 2010. The difference in the numbers of Sop-Cast clients discovered by the crawlers in the 7 experiments is under 1%. Each PlanetLab computer collects data from CCTV-1 chinese channel at local prime time (8pm), and stored all the data exchanged with its partners during a 60-minute transmission. CCTV-1 transmits live content at 600 Kbps.

The data crawling methodology consists of two main steps. During the first step, the PlanetLab computers were configured without storage or bandwidth constraints. Wire-shark[5] (*tcpump*) was used to capture the network traffic observed during the monitoring period. We configured Wire-shark to capture only SopCast traffic. We stored the data collected during each experiment in a log file containing the date and time of each packet sent/received. We used the Network Time Protocol (NTP)[6] to synchronize the time of all PlanetLab computers as in [7] for ensuring that time differences between computers could be negligible.

The second step consists of the data collection. Crawlers join the SopCast channel and capture all data packets exchanged with their partners. Crawlers joining times are normally distributed during an initial startup period of 10 minutes. After this, we monitor the channel transmission for 40 more minutes.

Once the monitoring period is finished, we merge all log files. Based on the time information as well as the source and destination IP addresses of each packet, we reconstruct the SopCast overlay network dynamics taking 118 consecutive snapshots of the network, each one built with data collected during a window of 60 seconds. We take the snap-

shots by considering a sliding window that moves along the transmission 20 seconds at a time.

Our collection methodology captures a partial view of the CCTV-1 network. We used a large number of Planet-Lab crawlers to cover the largest number of SopCast clients. Our experiments evidence that we were able to cover a very large fraction of SopCast: considering the client population discovered by all 465 crawlers during each snapshot of one monitored transmission, using 200 crawlers is enough to cover more than 98% of those clients. In other words, using more than 200 crawlers contribute only marginally to discover more clients, whereas a significant fraction of the clients may be missed if fewer crawlers are used.

4 Characterization Metrics

We model each SopCast network snapshot as a directed graph $G = (V, E)$, where V is the set of vertices and E is the set of edges. A directed edge $(v_i, v_j) \in E$ indicates a partnership between the clients represented by vertices v_i and v_j.

We use complex network metrics and analyze how these metrics change over time. We perform a characterization from two perspectives: the local view of individual peers and the global view of the whole network.

4.1 Peer Metrics

We characterize the dynamics of peers within the structure of the overlay network using the following node centrality metrics:

Degree: The degree $d(v_i)$ of v_i is a simple metric which represents the total number of partners of a given peer.

Betweenness: The betweenness of v_i is the fraction of all shortest paths, computed using breadth-first search, connecting pairs of vertices that pass through v_i. In other

---

[5]www.wireshark.org
[6]www.ntp.org

Table 1. *Peer Centrality Profiles for Each Experiment*

			% of Peers	Degree		Betweenness		Closeness	
				Average	CV	Average	CV	Average	CV
Experiment	1	HC	4.76%	282.83	0.17	3312.52	0.45	0.005	1.25
		IC	32.69%	257.99	0.20	1212.52	0.30	0.008	1.20
		LC	62.53%	86.94	0.82	129.45	1.60	0.005	1.16
	2	HC	6.98%	334.89	0.21	6585.63	0.33	0.005	0.61
		IH	45.85%	224.13	0.17	2258.33	0.32	0.006	0.64
		LC	47.16%	36.10	1.22	58.73	3.22	0.003	0.88
	3	HC	3.81%	361.95	0.18	8604.33	0.30	0.003	0.52
		IC	17.55%	240.43	0.21	2556.77	0.36	0.004	1.06
		LC	78.62%	56.83	1.14	135.47	2.30	0.003	0.90
	4	HC	2.92%	298.33	0.24	11430.93	0.43	0.004	0.56
		IC	14.61%	230.04	0.20	3443.58	0.38	0.007	0.82
		LC	82.46%	56.83	1.14	157.44	2.64	0.004	0.97
	5	HC	6.76%	322.39	0.24	6898.68	0.56	0.004	0.79
		IC	40.97%	230.42	0.19	2134.54	0.36	0.006	0.89
		LC	52.25%	56.83	1.14	120.73	2.24	0.003	1.04
	6	HC	2.69%	241.23	0.20	11865.11	0.63	0.004	0.12
		IC	8.99%	252.24	0.24	3159.50	0.42	0.007	0.75
		LC	88.31%	86.20	1.04	282.39	1.61	0.007	1.08
	7	HC	19.79%	298.42	0.16	1751.81	0.31	0.006	1.42
		IC	30.03%	261.86	0.16	833.93	0.22	0.007	1.37
		LC	50.17%	123.47	0.68	159.52	1.13	0.005	1.61

words, let σ_{v_j,v_k} represent the number of shortest paths between v_j and v_k, and $\sigma_{v_j,v_k}(v_i)$ the number of those paths that pass through v_i. The betweeness of v_i is defined as:

$$Betweenness(v_i) = \sum_{v_j \neq v_i \neq v_k \in V} \frac{\sigma_{v_j,v_k}(v_i)}{\sigma_{v_j,v_k}}$$

Closeness: The closeness of v_i captures how close it is from all other vertices reachable from it in the network. Given $l(v_i, v_j)$, the shortest path between v_i and any other vertex v_j, reachable from v_i, the closeness of v_i is defined as:

$$Closeness(v_i) = \frac{|V| - 1}{\sum_{v_j \neq v_i, v_j \in V} l(v_i, v_j)}$$

4.2 Network Metrics

The structure of the overlay network as a whole was evaluated using the following metrics:

Diameter: Defined as the longest shortest paths between any 2 nodes. It provides an idea of the dispersion of the graph and thus, can be used to estimate the latency experienced by clients in the P2P network. The diameter of a graph G is defined as:

$$Diameter(G) = max_{\forall v_i, v_j \in V} m(v_i, v_j)$$

Average Shortest Path: This metric is computed by taking the average across all values of $l(v_i, v_j)$ for all pairs of

vertices $v_i, v_j \in V$.

Clustering Coefficient: The clustering coefficient of a graph G is a measure of degree to which the vertices of G tend to cluster together. The clustering coefficient of v_i, $CC(v_i)$ is defined as the proportion of links between vertices that are neighbors of v_i divided by the number of links that could possibly exist between them. Let $nEdges(v_i)$ be the number of edges connecting neighbors of v_i in G, $CC(v_i)$ is defined as:

$$CC(v_i) = \frac{nEdges(v_i)}{d(v_i)(d(v_i) - 1)}$$

Finally, $ClusteringCoefficient(G) = \frac{\sum_{v \in V} CC(v)}{|V|}$

Maximum Degree: The maximum degree of a graph G is given by the largest degree of any vertex $v_i \in V$:

$$MaximumDegree(G) = max_{\forall v_i \in V} d(v_i)$$

5 Characterization Results

In this section we present the most relevant results from our characterization. We discuss the characterization of dynamic properties of individual peers (Section 5.1), and of the complete network (Section 5.2).

5.1 Individual Peers

We first identify a number of *centrality profiles* and then analyzed how a single node changes its profile during a transmission. In order to identify centrality profiles during one experiment, we represented each peer/node v_i during a snapshot t of the network by a vector $V(v_i, t) = (d, b, c)$, where d, b and c are the degree, betweenness and closeness of v_i in t. We then ran the *k-means* [14] clustering algorithm to group vectors into a number of clusters, each one representing a different centrality profile.

We selected the number of clusters k according to the ratio of the Coefficient of Variation (CV) of the distances between vectors within the same cluster to the coefficient of variation of the distances between vectors belonging to different clusters. We used the Euclidean distance computed in the 3-dimensional space as a measure of the distance between 2 vectors. As described in [6], the number of clusters should be the smallest once the aforementioned ratio stabilizes.
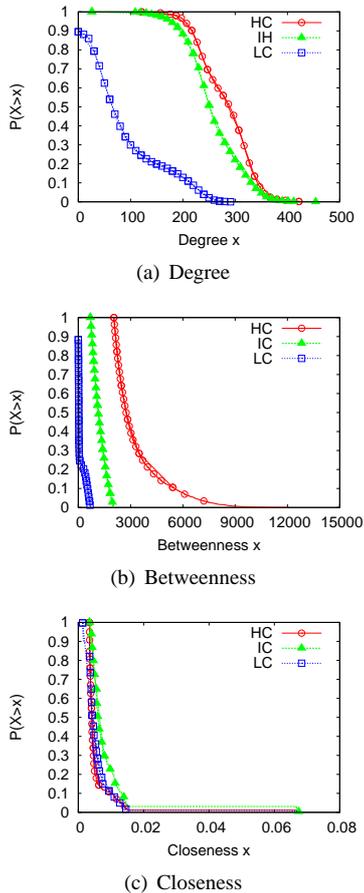

(a) Degree

(b) Betweenness

(c) Closeness

Figure 2. Distributions of Centrality Measures in each Profile

We found 3 peer centrality profiles in all 7 experiments. We refer to them as *High Centrality* (HC), *Intermediate Centrality* (IC), and *Low Centrality* (LC). Table 1 summarizes the characteristics of each profile for each experiment, presenting average values and CVs for the centrality metrics. The exact values that define each cluster centroid vary depending on the experiment.

HC peers have higher degree and betweenness than the other profiles. The *HC* cluster correspond to peers that are more central in the network. They represent a small fraction of the peers, varying from 2.5% to 7% of all peers (occasionally 20%). *LC* peers have the lowest degree and betweenness values, corresponding to peers that are located in the periphery of the network. *LC* is the most common profile, containing from 47% to 82% of all peers. *IC* peers have intermediate values. Closeness does not present a considerable variation, thus it cannot clearly distinguish profiles.

Figure 2 shows the Complementary Cumulative Distribution Function (CCDF) of each centrality metric for each profile (experiment 1 data). These curves are representative for all other experiments. Figure 2-a shows that *HC* peers tend to have much larger degrees: only 4.76% of the peers have degrees below 200. In contrast, around 62.53% of the *LC* peers have degrees less than 2. The distributions of betweenness also show clear differences between the profiles: *HC* peers have significant larger betweeness values than *IC* and *LC* peers.

(a) Experiment 1
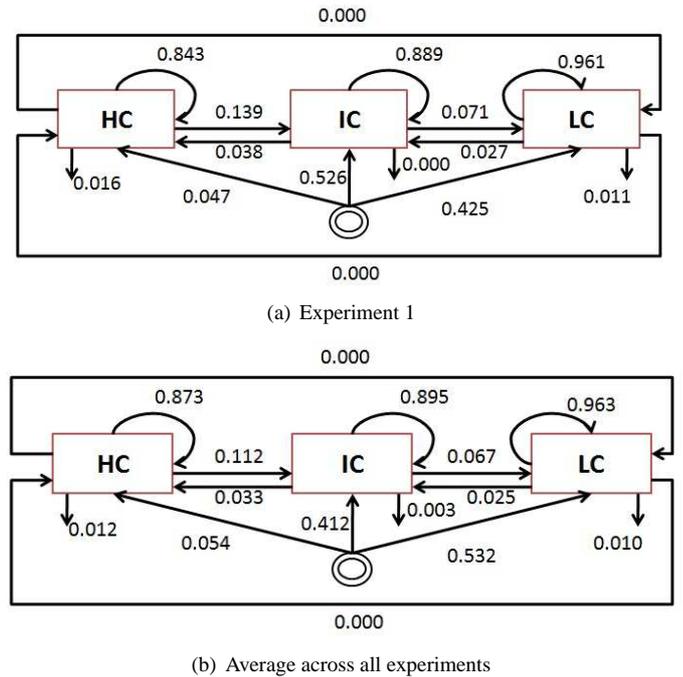

(b) Average across all experiments

Figure 3. Changes in a Peer's Centrality Profile During a Transmission

During one transmission peers change their partnerships and the network structure modifies. At a given time, a certain peer may be located very centrally in the network, having large centrality measures. Sometime later, after changes in the partnerships established in the network, the same peer may have moved towards the periphery of the network, exhibiting lower centrality measures.

To characterize these dynamic changes in centrality profile, we apply a modeling technique called *Customer Behavior Model Graph* (CBMG) [6]. A CBMG is a state-transition diagram commonly used to model a customer's dynamic behavior. We build one CBMG for each experiment. Each state of a CBMG represents a centrality profile, and the transitions between 2 profiles a and b are labeled with the probability of a peer changing its profile from a to b in 2 consecutive network snapshots of the given experiment. Figure 3-a shows the CBMG built for Experiment 1 whereas an aggregated CBMG, built for all 7 experiments, is shown in Figure 3-b.
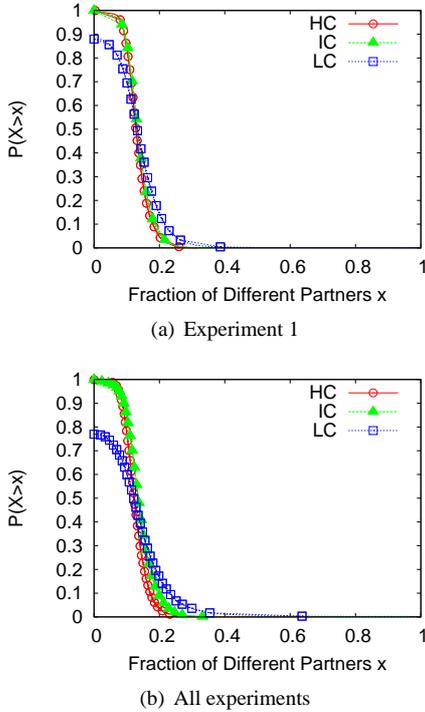

(a) Experiment 1

(b) All experiments

Figure 4. Distributions of the Fraction of Different Partners of a Peer in Consecutive Network Snapshots

Both CBMG's are similar, indicating similar results for all 7 experiments. Peers tend to remain in the same profile in consecutive snapshots. Figure 3-b shows that the probability that a peer that has low (high) centrality remains with the same profile is 0.963 (0.873). Moreover, the probability

of transitions between extreme profiles is practically zero. Figure 3 also shows non-negligible probabilities of peers changing their profiles, increasing/reducing their centralities. We find that the probability of a peer reducing its centrality always tends to be larger than the opposite. The transition between HC and IC occurs with probability 0.112 in the average CBMG, whereas the reverse transition has a probability of only 0.033.

A peer may change its partners without affecting its centrality profile. We analyze this issue by quantifying, for a given peer i, the percentage of different partners in different snapshots of the same transmission. Let p_i^t be the list of partners of peer i in snapshot t. We quantify the difference between $p_i^{t_1}$ and $p_i^{t_2}$ by taking the ratio of the intersection of $p_i^{t_1}$ and $p_i^{t_2}$ to the union of $p_i^{t_1}$ and $p_i^{t_2}$. In other words, we measure $\frac{p_i^{t_1} \cap p_i^{t_2}}{p_i^{t_1} \cup p_i^{t_2}}$.
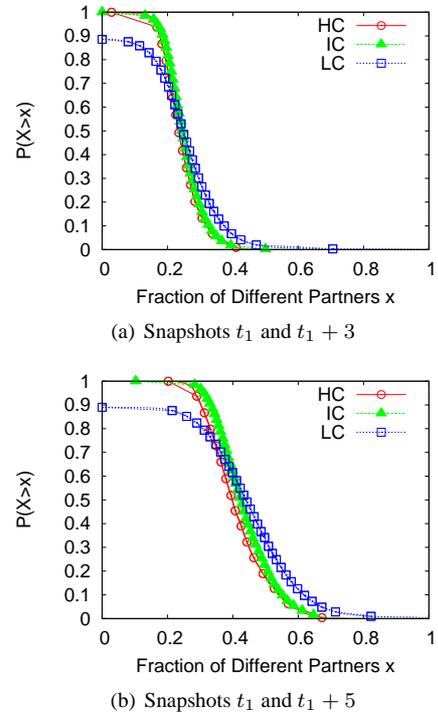

(a) Snapshots t_1 and $t_1 + 3$

(b) Snapshots t_1 and $t_1 + 5$

Figure 5. Distributions of the Fraction of Different Partners of a Peer for Non-Consecutive Snapshots

Figure 4 shows the CCDF of the fraction of different partners of a peer for consecutive snapshots for each centrality profile. Results shown for experiment 1 (Figure 4-a) and for all experiments (Figure 4-b) are very similar. The fraction of different partners in consecutive snapshots tends the same for all profiles. For HC and IC peers, this fraction is less than 20%. Figure 5 shows results computed for

		Diameter		Average Shortest Path		Clustering Coefficient		Maximum Degree	
		Average	CV	Average	CV	Average	CV	Average	CV
Experiment	1	4.11	0.07	1.98	0.03	0.24	0.34	361.47	0.08
	2	4.13	0.09	2.08	0.03	0.15	0.28	431.51	0.11
	3	4.33	0.13	2.09	0.02	0.15	0.46	429.71	0.06
	4	4.27	0.10	2.16	0.03	0.12	0.56	389.88	0.09
	5	4.22	0.09	2.14	0.02	0.17	0.21	398.68	0.12
	6	4.23	0.10	2.10	0.07	0.18	0.64	339.83	0.11
	7	4.05	0.05	2.00	0.02	0.32	0.11	377.00	0.05

snapshots separated by 3 and 5 snapshots, for experiment 1. That is, for $t_2 = t_1 + 3$ (Figure 5-a) and for $t_2 = t_1 + 5$ (Figure 5-b). These results indicate greater changes in peer partnership. For instance, around 20% of the HC peers have at least 50% different partners after a time interval corresponding to 5 snapshots (i.e., 140 seconds). For IC and LC peers, this fraction is even larger.

5.2 Overlay Network

We now turn to the characterization of the structural properties of the overlay network. Table 2 presents average and CV values for each metric, computed for each experiment. Figure 6 shows the results for each snapshot, considering Experiment 1.

In all experiments, the network diameter, average shortest path and maximum degree remained roughly stable. The small CV values in Table 2 corroborate this finding. Network diameter is 4 through most of the transmissions, occasionally increasing slightly to 5. The average shortest path remained around 2 throughout all experiments. The maximum degree remained also roughly stable through the experiments. Note the large average values, falling between 339.8 and 431.5, in Table 2, indicating that, during one snapshot, a peer may exchange data with hundreds of other clients.
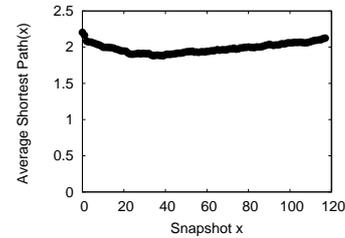
The network clustering coefficient is an exception and decreases with time. We conjecture that this is due to peers establishing new partnerships as time progresses (Figure 5). This result is also consistent with our previous finding that the probability of a peer reducing its centrality tends to be larger than the other way around (Figure 3).

6 Conclusions and Future Work

We have presented a characterization of the dynamic structural properties of the SopCast overlay network. Our analyses were performed from 2 perspectives: the local view of individual nodes and the global view of the complete network. They were based on data collected using at least 200 PlanetLab computers during 7 experiments.
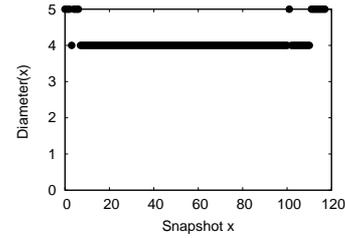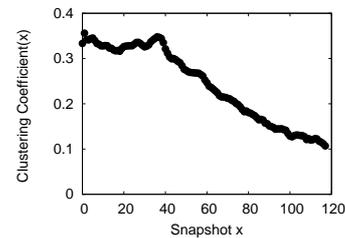

(a) Diameter

(b) Average Shortest Path
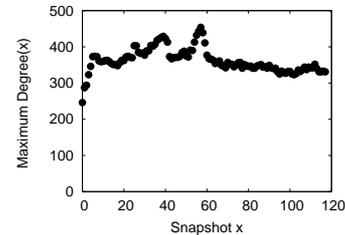
(c) Maximum Degree

(d) Clustering Coefficient

Figure 6. Network metrics during Experiment 1

We found that peers can be clustered into 3 profiles in terms of their centrality in the network. We found that a peer tends, with very high probability, to remain with the same profile throughout the transmission. When a peer changes its profile, the probability of it reducing its centrality is higher than the probability of increasing centrality. Moreover, we found that peers tend to keep most of the same partners within 2 snapshots period, the fraction of different partners tend to increase fast as time progresses. In terms of the network structure, we found that it remains roughly stable throughout a transmission, when it comes to diameter, average shortest path and maximum degree. The network clustering coefficient tends to decrease with time, possibly in response to changes in the partnerships of individual nodes.

As future work, we intend to validate these results with other applications as well as exploiting these findings to build realistic P2P live streaming simulation environments.

References

[1] I. Bermudez, M. Mellia, and M. Meo. Passive characterization of sopcast usage in residential isps. *IEEE International Conference on Peer-to-Peer Computing(P2P'11)*, 2011.

[2] D. Easley and J. Kleinberg. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge University Press, New York, NY, USA, 2010.

[3] B. Fallica, Y. Lu, F. Kuipers, R. Kooij, and P. V. Mieghem. On the quality of experience of sopcast. In *Proceedings of the 2008 The Second International Conference on Next Generation Mobile Applications, Services, and Technologies*, pages 501–506, Washington, DC, USA, 2008. IEEE Computer Society.

[4] X. Hei, C. Liang, J. Liang, Y. Liu, and K. W. Ross. A measurement study of a large-scale p2p iptv system. *IEEE Transactions on Multimedia*, 9, 2007.

[5] X. Hei, Y. Liu, and K. Ross. Iptv over p2p streaming networks: the mesh-pull approach. *IEEE Communications Magazine*, 2008.

[6] D. Menasce and V. Almeida. *Scaling for E-Business: Technologies, Models, Performance, and Capacity Planning*. Prentice Hall, 2007.

[7] A. Pathak, H. Pucha, Y. Zhang, Y. C. Hu, and Z. M. Mao. A measurement study of internet delay asymmetry. In *Proceedings of the 9th international conference on Passive and active network measurement*, PAM'08, pages 182–191, Berlin, Heidelberg, 2008. Springer-Verlag.

[8] L. Peterson, A. Bavier, M. E. Fiuczynski, and S. Muir. Experiences building planetlab. In *In Proceedings of the 7th USENIX Symp. on Operating Systems Design and Implementation (OSDI*, 2006.

[9] A. Sentinelli, G. Marfia, M. Gerla, L. Kleinrock, and S. Tewari. Will IPTV ride the peer-to-peer stream? [Peer-to-Peer Multimedia Streaming]. *IEEE Communications Magazine*, 45(6):86–92, June 2007.

[10] T. Silverston and O. Fourmaux. Measuring p2p iptv systems. *NOSSDAV'07: Proceedings of Network Operating Systems Support for Digital Audio Video*, 2007.

[11] T. Silverston, O. Fourmaux, K. Salamatian, and K. Cho. On fairness and locality in p2p-tv through large-scale measurement experiment. In *GLOBECOM'10*, 2010.

[12] D. Stutzbach, R. Rejaie, and S. Sen. Characterizing unstructured overlay topologies in modern p2p file-sharing systems. *IEEE/ACM Trans. Netw.*, 16:267–280, April 2008.

[13] S. Tang, Y. Lu, J. M. Hernández, F. Kuipers, and P. Mieghem. Topology dynamics in a p2ptv network. In *Proceedings of the 8th International IFIP-TC 6 Networking Conference*, NETWORKING '09, pages 326–337, Berlin, Heidelberg, 2009. Springer-Verlag.

[14] S. J. Wan, S. K. M. Wong, and P. Prusinkiewicz. An algorithm for multidimensional data clustering. *ACM Trans. Math. Softw.*, 14:153–162, June 1988.

[15] D. J. Watts and S. H. Strogatz. Collective dynamics of small-world networks. *Nature*, 393(6684):440–442, June 1998.

[16] C. Wu, B. Li, and S. Zhao. Magellan: Charting large-scale peer-to-peer live streaming topologies. In *In Proc. of the 27th International Conference on Distributed Computing Systems (ICDCS)*, 2007.