

SocialTube: P2P-Assisted Video Sharing in Online Social Networks

Haiying Shen, *Senior Member, IEEE*, Ze Li, *Student Member, IEEE*, Yuhua Lin, and Jin Li, *Fellow, IEEE*

Abstract—Video sharing has been an increasingly popular application in online social networks (OSNs). However, its sustainable development is severely hindered by the intrinsic limit of the client/server architecture deployed in current OSN video systems, which is not only costly in terms of server bandwidth and storage but also not scalable with the soaring amount of users and video content. The peer-assisted Video-on-Demand (VoD) technique, in which participating peers assist the server in delivering video content, has been proposed recently. Unfortunately, videos can only be disseminated through friends in OSNs. Therefore, current VoD works that explore clustering nodes with similar interests or close location for high performance are suboptimal, if not entirely inapplicable, in OSNs. Based on our long-term real-world measurement of over 1,000,000 users and 2,500 videos on Facebook, we propose SocialTube, a novel peer-assisted video sharing system that explores social relationship, interest similarity, and physical location between peers in OSNs. Specifically, SocialTube incorporates four algorithms: a social network (SN)-based P2P overlay construction algorithm, an SN-based chunk prefetching algorithm, chunk delivery, and scheduling algorithm, and a buffer management algorithm. Experimental results from a prototype on PlanetLab and an event-driven simulator show that SocialTube can improve the quality of user experience and system scalability over current P2P VoD techniques.

Index Terms—Video-on-demand (VoD), on-line social networks, peer-to-peer networks, peer-to-peer assisted VoD

1 INTRODUCTION

ONLINE Social Networks (OSNs) (e.g., Facebook, Twitter) are now among the most popular sites on the Web. An OSN provides a powerful means of establishing social connections and sharing, organizing, and finding content. For example, Facebook presently has over 500 million users. Unlike current file or video sharing systems (e.g., BitTorrent and YouTube), which are mainly organized around content, OSNs are organized around users. OSN users establish friendship relations with real-world friends or virtual friends, and post their profiles and content such as photos, videos, and notes to their personal pages.

Video sharing has been an increasingly popular application in OSNs, enabling users to share their personal videos or interesting videos they found with their friends. Indeed, according to comScore Releases in August 2010, Facebook is now the second-largest online video viewing platform. The total time spent on video viewing on Facebook increased 1840 percent year-over-year, from 34.9 million minutes in October 2008 to 677.0 million minutes in October 2009. During the same time period, the number of unique video viewer increased by 548 percent and total number of streams grew by 987 percent [1].

The recent rapid development of OSN video sharing applications illustrates the evolution of OSNs from simply communication focused tools to a media portal. OSNs are

transforming from a platform for catching up with friends to a venue for personal expression and for sharing a full variety of content and information. However, OSN's further advancement is severely hindered by the intrinsic limits of the conventional client/server architecture of its video sharing system, which is not only costly in terms of server storage and bandwidth but also not scalable with the soaring amount of users and video content in OSNs. For example, the world's largest video sharing website, YouTube, spends roughly \$1,000,000/day for its server bandwidth [2]. This high and ever rising expense was one of the major reasons that YouTube was sold to Google. OSNs are now facing the same formidable challenge as YouTube, as more and more users rely on Facebook for video sharing. Though OSNs can depend on content delivery networks (CDNs) for video content delivery (e.g., Facebook depends on Akamai for video delivery), the CDN service is costly.

In recent years, much effort has been devoted to improving the client/server architecture for video sharing, with the peer-to-peer (P2P) architecture being the most promising. P2P-based video sharing has been used in on-demand video streaming (e.g., GridCast and Vanderbilt VoD). With each peer contributing its bandwidth to serving others, the P2P architecture provides high scalability for large user bases. Previous P2P VoD systems either randomly cluster peers for video inquiry [3], [4], [5], [6], [7], [8], [9], [10], [11], [12], [13], [14], [15], [16], [17], [18] or form certain peers into a distributed hash table (DHT) for chunk indexing [19], [20], [21]. In order to reduce the video transmission and/or prefetching delay, some works cluster nodes with close physical proximity [7], [8], [19] or similar interests [9], [21]. However, those mechanisms are suboptimal, if not entirely inapplicable, in OSNs. Unlike VoD systems that provide system-wide video searching and

• H. Shen, Z. Li, and Y. Lin are with the Department of Electrical and Computer Engineering, Clemson University, Clemson, SC 29634 USA. E-mail: {shenh, zel, yuhual}@clemson.edu.

• J. Li is with Microsoft Research, Redmond, WA 9805 USA. E-mail: jinli@microsoft.com.

Manuscript received 27 Oct. 2012; revised 12 Mar. 2013; accepted 6 May 2013. Date of publication 22 May 2013; date of current version 13 Aug. 2014. Recommended for acceptance by X. Tang.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.

Digital Object Identifier no. 10.1109/TPDS.2013.139

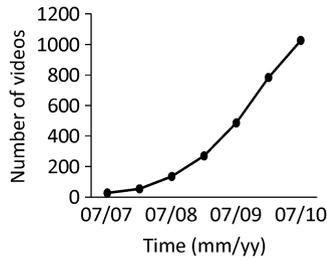


Fig. 1. Number of videos uploaded.

sharing, where a peer can access any other peer's content, OSNs do not provide video search functionality. In an OSN, videos are visited and spread by the users' friends through the Friend-of-Friend (FOF) relationship. Specifically, a friend of user *X*, say user *Y*, receives notification when user *X* uploads a video, and user *Y*'s friend *Z* is informed about the video only when user *Y* "shares" the video, and so on. Facebook has options to share videos only between friends, FOFs (default), or everyone. Therefore, users in an OSN watch videos driven much more by the friendship relation than video content.

In order to investigate the video watching behaviors of users in OSNs, we queried more than 1,000,000 users and retrieved about 2,500 public visible videos metadata on Facebook. Our measurement reveals that 1) most of the viewers of a user's videos are the user's close friends, 2) most video views are driven by social relationships, and the rest are driven by interests, and 3) viewers of the same video tend to reside in the same location. Based on our observations, we propose SocialTube, a system that explores the social relationship, interest similarity and location to enhance the performance of video sharing in OSNs. Specifically, an OSN has a social network (SN)-based P2P overlay construction algorithm that clusters peers based on their social relationships and interests. Within each cluster, nodes are connected by virtue of their physical location in order to reduce video transmission latency. SocialTube also incorporates an SN-based chunk prefetching algorithm to minimize video playback startup delay. We have conducted extensive experiments in an event-driven simulator and implemented a prototype on PlanetLab to evaluate the performance of SocialTube. Performance results show that SocialTube greatly reduces the workload of the server, improves the quality of playback, and scales well to a large client population. In the supplemental material, we present a chunk delivery and scheduling algorithm and a buffer management algorithm to further enhance the performance of SocialTube.

To our knowledge, this work is the first that studies the distinct characteristics of OSN video sharing that vary from other content-based system-wide video sharing, and builds a P2P-based video sharing system in an OSN by leveraging those characteristics for higher performance. Our previous conference version of this article [22] introduces the basic trace data analysis and design of SocialTube. This article presents more trace data analytical results. It also presents new SocialTube mechanisms including locality-aware video prefetching mechanism, two policies to increase the chunk delivery abilities, and buffer management algorithm. This article further presents more simulation results

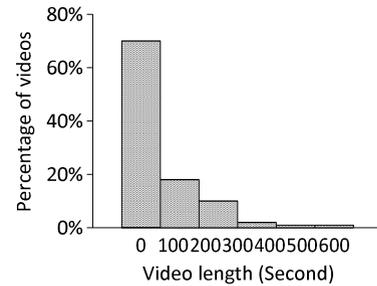


Fig. 2. Distribution of video length.

and the experimental results for the SocialTube prototype on the real-world PlanetLab testbed.

The rest of the paper is organized as follows. Sections 2 and 3 present the Facebook trace data analysis and the design of SocialTube. Sections 4 and 5 present the trace-driven simulation results and experimental results on PlanetLab. We conclude this paper with remarks on future work in Section 6. The supplemental material presents two enhanced methods including chunk delivery and scheduling and buffer management algorithm. It also presents additional experimental results and an overview of related work.

2 FACEBOOK MEASUREMENT AND ANALYSIS

In this section, we present our Facebook trace measurement results and give an in-depth perspective of Facebook video viewing patterns, that shows the necessity of peer assistance in OSN video sharing and provides a direction for the design of a P2P video sharing system in OSNs.

We used breadth-first-search [23] to query over 1,000,000 users seeded by 5 users in the USA. In order to avoid overloading the Facebook, we sent a query to Facebook every 5 s. We can only see the video activities of the users who are friends or FOFs of the crawler and the users that chose "everyone" as their video access option. Because of this access limit, we only found about 2,500 videos and 12,000 users who watched these videos during the time period from July 2007 to Aug. 2010, which is used as a sample for the video sharing and watching activities. The collected dataset includes the information about user friendship relations, interests, locations, and videos uploaded and shared by users. For each video, we retrieved the video metadata such as its title, length, and viewers when available. To respect the privacy of the users, we anonymize the user names before storing the data in our database. We only crawled the video metadata of these users, with other personal information untouched.

2.1 Popularity of Videos on Facebook

First, we investigate the popularity of videos on Facebook over the years. Fig. 1 plots the number of videos corresponding to the time they are uploaded in our collected video pool. It shows that the number of videos uploaded to Facebook increases sharply along with time. Since Facebook launched its video service in 2007, the increasing trend of video uploading has never slowed down, making it one of most popular applications on Facebook. Fig. 2 shows the video length distribution in our collected video pool. From the figure, we can see that about

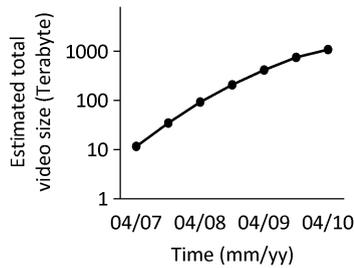


Fig. 3. Estimated total size.

70 percent of the videos are less than 100 s. Videos longer than 200 s account for less than 10 percent of all videos. It may be because users generally share short user-generated videos of their lives with their friends in OSNs. Based on the measured results, we make an observation (O).

2.1.1 O1: Video Sharing in Facebook Is Becoming Increasingly Popular, and Most Are Very Short

Typical Facebook videos have a bitrate of [300, 600] kbps. Based on the statistics from [1], the total number of videos on Facebook was 41,537,000 in April, 2009. Using the increasing trend derived from Fig. 1, we calculated the approximate number of videos on Facebook at different times. Using the average video length calculated from Fig. 2, we obtain an estimate of the server storage cost for videos as shown in Fig. 3. The figure indicates that the storage cost of videos has increased 100 times from 10 Terabytes to 1000 Terabytes, which is a great burden for the server considering the high bandwidth cost caused by uploading and downloading videos.

2.1.2 O2: The Bandwidth and Storage Burden on Video Servers for Providing Video Services Is High and Has Been Increasing at a Rapid Rate

According to the comScore [1] study, users on average spent 9.9 percent of their total online time on Facebook. Fig. 4 shows the total time users spent on Facebook over years according to the comScore statistics. We see that the total time users spent on Facebook has been increasing quickly. It was reported that web users spend more time on Facebook than Google sites. On average, more than 8 billion minutes are spent on Facebook every day. As long as a user is online, his/her cached videos can be fetched for P2P video sharing. This makes a P2P video sharing system very suitable for Facebook since a fundamental prerequisite of P2P video sharing systems is that there are enough peers online to participate in video sharing.

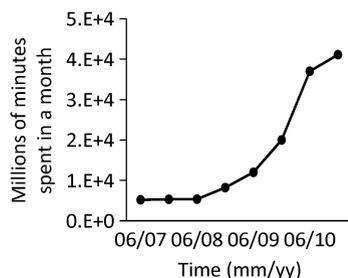


Fig. 4. Time spent on Facebook.

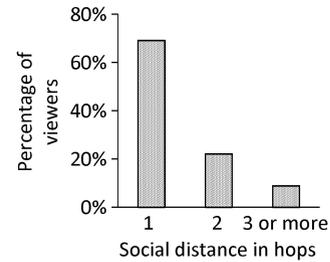


Fig. 5. Social distance.

2.1.3 O3: Facebook Users Normally Spend Long Time Periods Online, Which Makes P2P Video Sharing in Facebook Feasible

2.2 Effect of Social Distance on Video Viewing Patterns

Social distance between two users in the social network graph represents the closeness of their relationship. If two users are directly connected in the social network, their social distance is 1; if one user is a friend of another user's friend, then the social distance between them is 2, and so on. Next, we investigate the impact of social distance on user video viewing patterns.

Among 52,500 video watching activities involving 12,000 users, we measured the social distance of a video viewer from the video owner, and show the distribution in Fig. 5. We find that most of the viewers (around 70 percent) are 1-hop friends of the video owner, 2-hop viewers account for a portion of about 20 percent, and the remaining 10 percent of viewers watched videos of video owners are more than 2 hops away.

2.2.1 O4: In Facebook, More Than 90 Percent of the Viewers of a Video Are Within 2 hops in the Video Owner's Social Network

We define a video viewer group as all users who have watched the video owner's videos. From O4, we obtain the inference (I):

2.2.2 I1: A Video Viewer Group of a Video Owner in Facebook is Mostly Within the 2-Hop Friend Circle of the Owner

Note that a user may own more than one video. To further identify the impact of social relationships on video viewing patterns, we selected the users who have multiple videos from our dataset and inspected the viewer group of each

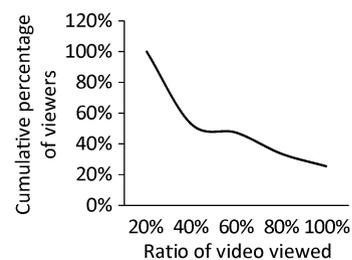


Fig. 6. Different viewer types.

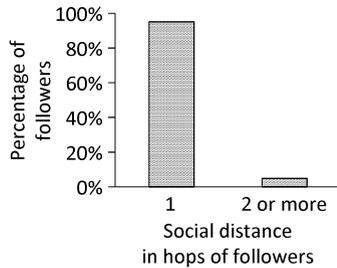


Fig. 7. Social distance between follower and video owners.

video owner. We classified the viewers of a video owner based on the percentage of the owner’s videos they watched and calculated the distribution of different viewer classes in a viewer group. Fig. 6 shows the average values of the percent versus the ratio of videos watched from the video owners. We observe that:

2.2.3 O5: On Average, in a User’s Viewer Group, 25 Percent of Viewers Watched All, 33 Percent of Viewers Watched 80 Percent, and All Viewers Watched 20 Percent of the User’s Videos

We call the viewers who have watched almost all videos of a user the user’s *followers*, and call other viewers *non-followers*. We use a threshold T_h for the percent of all the videos of a user that a viewer watches in order to become a follower, and set $T_h = 80\%$ in this analysis. Figs. 7 and 8 show the distribution of followers and nonfollowers in terms of the social distance with the video owner. We can see from the figures that:

2.2.4 O6: Viewers That Watch Almost All of a User’s Videos (i.e., Followers) Usually Are 1-hop Friends of the User, While Most of Other Viewers (i.e., Nonfollowers) Are 1-hop or 2-hop Friends of the User

2.3 Effect of Interest on Video Viewing Pattern

Next, we explore the correlation between user interests and video viewing patterns. We selected a sample of 118 distinct users that watched more than one video from our dataset and manually classified the videos they watched into 19 interest groups based on video content. The 19 interest groups were determined based on the video categories in YouTube such as gaming, rock music and action movie. For each user, we calculated the percentage of viewed videos of each interest group. Then, we ranked

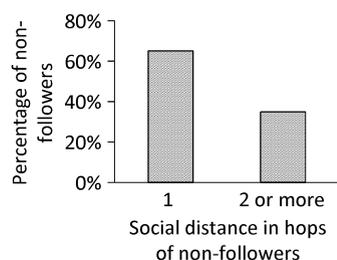


Fig. 8. Social distance between nonfollower and video owners.

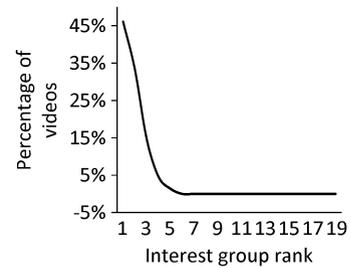


Fig. 9. Interest clustering effect.

these 19 interest groups in descending order of the percentage values. We calculated the average percentage value of the 118 users for each interest group rank and show the result in Fig. 9. We observe that, on average, 46 percent of videos a user watched are on his/her top 2 interests topics, 79 percent of videos a user watched are on his/her top 3 interests topics, and 94 percent are on his/her top 4 interests topics. The result implies that the videos each user watches are generally orientated towards his/her few primary interests.

2.3.1 O7: Users Tend to Watch the Videos of Their Interests and Each User Generally Has ≤ 4 Video Interests

A user can post on Facebook either self-uploaded videos or external video links from a third party video service provider such as YouTube. The video linking in Facebook is called “share”, by which users can share links to videos they find interesting with their friends. Fig. 10 shows the different sources of videos in our collected dataset. We see that the self-uploaded videos in Facebook account for about 14 percent of all videos. Others are external video links. YouTube, as the largest video site in the world, accounts for over 80 percent of all external links. Many other video sources such as TED and Hulu account for the remaining percentage.

For videos uploaded to Facebook, social relationships are the primary consideration when viewers decide whether to watch a video. For external videos, interest gains more weight in influencing the watching probability. However, no matter if it is a video uploaded to Facebook or an external link, we assume that the video owner, the video uploader, or the one who shared the external video link would have the video in his/her local cache. This assumption is reasonable because in the case of uploaded videos, the uploaders have the original videos, while in the

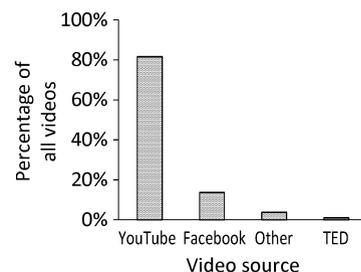


Fig. 10. Video source in Facebook.

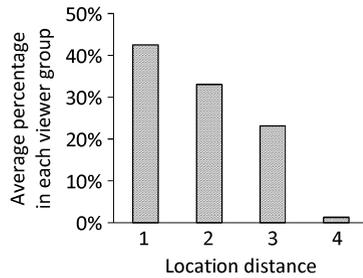


Fig. 11. Physical distance.

case of external links, a user usually shares a video after (s)he has watched the video (or at least part of it) and hence has the video in cache. Therefore, both types of videos are applicable in P2P video sharing, in which videos are prefetched from users' local cache instead of the video server.

2.3.2 O8: A Large Percentage of Videos in Facebook Are from YouTube, Where the User Video Viewing Patterns Are Driven by Interests

Combining O4-O8, we can find that different watching incentives can be applied to different types of viewers. The followers of a user watch most of the user's videos regardless of the video content because of their close social relationship (e.g., close friends and fervent admirers). For the viewers that watch only a few of the user's videos, interest in the video content is a more important incentive. Additionally, some show a mixed video watching incentive. Thus, we infer.

2.3.3 I2: Followers Are Primarily Driven by Social Relationship to Watch Videos, While Nonfollowers Are Driven Mainly by Interest

2.4 Effect of Physical Location on Video Viewing Patterns

We also analyze the geographical locations of users who view the same videos in order to see whether location can also be leveraged for video sharing in OSNs. In Facebook, some users input their current resident city in their profiles. The location is in the format of "city, state" (e.g., Los Angeles, CA). For international users, the location is in the format of "city, country" (e.g., Tokyo, Japan). We denote two users located in the same city, state, or country as 1, 2, and 3, respectively, and two users located in different countries as 4. To investigate the location distribution of viewers, we calculated the percentage of viewers in each viewer group corresponding to different location distances between the viewer and video owner. We plot the average value of all viewer groups in Fig. 11. From the figure, we can see that most users watching the same video are physically close to each other. Because many friend relationships in Facebook are connected by offline relationship, such as classmates or colleagues, this produces a strong location clustering effect. This result conforms to the observation in [24] that most of the wall posts are sent within local physical region. This effect could make P2P video sharing systems in

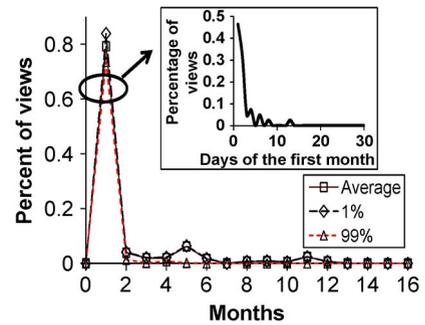


Fig. 12. Video active period.

OSNs more efficient by enabling geographically close nodes to share videos between each other.

2.4.1 O9: Most Users Watching the Same Video Are Physically Close, and on Average About 40 Percent of Users Watching the Same Video Are in the Same City

2.5 Active Life Period of Videos

We measured the percentage of views of a video in each month after the video is uploaded out of all views. Fig. 12 shows the average of these values. We found that videos in Facebook have an active life period of about one month. Views in this period account for more than 90 percent of all views. After one month, there are only occasional views. The small figure inside Fig. 12 more clearly shows the decreasing active life over days in the first month. We find that it follows a power-law distribution.

2.5.1 O10: The Videos in Facebook Have an Active Life Period of About One Month on Average. The Decrease in the Number of Views of a Video Follows an Exponential Distribution

3 THE DESIGN OF SOCIALTUBE

The Facebook measurement in Section 2 shows that video sharing in Facebook is increasingly popular (O1) and may generate a heavy burden on the video server (O2). Fortunately, the length of time that the users stay online in Facebook is rapidly increasing (O3), enabling the possibility for P2P video sharing among the online users themselves. Therefore, P2P-assisted video sharing is a promising strategy in OSNs. Based on observations O4-O10 on the characteristics of video viewer behavior in Facebook, we propose SocialTube, a P2P video sharing system for OSNs.

We first introduce the basic concepts and strategies used in SocialTube. In Facebook, each node can upload a video to the Facebook video server or an external link to a video from an external server. In this paper, we use *server* to represent all video source servers, including both Facebook and external video servers. Similar to current peer-assisted content delivery mechanisms, the peers in SocialTube

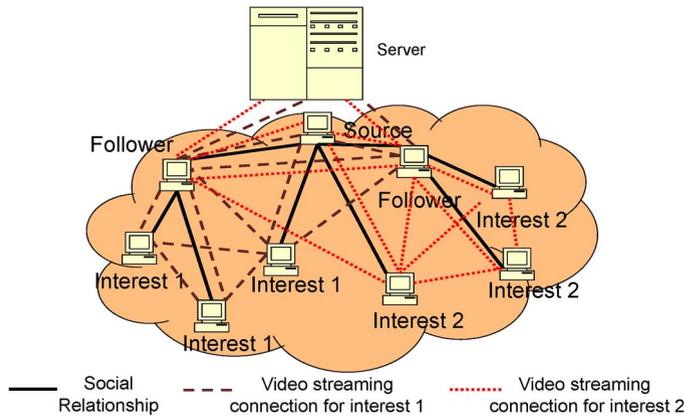


Fig. 13. Structure of SocialTube.

store videos they have watched previously for video redistribution. In SocialTube, a video is divided into small chunks with a fixed size. Thus, a video viewer only needs to download the corresponding chunks of the video segment (s)he wants to watch.

3.1 Social Network-Based P2P Overlay Construction Algorithm

To identify followers and nonfollowers of a source node for structure construction, SocialTube pre-defines two thresholds, T_h and T_l , for the percent of videos in the source node that a viewer has watched during a time unit, say one week. If the percent value of a viewer is $\geq T_h$, the viewer is a follower. If the percent is $T_l < x \leq T_h$, the viewer is a nonfollower.

Video sharing in Facebook distinguishes itself from other video sharing websites (e.g., YouTube) in two aspects: video sharing scope and video watching incentives. First, other websites provide system-wide video sharing where a user can watch any video, while in Facebook, videos are usually shared in a 2-hop small circle of friends (I1). Second, users in other video sharing websites are driven to watch videos by interests, while in Facebook, the followers of a source node (i.e., video owner) are driven to watch almost all of the source's videos primarily by social relationship, and nonfollowers watch a certain amount of videos mainly driven by interest (I2). According to these differentiating aspects, we design the P2P overlay structure, which is shown in Fig. 13.

Based on I1, SocialTube establishes a *per-node* (in contrast to *per-video* in YouTube) P2P overlay for each source node. It consists of peers within 2 hops to the source that watch at least a certain percentage ($> T_l$) of the source's videos. Other peers can still fetch videos from the server. As shown in the figure, such peers of a source node S in the social network constitute a P2P overlay for the source node. We aim to achieve an optimal tradeoff between P2P overlay maintenance costs and video sharing efficiency. Some nodes within 2 hops may watch only a few videos in a source. Including these nodes and users beyond 2-hops into the overlay generates a greater structure maintenance cost than video sharing benefits. Based on I2, we build a hierarchical structure that connects a source node with its socially-close followers, and connects the followers with other non-

followers. Thus, the followers can quickly receive chunks from the source node, and also function as a pseudo-source to distribute chunks to other friends. The source pushes the first chunk of its new video to its followers. The chunk is cached in each follower and has high probability of being used since followers watch almost all videos of the source. Further, nonfollowers sharing the same interest are grouped into an interest cluster for video sharing. We call peers in an interest cluster *interest-cluster-peers*. A node that has multiple interests is in multiple interest clusters of the source node.

Because the source node and followers are involved in every interest cluster for providing video content, we call the group formed by the source, followers, and interest-cluster-peers in an interest cluster *swarm*, and call all nodes in a swarm *swarm-peers*. As I1 indicates, the cluster size of each interest cluster should be small. O9 indicates that many viewers of a video are physically close peers. Therefore, in order to reduce delay, physically close interest-cluster-peers are randomly connected with each other. The peers find their physically close peers based on their ISP, subnet information [25]. To preserve the privacy protection on OSN, we can add a constraint in which peer A can connect to peer B only when peer A is peer B's friend or can access peer B's shared videos. In Fig. 13, the viewers of S form into two swarms. Because the nodes in each swarm have a high probability of owning chunks of the same video, they can retrieve chunks from their swarm-peers without relying on querying the server or large scale query flooding.

In current video sharing in Facebook, a node always requests the server for videos uploaded by source nodes. We let the server keep track of the video watching activities of viewers of a specific source node in order to identify and update its followers and nonfollowers based on SocialTube's pre-defined thresholds of T_l and T_h . This duty can be assigned to the source node itself if it has sufficient capacity. The nodes in the system will periodically report their video providing activities to the server. When the server determines that a peer is a follower of the source node, it notifies the source node, which notifies all nodes in its swarms about the follower. Consequently, the follower becomes a member of each of the swarms, and all swarm-peers in each of the swarms connect to it. When the server determines that a peer is a nonfollower of the source node, the server determines its interests based on the contents of videos the peer visited, and notifies the source node about the nonfollower along with its interests. The source node then notifies the peers in the clusters of the interests of that nonfollower, and notifies the nonfollower about the clusters. The nonfollower connects to all followers and the source and to a few physically close nodes in each cluster. Consequently, the nonfollower becomes a member of the swarm of each of the interest clusters. The server also periodically updates the roles of the followers and nonfollowers. If a node becomes neither follower nor nonfollower, the server removes it by notifying others to disconnect from the node. If a follower becomes a nonfollower, its connections are also updated accordingly. In an OSN, after a node logs out (i.e., leaves) the system, it will always logs in (i.e., joins in), so the SocialTube overlay does not update due to node departures. Only when the

TABLE 1
Bandwidth Capacity and Distribution of Users

Groups	Downloading bandwidth	Uploading bandwidth	Percentage of nodes
1.	768k/s	128k/s	21.4%
2.	1536k/s	384k/s	23.3%
3.	3072k/s	768k/s	55.3%

server notices that a node did not join in after a very long time departure, the node is removed from the overlay. Neighbors in the overlay periodically exchange messages. When a node notices that its neighbor is offline, it marks the connection and will unmark it when the neighbor is online. A node does not request videos from neighbors with marked connections. The nodes in a P2P structure, including the source, followers and nonfollowers, remember their roles and connections. Next time when a node goes online, it automatically connects to its previous neighbors and function based on its role.

As indicated in Fig. 13, the source node has two followers, and its videos can be divided into two interest categories based on video content. The 1-hop and 2-hop friends of the source node with interest 1 and interest 2 form into two clusters, respectively. The source node and the followers are in each interest cluster, all of which form a swarm.

3.2 Social Network Based Prefetching Algorithm

To reduce the video startup latency, we propose a push-based video prefetching mechanism in SocialTube. In SocialTube, when a source node uploads a new video to the server, it also pushes the prefix (i.e., first chunk) of the video to its followers and to the interest-cluster-peers in the interest clusters matching the content of the video. The prefix receivers store the prefix in their cache. Those interest-cluster-peers and followers who are not online when the source node pushes the prefix will automatically receive it from the source node or the server once they come online. After the source node leaves, the responsibility to push the prefix falls to the server. Since these followers and interest-cluster-peers are very likely to watch the video, the cached prefixes have a high probability of being used.

Once the nodes request the videos, the locally stored prefix can be played immediately without delay. Meanwhile, the node tries to retrieve the remaining video chunks from its swarm-peers. Similar to BitTorrent, SocialTube allows a requester to request 4 online nodes at the same time to provide the video content in order to guarantee provider availability and achieve low delay by retrieving chunks in parallel. It first contacts interest-cluster-peers, then followers, then the source node. If the requester still cannot find 4 providers after the source node is contacted, it resorts to the server as the only provider. Considering the high capacity of the server, the requester does not need to have 4 providers if it has the server as a provider. This querying order can distribute the load of chunk delivery among the swarm-peers while providing high chunk availability. The algorithm takes advantage of all resources for efficient video sharing without overloading specific nodes. The server can guarantee the availability of the

TABLE 2
Experiment Default Parameters

Parameter	Default value
Number of clients	5,000
Number of videos	2,000
Number of interest categories	19
Number of interests per client	2-4
Trace duration	40 days
Chunk size	3 MBytes
Prefix length	3 MBytes
Server uploading bandwidth	20Mbps
Video size	Distribution of YouTube videos
Cache size	300MBytes

video, even if the number of online users in a swarm is small.

4 PERFORMANCE EVALUATION

We evaluate the performance of SocialTube through both simulations on the event-driven simulator PeerSim [26] and PlanetLab [27] prototype implementation. We run each experiment for 10 times and report results within 95 percent confidential intervals.

4.1 Experiment Settings

Our simulation setup is based on a part of our crawled dataset, which contains approximately 2000 distinct videos. We selected 5000 nodes from the trace data. The link delay for nodes with location distance x as defined in Fig. 11 was set to x seconds. The connectivity degree of these selected nodes follows power-law distribution with skew factor equals to 0.5 [28]. We assigned each of the 2000 videos to a randomly chosen node in the system. To simulate the real Internet environment, where peers have heterogenous bandwidth, we used the statistics in Table 1 used in [9]. To simulate the geographic locations of nodes, each node has a location ID in [1], [2], [3], [4], [5], [6], [7], [8], [9], [10]. The nodes with numerically closer ID are geographically closer nodes. We used the distribution of friendship in spatial proximity from [29] in our experiments.

Table 2 shows the default parameters unless otherwise specified. The video bitrate was set to 330 kbps. Based on empirical observations, we assume that whenever a node cannot receive a chunk in time for viewing, it pauses for 3 s and then resumes playback. The file size of a video was randomly chosen from [20], [21], [22], [23], [24], [25], [26], [27], [28], [29], [30] MB, which is a normal size for a video with a length of 2-3 min.

Since every person spends about 420 min on Facebook per month [30] on average, we can infer that every person spends about 14 min and hence watches at most 4 videos per day on average. We used simulation cycle to simulate one day. In each simulation cycle, the number of videos a peer watched was randomly chosen from [1], [2], [3], [4]. One simulation lasted for 100 simulation cycles. Based on O7, each node in the system randomly selected [2], [3], [4] interest categories as its interests.

Based on O6, every source node randomly selected 20 percent of its friends as its followers. As the practical user

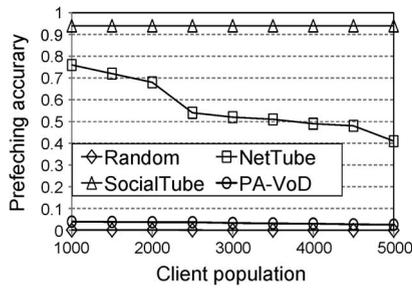


Fig. 14. Prefetching accuracy vs. node population.

video viewing behaviors in Facebook, every node can see the names of the videos owned or shared by its friends. Based on O6 and O7, unless otherwise indicated, from these videos, a node randomly selected videos in its interests or its followers to view. Recall that a node can access videos either owned or shared by its friends in Facebook. In practice, after a node watches a video, it may or may not share the video. To simulate this behavior, each node first randomly selected a number $x \in [1, 40]$ and shared its watched x th video, and repeated this process until the simulation completed.

We compare the performance of SocialTube with other representative works in peer-assisted video streaming, PA-VoD [8], NetTube [9] and Random (as a baseline). In PA-VoD, physically close peers with the same location ID are clustered for video sharing between each other. In NetTube, peers that have similar interests are clustered together for video sharing. The details of PA-VoD and NetTube are presented in Section 9. Random clusters peers randomly into 10 groups regardless of their friendships, locations and interests. Nodes in a cluster connect with each other randomly, and use flooding for video search. In each simulation cycle, a source node pushes one video prefix to its followers and the interest-cluster-peers in the interest cluster matching the video in SocialTube; the server pushes the prefix of one video to the nodes in the interest cluster of the video's uploader in PA-VoD and Random. After a node selects a video to view, it searches its interest cluster in NetTube, and searches its own cluster and a neighboring cluster in PA-VoD and Random. It finds one video provider to fetch video chunks; if it cannot find or access the requested video, it resorts to the server. To make the methods comparable, SocialTube does not include its chunk delivery and scheduling algorithm and buffer management unless otherwise indicated.

We used the following metrics in the experiments:

- *Prefetching accuracy.* This is the probability that a user requests a video whose prefix is in its cache and it can access the prefix's video. This metric shows the effectiveness of the prefetching and enhancing the video playback continuousness.
- *Chunk transmission delay.* This is the chunk transmission time between two peers. This metric shows the delay in retrieving video chunks.
- *The number of searched clients for a video.* This is the number of unique nodes that are queried before a

node finds a video provider or fails to find a video from clients. This metric shows the efficiency in video provider searching.

- *Percent of server contribution.* This is the ratio of server bandwidth consumed in a system over the total bandwidth consumed in the client/server system. The bandwidth is measured by the total size of served videos. This metric shows the effectiveness of reducing the bandwidth burden of the server.
- *Playback continuousness probability.* This is the percent of nodes that have not experienced freezing caused by waiting for a missing chunk in video playback due to a user's forward skipping when watching video.
- *Startup delay.* This is the time elapsed after a node selects a video and before the video starts to play. This metric reflects the effectiveness of a prefetching mechanism.
- *Buffering delay.* This is the total time for a user to receive a certain number of chunks after sending out a video request.
- *Average overlay maintenance cost.* This is the number of communication messages between neighboring nodes for overlay maintenance.

4.2 Effectiveness of the Prefetching Algorithm

In this experiment, we let each node randomly request videos whose prefixes are in its cache. A node cannot access a video not owned or shared by its friends. We varied the number of nodes in the system from 1000 to 5000 with 500 increase in each step. Fig. 14 shows the average prefetching accuracy versus client (i.e., node) population. We see that the prefetching accuracy of SocialTube remains higher than 0.9, those of PA-VoD and Random remain at 0.03 and nearly 0, respectively, and that of NetTube lies in the middle. Also, as the client population increases, unlike other methods, the prefetching accuracy of NetTube decreases significantly. SocialTube explores the social relationship and clusters the follower and similar-interest peers within 2-hops from a source together. A source node pushes its video prefixes to its followers and the cluster nodes of the video's interest. In a source's per-node overlay, nodes 2 hops away from the source cannot access its videos unless the videos are shared by their friends, and 1-hop friends of the source can always access the videos. Thus, SocialTube achieves high prefetching accuracy but it is not 100 percent. As the client population increases, the size of interest cluster increases. Then, more cluster-peers within 2-hops of the source can visit shared videos from their friends, thus the video prefetching accuracy remains nearly constant. NetTube only clusters peers with similar interests without considering the social relationship. When a video prefix is pushed to the cluster of the video's interest, some nodes are unable to access the video since they are not the friends of the source or the video sharers. Given a certain number of videos in the system, as the client population increases, the number of videos in a node decreases and then the number of accessible videos from friends of a node decreases, leading to a decreasing prefetching accuracy.

PA-VoD clusters physically close nodes without considering interests or social relationship. A node may not be in the same cluster with its friends or the video sharers, leading to low prefetching accuracy. In Random, nodes are randomly clustered. Thus, nodes in a cluster receiving a prefix of a video are unlikely to be the friends of the source node or the video sharers the video, leading to very low prefetching accuracy.

In this experiment, we assigned 6 videos to each source node. We randomly selected 5 per-node overlays and chose a client in each overlay. We recorded each node's prefetching accuracy when it has x ($x = 1, 2, \dots, 5$) video prefixes from the source node. Fig. 15 shows the average prefetching accuracy of the 5 nodes versus the number of prefetched videos of each node. As the number of prefetched videos increases, the prefetching accuracy of SocialTube, NetTube and PA-VoD increases. This is because as a peer receives more video prefixes, when it chooses a video to view, it has a higher probability to have the prefix of the chosen video in its cache. We also see that NetTube leads to lower prefetching accuracy than SocialTube. Also, PA-VoD and Random lead to extremely low prefetching accuracy. In SocialTube, a source pushes its video prefix to its followers and the interest-cluster of the video's interests in its per-node overlay. In NetTube, a video prefix is pushed to all nodes that share the same interest of the video's interest. Thus, when a follower chooses a video in its followee to view, the prefix of this video is in its cache in SocialTube but may not be in its cache in NetTube. In PA-VoD, a video prefix is pushed to the physically close nodes of the video uploader, which have low probability to be the source's friends or the video sharers, thus generating very low prefetching accuracy. We also find that the prefetching accuracy in Random maintains at around 0 due to the same reason explained previously.

In this experiment, we randomly selected 5 per-node overlays and chose a client in each overlay. We let each source node have 30 videos in 5 interests. A client initially has [2], [3], [4] interests. In every simulation cycle, each client receives a prefix from its source node and chooses [1], [2], [3], [4] videos to watch. If the client chooses a video not in its current interests, the video's interest is added to the client's interest list and it joins in the interest-cluster of this interest and receives the pushed prefix of videos in this interest in SocialTube and NetTube. Fig. 16 shows the average value of the prefetching accuracies of the 5 clients versus the number of watched videos of a client over time.

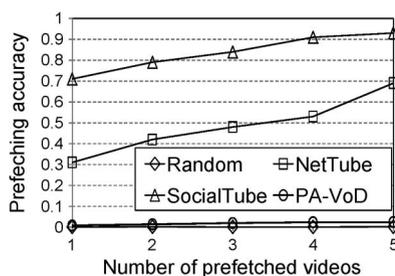


Fig. 15. Prefetching accuracy vs. number of prefetched videos.

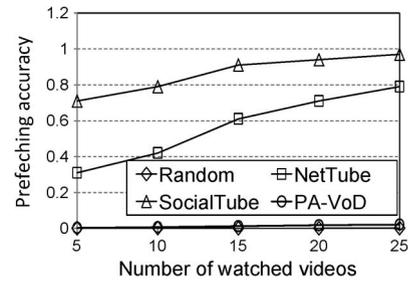


Fig. 16. Prefetching accuracy vs. number of watched videos.

The figure illustrates that as the number of watched videos increases, the prefetching accuracy of SocialTube and NetTube increases. This is because in SocialTube and NetTube, as a client watches more videos, it joins in more interest-clusters and receives more prefixes, and thus its prefetching accuracy increases. Again, NetTube generates lower prefetching accuracy than SocialTube due to the same reason as in Fig. 15. The prefetching accuracy of PA-VoD and Random is not affected by the number of watched videos as they do not have interest clusters. Also, they produce low prefetching accuracy due to the same reasons as explained before.

4.3 Scalability Performance

Fig. 17 shows the average number of searched clients for a video provider versus the client population. We see that Random and PA-VoD need to search significantly more clients than NetTube and SocialTube. In Random, as the cluster nodes are unlikely to have the same interests or friend relationship, the viewers need to search much more nodes for the requested videos. In PA-VoD, cluster nodes are physically close nodes that may have friendship, thus it has fewer searched clients for a video than Random. In contrast, in SocialTube and NetTube, the nodes that are likely to watch the same video are clustered. Therefore, they have high probability to find the requested video in their cluster. We also find that NetTube generates more searched clients than SocialTube. This is because in NetTube, nodes in one interest cluster may not have the permission to access videos in some of the nodes in the cluster due to the constraint in OSNs. Also, when a node visits its followee's video, the followee may not be in the interest cluster of the video. In SocialTube, the common-interest nodes within two hops of a source node and its followers are clustered. Thus, a node has a high probability to find and access the requested video from a cluster node.

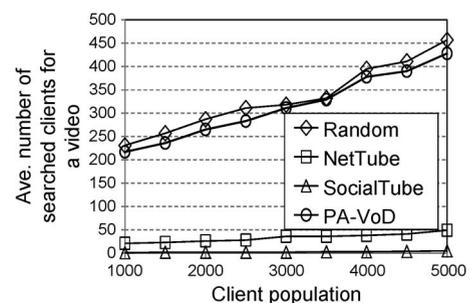


Fig. 17. Number of searched clients vs. client population.

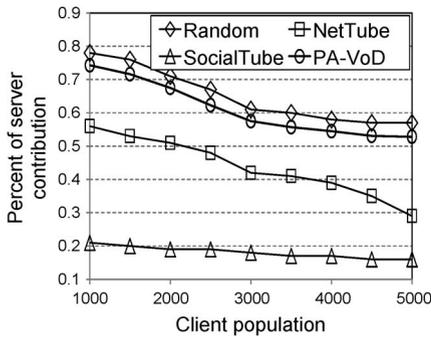


Fig. 18. Percent of server contribution vs. client population.

The figure also shows that as the client population increases, the video search overhead increases more rapidly in Random and PA-VoD than in NetTube and SocialTube. A larger client population increases the cluster size, so a node needs to search more cluster nodes in Random and PA-VoD. We notice a slight increase in the search overhead in NetTube, because when more nodes are included in a cluster in NetTube, the probability that a node can find a friend decreases. In SocialTube, as only the common-interest nodes within 2-hops and followers of a source node are clustered, the client population increase does not greatly increase a cluster size, so both the video access probability and provider location probability remain approximately the same.

In the following experiments in this section, we let each requester find four video providers to fetch video chunks at the same time. Fig. 18 illustrates the percent of server contribution versus client population. It shows that as the client population increases, the percent of server contribution in all systems decreases. Higher client population increases the number of friends of a node that can provide services. It also causes more clients viewing videos, leading to faster video dissemination to a wider friend circle around a video owner. Then nodes have higher probability of finding videos from peers, thus reducing the bandwidth consumption of the server. We also see that the percent of server contribution follows $\text{SocialTube} < \text{NetTube} < \text{PA-VoD} < \text{Random}$. This is because the nodes in SocialTube can locate the video chunks from other peers more efficiently, since it considers both social relationship and interest. Common-interest nodes are more likely to be friends than physically close nodes, so NetTube can locate

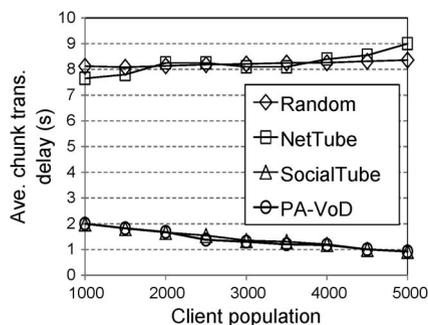


Fig. 19. Chunk transmission delay vs. client population.

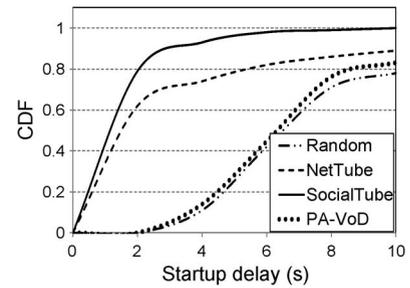


Fig. 20. CDF of node startup delay.

video providers more efficiently than PA-VoD. Randomly clustered nodes in Random are unlikely to be friends, so Random produces the lowest video provider location efficiency.

Fig. 19 plots the average chunk transmission delay versus client population. We did not consider the transmission from the central server in calculating this metric. We see that SocialTube and PA-VoD have almost the same chunk transmission delay, which is dramatically lower than those of NetTube and Random. In SocialTube and PA-VoD, the nodes contact the physically close nodes for the video chunks. NetTube and Random do not consider the physical locations of nodes. We also see that as the client population grows, the chunk transmission delay of SocialTube and PA-VoD decreases slightly, while that of NetTube and Random exhibit a marginal increase. As client population increases, a video requester has a higher probability to find physically nearby video holder in SocialTube and PA-VoD, while the average distance between video requester and video provider increases in NetTube and Random.

5 EXPERIMENTAL RESULTS ON PLANETLAB

We implemented SocialTube based on the Facebook video trace on the PlanetLab real-world testbed [27]. We randomly chose 268 online nodes in the PlanetLab and installed the server at 128.112.139.26 located in Princeton University. We used a smaller set of 50 distinct videos from our dataset for testing.

Fig. 20 shows the CDF of the average startup delay. It shows that about 80 percent of the nodes in SocialTube and 65 percent of the nodes in NetTube have a startup delay of less than 2 s, but the startup delay of all nodes in PA-VoD and Random is larger than 2 s and the startup delay of 20 percent of their nodes is even larger than 10 s. This is

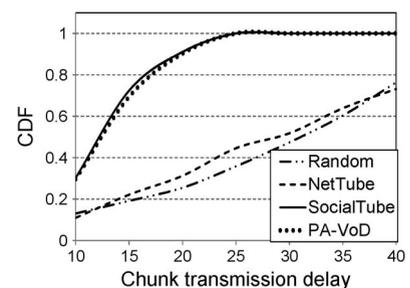


Fig. 21. CDF of chunk transmission delay.

mainly due to the high effectiveness of the prefix prefetching mechanism in SocialTube and NetTube and its low effectiveness in PA-VoD and Random as explained previously.

Fig. 21 shows the CDF of the average chunk transmission delay between nodes. We find that NetTube and Random generate much higher delay than SocialTube and PA-VoD. Since NetTube and Random do not consider the node geographic locations, a node may contact geographically distant nodes for video chunks, which increases the chunk transmission delay. Both SocialTube and PA-VoD try to fetch chunks from geographically close nodes, thus their chunk delivery delays remain low and are almost the same.

Fig. 22 shows the CDF of the percent of server contribution for video transmission. We see that about 90 percent of the nodes in SocialTube and 80 percent of the nodes in NetTube have no more than 0.6 server contribution. Around 60 percent of the nodes have more than 0.7 and 0.78 server contribution in PA-VoD and Random, respectively. Since SocialTube considers social relationship and interests in node clustering, the probability that nodes find video providers is high. Therefore, SocialTube has much less bandwidth demand from server than other methods. NetTube clusters the nodes with the same interests. As followers may not find requested videos of their followees in their interest clusters in NetTube, so it produces higher server contribution than SocialTube. PA-VoD clusters the physically close nodes that may not be friends, leading to lower probability of finding video providers hence higher server contribution than NetTube. Random only randomly clusters nodes, so it incurs the highest server bandwidth contribution. The result shows that SocialTube is most effective in reducing the bandwidth burden on the central video server.

Fig. 23 shows the CDF of the accumulated freezing delay, which is defined as the total freezing delay when a video is playing. We see that the percentage of the users that did not experience any freezing delay in SocialTube, PA-VoD, Random and NetTube is 42 percent, 30 percent, 28 percent and 15 percent, respectively. Also, 80 percent of the users experienced accumulated freezing delay less than 5 s, 7 s, 9 s, and 12 s in SocialTube, PA-VoD, Random and NetTube, respectively. That is, the users experience more video freezing delay in NetTube than in PA-VoD and Random, and experience the smallest freezing delay in SocialTube. In SocialTube, the video chunks can be more accurately prefetched since nodes push video chunks to their friends that are very likely to watch the videos. The interest-based prefetch mechanism in NetTube and locality-

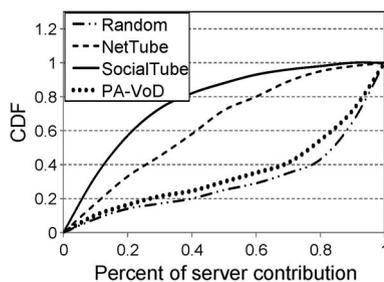


Fig. 22. CDF of the percent of server contribution in different systems.

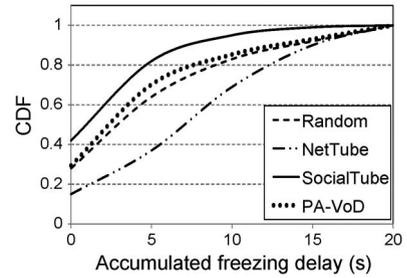


Fig. 23. CDF of accumulated freezing delay.

based prefetch mechanism in PA-VoD are not as accurate as SocialTube. In addition, in SocialTube, when a user watches the first chunk of the video, the other video chunks are transmitted from other users that are physically close to the user. Therefore, the nodes are unlikely to experience playback freezing. In PA-VoD, nodes receive chunks from either physically close nodes or the server, both of which provide quick transmission. Therefore, PA-VoD produces the second smallest freezing time. In Random, nodes receive chunks mainly from the server, so chunk transmission delay is relatively low. As NetTube can locate video providers with high probability but video providers are not physically close to video requesters, the users in NetTube experience the longest freezing delay.

Fig. 24 shows the CDF of the percent of server contribution of each node with different client population in SocialTube. It shows that in the system with 268 nodes, 80 percent of the nodes require no more than 40 percent of the traffic from the server, and in the system with 200 and 100 nodes respectively, 80 percent of the nodes require no more than 50 percent and 56 percent of the traffic from the server, respectively. As the system has more nodes, more nodes can share videos after they watch the videos. Then, the total uploading capacity of the P2P transmission increases, which reduces the burden on the server. This result implies that P2P structure is an optimal architecture for a large-scale video sharing system.

6 CONCLUSION

Video sharing is an increasingly popular application in OSNs. However, the client/server architecture deployed by current video sharing systems in OSNs costs a large amount of resources (i.e., money, server storage) for the service provider and lacks scalability. Meanwhile, because of the privacy constraints in OSNs, the current peer-assisted Video-on-Demand (VoD) techniques are suboptimal if not entirely

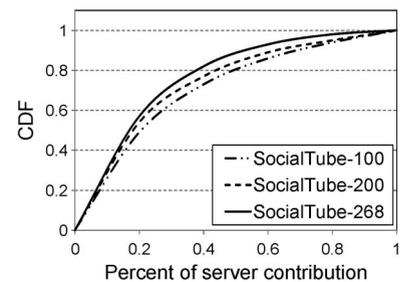


Fig. 24. CDF of the percent of server contribution in SocialTube.

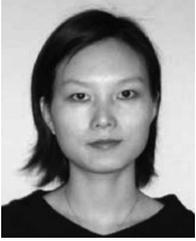
applicable to the video sharing in OSNs. In this paper, we crawled video watching trace data in one of the largest OSN websites Facebook, from July 2007 to Aug. 2010 and explored the users' video viewing patterns. We found that in a user's viewer group, 25 percent viewers watched all videos of the user driven by social relationship, and the viewing pattern of the remaining nodes is driven by interest. Based on the observed social and interest relationship in video watching activities, we propose SocialTube, which provides efficient P2P-assisted video sharing services. Extensive simulation results show that SocialTube can provide a low video startup delay and low server traffic demand. We also implemented a prototype in PlanetLab to evaluate the performance of SocialTube. The experimental results from the prototype further confirm the efficiency of SocialTube.

ACKNOWLEDGMENT

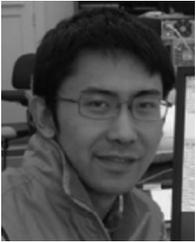
This research was supported in part by U.S. NSF Grants CNS-1254006, CNS-1249603, CNS-1049947, CNS-1156875, CNS-0917056 and CNS-1057530, CNS-1025652, CNS-0938189, CSR-2008826, CSR-2008827, Microsoft Research Faculty Fellowship 8300751, and U.S. Department of Energy's Oak Ridge National Laboratory including the Extreme Scale Systems Center located at ORNL and DoD 4000111689. An early version of this work was presented in the Proceedings of Infocom-mini'12 [22].

REFERENCES

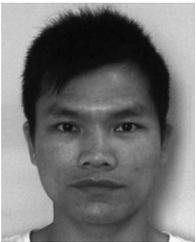
- [1] Facebook Passes Google in Time Spent on Site for First Time Ever. [Online]. Available: <http://www.businessinsider.com/Social-Media-Web-2.0-and-Internet-Stats>. [Online]. Available: <http://thefuturebuzz.com/2009/01/12/social-media-web-20-internet-numbers-sta>
- [2] K. Wang and C. Lin, "Insight into the P2P-VoD system: Performance Modeling and Analysis," in *Proc. ICCCN*, 2009, pp. 1-6.
- [3] Y. Huang, Z. Fu, D. Chiu, C. Lui, and C. Huang, "Challenges, Design and Analysis of a Large-Scale P2P VoD System," in *Proc. SIGCOMM*, 2008, pp. 375-388.
- [4] B. Cheng, L. Stein, H. Jin, X. Liao, and Z. Zhang, "Gridcast: Improving Peer Sharing for P2P VoD," *ACM Trans. Multimedia Comput., Commun., Appl.*, vol. 4, no. 4, p. 26, Oct. 2008.
- [5] C.-P. Ho, S.-Y. Lee, and J.-Y. Yu, "Cluster-Based Replication for P2P-Based Video-on-Demand Service," in *Proc. ICEIE*, 2010, pp. V2-49-V2-53.
- [6] J. Wang, C. Huang, and J. Li, "On ISP-Friendly Rate Allocation for Peer-Assisted VoD," in *Proc. MM*, 2008, pp. 279-288.
- [7] C. Huang, J. Li, and K.W. Ross, "Can Internet Video-on-Demand be Profitable?" in *Proc. SIGCOMM*, 2007, pp. 133-144.
- [8] X. Cheng and J. Liu, "NetTube: Exploring Social Networks for Peer-to-Peer Short Video Sharing," in *Proc. IEEE INFOCOM*, 2009, pp. 1152-1160.
- [9] B. Li, M. Ma, Z. Jin, and D. Zhao, "Investigation of a Large-Scale P2P VoD Overlay Network by Measurements," *Peer-to-Peer Netw. Appl.*, vol. 5, no. 4, pp. 398-411, Dec. 2012.
- [10] C. Wu, Z. Li, X. Qiu, and F. Lau, "Auction-Based P2P VoD Streaming: Incentives and Optimal Scheduling," *ACM Trans. Multimedia Comput., Commun., Appl.*, vol. 8, no. S1, p. 14, Feb. 2012.
- [11] T. Fujimoto, R. Endo, and H. Shigeno, "P2P Video-on-Demand Streaming using Caching and Reservation Scheme Based on Video Popularity," *Int. J. Grid Utility Comput.*, vol. 3, no. 2/3, pp. 188-199, 2012.
- [12] W. Wu and J. Lui, "Exploring the Optimal Replication Strategy in P2P-VoD Systems: Characterization and Evaluation," *IEEE Trans. Parallel Distrib. Syst.*, vol. 23, no. 8, pp. 1492-1503, Aug. 2012.
- [13] Y. Zhou, T. Fu, and D. Chiu, "A Unifying Model and Analysis of P2P VoD Replication and Scheduling," in *Proc. IEEE INFOCOM*, 2012, pp. 1530-1538.
- [14] D. Niu, Z. Liu, B. Li, and S. Zhao, "Demand Forecast and Performance Prediction in Peer-Assisted on-Demand Streaming Systems," in *Proc. IEEE INFOCOM*, 2011, pp. 421-425.
- [15] L. Chang and J. Pan, "Towards the Optimal Caching Strategies of Peer-Assisted VoD Systems with HD Channels," in *Proc. IEEE ICNP*, 2012, pp. 1-10.
- [16] Z. Wang, L. Sun, S. Yang, and W. Zhu, "Prefetching Strategy in Peer-Assisted Social Video Streaming," in *Proc. MM*, 2011, pp. 1233-1236.
- [17] A. Barrios, M. Barrios, D. Vera, P. Bocca, and C. Rostagnol, "GoalBit: A Free and Open Source Peer-to-Peer Streaming Network," in *Proc. ACM Multimedia*, 2011, pp. 727-730.
- [18] W.P.K. Yiu, X. Jin, and S.H.G. Chan, "VMesh: Distributed Segment Storage for Peer-to-Peer Interactive Video Streaming," *IEEE J. Sel. Areas Commun.*, vol. 25, no. 9, pp. 1717-1731, Dec. 2007.
- [19] H. Shen, L. Zhao, Z. Li, and J. Li, "A DHT-Aided Chunk-Driven Overlay for Scalable and Efficient P2P Live Streaming," in *Proc. ICPP*, 2010, pp. 248-257.
- [20] H. Shen, L. Zhao, H. Chandler, J. Stokes, and J. Li, "P2P-Based Multimedia Sharing in User Generated Contents," in *Proc. IEEE INFOCOM*, 2011, pp. 667-675.
- [21] Z. Li, H. Shen, H. Wang, G. Liu, and J. Li, "SocialTube: P2P-Assisted Video Sharing in Online Social Networks," in *Proc. IEEE INFOCOM*, 2012, pp. 2886-2890.
- [22] M. Gjoka, M. Kurant, C.T. Butts, and A. Markopoulou, "Walking in Facebook: A Case Study of Unbiased Sampling of Osns," in *Proc. IEEE INFOCOM*, 2010, pp. 1-9.
- [23] M. Wittie, V. Pejovic, L. Deek, K. Almeroth, and B. Zhao, "Exploiting Locality of Interest in Online Social Networks," in *Proc. Co-NEXT*, 2010, p. 25.
- [24] H. Wang, J. Liu, and K. Xu, "On the Locality of Bittorrent-Based Video File Swarming," in *Proc. P2P*, 2009, p. 12.
- [25] Peersim: A Peer-to-Peer Simulator. [Online]. Available: <http://peersim.sourceforge.net/>
- [26] PlanetLab. [Online]. Available: <http://www.planet-lab.org/>
- [27] A. Nazir, S. Raza, and C. Chuah, "Unveiling Facebook: A Measurement Study of Social Network Based Applications," in *Proc. SIGCOMM*, 2008, pp. 43-56.
- [28] L. Backstrom, E. Sun, and C. Marlow, "Find Me if You Can: Improving Geographical Prediction with Social and Spatial Proximity," in *Proc. WWW*, 2010, pp. 61-70.
- [29] Facebook Users Average 7 hrs a Month in January as Digital Universe Expands. [Online]. Available: <http://blog.nielsen.com/nielsenwire>
- [30] L. Zhou, Y. Zhang, K. Song, W. Jing, and V. Vasilakos, "Distributed Media Services in P2P-Based Vehicular Networks," *IEEE Trans. Veh. Technol.*, vol. 60, no. 2, pp. 692-703, Feb. 2011.
- [31] P. Salvador and A. Nogueira, "Study on Geographical Distribution and Availability of Bittorrent Peers Sharing Video Files," in *Proc. IEEE ISCE*, 2008, pp. 1-4.
- [32] T. Nair and P. Jayarekha, "Strategic Prefetching of VoD Programs Based on ART2 Driven Request Clustering," *Int. J. Inf.*, vol. 1, no. 2, pp. 13-21, Sept. 2011.
- [33] D. Wu, C. Liang, Y. Liu, and K. Ross, "View-Upload Decoupling: A Redesign of Multi-Channel P2P Video Systems," in *Proc. IEEE INFOCOM*, 2009, pp. 2726-2730.
- [34] T. Silva, J. Almeida, and D. Guedes, "Analysis of P2P Streaming Based on the Characterization of Live User-Generated Video," in *Proc. IEEE ICC*, 2010, pp. 1-5.
- [35] C. Wu, B. Li, and S. Zhao, "Multi-Channel Live P2P Streaming: Refocusing on Servers," in *Proc. IEEE INFOCOM*, 2008, pp. 1355-1363.
- [36] Y.-Y. Ahn, S. Han, H. Kwak, S. Moon, and H. Jeong, "Analysis of Topological Characteristics of Huge Online Social Networking Services," in *Proc. WWW*, 2007, pp. 835-844.
- [37] J. Liu, X. Cheng, and C. Dale, "Statistics and Social Network of YouTube Videos," in *Proc. IEEE IWQoS*, 2008, pp. 229-238.
- [38] H. Yu, D. Zheng, B.Y. Zhao, and W. Zheng, "Understanding User Behavior in Large-Scale Video-on-Demand Systems," *SIGOPS Oper. Syst. Rev.*, vol. 40, no. 4, pp. 333-344, Oct. 2006.
- [39] J.A. Pouwelse, P. Garbacki, J. Wang, A. Bakker, J. Yang, A. Iosup, D. Epema, M. Reinders, M.R. Van, and H.J. Sips, "Tribler: A Social-Based Peer-to-Peer System," *Concurr. Comput., Pract. Exp.*, vol. 20, no. 2, pp. 127-138, Feb. 2008.
- [40] D.N. Kalofonos and Z. Antoniou, "A Hybrid P2P/Infrastructure Platform for Personal and Social Internet Services," in *Proc. IEEE PIMRC*, 2008, pp. 1-6.



Haiying Shen received the BS degree in computer science and engineering from Tongji University, Shanghai, China, in 2000, and the MS and PhD degrees in computer engineering from Wayne State University, Detroit, MI, USA in 2004 and 2006, respectively. Currently, she is an Assistant Professor in the Department of Electrical and Computer Engineering at Clemson University, Clemson, SC, USA. Her research interests include distributed computer systems and computer networks, with an emphasis on P2P and content delivery networks, mobile computing, wireless sensor networks, and cloud computing. She is a Microsoft Faculty Fellow of 2010 and a Senior Member of the IEEE and ACM.



Ze Li received the BS degree in electronics and information engineering from Huazhong University of Science and Technology, Wuhan, China, in 2007, and the PhD degree in computer engineering from Clemson University, Clemson, SC, USA. His research interests include distributed networks, with an emphasis on peer-to-peer and content delivery networks. Currently, he is a Data Scientist in the MicroStrategy Incorporation, Towers Crescent Plaza Tysons Corner, VA, USA. He is a Student Member of the IEEE.



Yuhua Lin received both the BS and MS degrees in computer engineering from Sun Yat-sen University, Guangdong, China, in 2009 and 2012, respectively. Currently, he is a PhD student in the Department of Electrical and Computer Engineering of Clemson University, Clemson, SC, USA. His research interests include social networks and reputation systems.



Jin Li received the PhD degree (with honor) from Tsinghua University, Beijing, China, in 1994. He is a Research Manager and Principal Researcher at Microsoft Research, Redmond, WA, USA. He manages the Compression, Communication, and Storage group. Blending theory and system, he excels at interdisciplinary research and is dedicated to advance communication and information theory and apply it to practical system building. After brief stints at USC and Sharp Labs, he joined Microsoft Research in 1999, first as one of the founding members of Microsoft Research Asia and then moved to Microsoft Research (Redmond, WA) in 2001. From 2000, Dr. Li has also served as an Affiliated Professor in Tsinghua University. His invention has been integrated into many Microsoft products. Recently, he and his group members have made key contributions to Microsoft product line (e.g., RemoteFX for WAN in Windows 8, Primary Data Deduplication in Windows Server 2012, and Local Reconstruction Coding in Windows Azure Storage), which leads to commercial impact in the order of hundreds of millions of dollars. Dr. Li was the recipient of the Young Investigator Award from Visual Communication and Image Processing'98 (VCIP) in 1998, ICME 2009 Best Paper Award, and USENIX ATC 2012 Best Paper Award. He is/was the Associate Editor/Guest Editor of IEEE TRANSACTIONS ON MULTIMEDIA, *Journal of Selected Area of Communication*, *Journal of Visual Communication and Image Representation*, P2P networking and applications, *Journal of Communications*. He was the General Chair of PV2009, the Lead Program Chair of ICME 2011, and the Technical Program Chair of CCNC 2013. He is a Fellow of the IEEE.

▷ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.