Stefan Wilk, Julius Rückert, Timo Thräm, Christian Koch, Wolfgang Effelsberg and David Hausheer
The Potential of Social-aware Multimedia Prefetching on Mobile Devices
In Proc. of the IEEE Networked Systems (NetSys), 2015, ISBN: 978-1-4799-5804-7

# The Potential of Social-aware Multimedia Prefetching on Mobile Devices

Stefan Wilk*, Julius Rückert†, Timo Thräm, Christian Koch†, Wolfgang Effelsberg* and David Hausheer†
* Distributed Multimedia Systems, Technische Universität Darmstadt
Email: {swilk|effelsberg}@cs.tu-darmstadt.de
† Peer-to-Peer Systems Engineering Lab, Technische Universität Darmstadt
Email: {rueckert|ckoch|hausheer}@ps.tu-darmstadt.de

***Abstract*—The access to Online Social Networks (OSN) and to media shared over these platforms account for around 20% of today's mobile Internet traffic. For mobile device users, the access to media content and specifically videos is still challenging and costly. Mobile contracts usually have a data cap and connection qualities can vary greatly, depending on the cellular network coverage. Prefetching mechanisms that fetch content items beforehand, in times when the mobile device is connected to a WiFi network, have a high potential to address these problems. Yet, such a mechanism can only be effective if relevant content can be predicted with a high accuracy. Therefore, in this paper, an analysis of content properties and their potential for prediction are presented. An initial user study with 14 Facebook users running an app on their mobile device was conducted. The results show that video consumption is very diverse across the users. This work discusses the evaluation setup, the data analysis, and their potential to define an effective prefetching algorithm.**

## I. Introduction

Online Social Networks (OSN) are being increasingly accessed from mobile devices. Recent studies show that OSNs accounts for roughly 20% of the overall traffic on mobile devices in Europe as well as North America [9]. A big part of the shared content is multimedia such as photos or videos. Today, OSNs act as Internet-based communication and content sharing platforms that became a central element of many people's daily life.

To connect to the Internet, mobile devices typically either rely on WiFi or cellular networks. Hereby, WiFi connections are preferred and, thus, used by the devices whenever they are available. In cases without WiFi connectivity, cellular network technologies like UMTS or LTE are used. Using these cellular networks usually has several drawbacks: First of all, typical mobile contracts include a data cap that limits the monthly data volume to be transferred. For WiFi-based access this is usually not the case. Secondly, as the network coverage of cellular technologies highly differs depending on the network provider and location, the quality and speed of connections can be very limited and fluctuating. As a result, the access to media content and specifically to videos can be challenging using cellular networks. This may lead to long startup delays and recurring playback stalls as some video segments are not delivered in time. To address these problems, the use of prefetching mechanisms has been studied for a while now. Here, the idea is that content is downloaded prior to its consumption in a WiFi network, e.g. at night when the mobile device is connected to the user's home network.

While the idea of prefetching seems straightforward, realizing an effective prefetching mechanism is challenging as it requires precise predictions on which content items are likely to be consumed in future. Content items that an OSN presents to a user, e.g. on its individual newsfeed, can be used as candidates that, with a high probability, are accessed by the user in near future. Thus, they are promising to be considered for prefetching. Yet, some of the items might be more interesting for the user than others and, consequently, also accessed more or less likely. Facebook offers a rich set of meta data, that might be helpful to derive the relevance of an item for the user. Properties, e.g. related to the social graph or the number of likes and comments affect this relevance.

To this end, this paper presents an analysis of content properties on Facebook on their potential to be used for user-based content prediction. An initial pilot study including 14 subjects has been conducted, spanning several weeks. All subjects agreed on using a specialized app on their mobile devices that mimics the functionality of the native Facebook app but allows collecting media- and OSN-related meta data in a privacy-conserving way. The results indicate that mobile video consumption on Facebook is very diverse and that general prediction rules are hard to obtain. Yet, a limited number of attributes exists that could very well help to design effective video prefetching algorithms for mobile devices. The insights provided by this work can help to better understand the potential of content prediction for mobile prefetching based on OSN information.

## II. Background and Related Work

Media prefetching allows to shift data traffic from cellular networks to WiFi connections. Latencies in cellular networks and lack of coverage in rural areas may dramatically increase the loading time of videos. Krishnan et al. [6] show that for short video clips, which are predominately shared in OSNs a startup delay of two seconds already infers that many users quit video streaming. Effective prefetching can decrease the video startup time which enhances the quality of a user in a video streaming session.

Different approaches for leveraging OSN information to improve media access have been proposed. Kaafar et al. [4] present a recommendation-aware content placement strategy for Content Delivery Networks. Bai et al. [1] investigates OSN-related caching mechanisms for Yahoo and Facebook. These approaches show sources for designing social-aware prefetching strategies, but aim at improving caching and the delivery to thousands of users by placing the right information at the right place within the network, whereas prefetching concentrates on loading content to a device.

Prefetching in peer-to-peer (P2P) video streaming systems is a well-researched topic. The work of Wang et al. [10] investigates RenRen, a popular OSN used in China, which works similar to Facebook, on its suitability for P2P-assisted video streaming. Prefetching of the first video chunks helps in this domain to reduce the video startup time. Their results indicate that social relationships and users' preferences are suitable predictors for OSNs similar to RenRen.

Li et al. [7] add insights on video dissemination in OSNs and identify very short life times of videos with low popularity. The resulting SocialTube [8] system demonstrates the efficiency of a peer-to-peer-based social network that is driven by social relationships and less by interests. In their analysis they show that 0.4% of the videos cause more than 80% of the traffic and the remaining 99.6% only cause 20% of the network traffic. But most, around 90% of the video views of a user can already be explained by shares of their 1- and 2-hop friends.Around 33% of the users watch 80% of videos and all viewers watch at least 20% of the videos on their newsfeed.

All of these studies and prefetching strategies do not consider the effect of mobile devices or behavior of mobile users. With the work by Gautam et al. [3] a mobile application prefetches whole video clips based on arbitrary sources such as OSNs or blogs. The goal of the work is reducing the energy consumption of a mobile device, but prefetching algorithms are not described in detail. The most promising approach is presented by Zhao et al. [2] in which a custom mobile Facebook application has been developed that integrates social network based algorithms for prefetching. They conduct a study on Facebook, similar to the work presented in this paper. The findings are limited and have to be supported by simulation experiments. The results indicate a significant energy savings, if the number of videos shared in OSNs would be higher. Yet, the designed application, their sorting of the posts, mapping of the real Facebook newsfeed and privacy-awareness is not described in detail. The work presented here addresses these issues by a detailed explanation of the mobile application, its suitability to preserve the privacy of users, a detailed description of the framework, and a study including the results and giving an outlook on future sources for media prediction in OSNs.

## III. Mobile App System Design

To gain valuable insights to media content consumption in OSNs, we conducted a user study with 14 participants. Each participant was asked to use an Android app, called SonNet, on his own mobile device(s). Additionally, the participants had to sign a privacy consent before the activation key for the app was handed out. The app's main functions are tracking of the posts shown on the user's Facebook feed, keeping track if the user watches the posts, and uploading the collected information to a database server for later analysis. The information is anonymized before uploading to ensure the privacy of the user. This was done by hashing all personal data-related fields with a hash algorithm (SHA-1). This way, e.g., multiple datasets by the same user could be related, without revealing any information on the user identity itself.

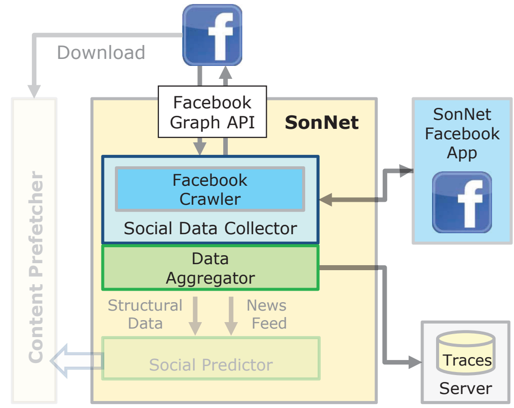


Fig. 1: SonNet Architecture

The architecture of SonNet is shown in Fig. 1. SonNet collects information about all posts on the user's feed. This information was retrieved using the Facebook Graph API, located in the *Facebook Crawler*, which is embedded in the *Social Data Collector* component. As the aim is to identify posts that the user is most interested in, a set of post features are collected as well. To this end, the *Facebook Crawler* fetches the number of likes and comments of the individual posts. Furthermore, the source of the post is identified to tell if direct friends of the user or friends of friends published or interacted with the posts. Additionally, the information if the post is created in a group the user is assigned to, is tracked.

The participants of the study access Facebook with the *SonNet Facebook App* instead of the native Facebook client. This allows to precisely keep track of interactions the user makes with the posts, e.g. watching, liking, or commenting. The posts presented to the user are the ones which were previously retrieved by the *Facebook Crawler* using the Facebook Graph API. This information is collected and filtered by the *Social Data Collector* and accessible from the *SonNet Facebook App*. The *Social Data Collector* periodically queries for new posts on the user's newsfeed. It is important to note that some limitations regarding the collection exists. First, the collected posts do not include any sponsored posts as they would appear in the native Facebook app. Second, the order of posts as retrieved by the *Facebook Crawler* is not necessarily the same as the one Facebook presents in the native app. As the sorting algorithm used by Facebook itself is not publicly available, the *SonNet Facebook App* follows a simplified approach and presents the posts ordered by creation time.

The *Social Data Collector* passes the received posts and user interaction information to the *Data Aggregator*, which filters and locally stores the data. The anonymization of all personal data-related fields takes place directly before storing the data. In a later step, the local database is transferred to the database server, where the datasets of all participants are stored.

In this work, the main focus is on understanding which features have an influence on the relevance of a post for a user and, therefore, might be predictive. Based on the results, it is planned to develop a *Social Predictor* component, which is ongoing work and was not part of the mobile app as used

| | Mean | Median | Stdev. | Min | Max |
|---|---|---|---|---|---|
| # Friends | 217.5 | 160 | 182.5 | 11 | 559 |
| # Posts | 152.71 | 61.5 | 174.41 | 18 | 631 |
| # Posts Watched | 15 | 3 | 174.4 | 0 | 90 |
| # Friend's list | 72.28 | 77 | 57.48 | 1 | 170 |
| # Groups | 31.29 | 13 | 46.71 | 0 | 180 |
| # Interests | 10.78 | 6.5 | 10.45 | 0 | 31 |

TABLE I: Overview on the mobile prefetching dataset.

during the study. This predictor would take the user's newsfeed as an input to train a machine learning model. The model is planned to be used to identify relevant newsfeed posts and use videos as well as images contained in these posts as prefetch candidates. These media files are then planned to be downloaded by the *Content Prefetcher* component.

## IV. Description of a Dataset for Prefetching

This pilot study contains the data gathered from mobile usage of Facebook by 14 participants over a time span of eight weeks. The data includes interactions of users with the posts, especially with those containing photo and video as well as social graph information such as social closeness of users or their interests. In total 2138 posts including 202 video posts have been gathered. The category video posts contains Facebook hosted videos, links to YouTube videos or other video sharing sites as well as the category 'flash' video.

The data gathered has to map real usage of the OSN as close as possible. Thus, all participants were motivated not to change their Facebook usage behavior. All users were asked to use their smart phone. Tablet devices or laptops - despite their mobility aspects - are not included in this study. No users or posts were filtered and no data clearance steps have been taken. In our dataset 42% of all posts include photos, 20% share links, 29% are status updates and videos account for 9% of the dataset. Table I lists statistical numbers on the gathered dataset. To ensure privacy and comply with national privacy laws, all data has been anonymized using a hashing algorithm on the mobile device.

## V. Analysis of the Mobile Facebook Dataset

The newsfeed within the OSN Facebook has been analyzed according to the questions: (1) How much media – especially video – that is being distributed in the OSN Facebook is watched by the participants of the study? (2) Can the global popularity of a video be used to tell, whether one of the participants is going to consume media? (3) Is the origin of a post a suitable predictor for media consumption?

The *first question* elaborates on the proportion of photo as well as video posts being consumed by the mobile users. As videos account for only 9% of the posts observed in our study it is questionable if Facebook is video-centric. Photos, in contrast, account for a large proportion of Facebook posts. Furthermore, an analysis of the share of watched versus non-watched content gives a first indication on the potential of media prefetching strategies for mobile devices.

The *second question* investigates the origin of a post. One hypothesis of this work is that the origin of a video or photo post significantly influences its consumption probability. To be more specific, this work postulates that videos shared by friends are more likely to be watched in comparison to videos from other sources.

The *third question* addresses likes, comments, and shares associated with a post. It investigates whether those interactions with a post have a significant influence on the probability of a single video being consumed.

### A. *Mobile media consumption & influence of content type*

The first research question addresses the specifics of the gathered mobile Facebook media consumption dataset. For the consumption probability touch/click events for photos or videos are monitored. For photos, the touch event maximizes the photo to full-screen size, for videos it starts the playback. In total, the average consumption rate of posts is 9.27% (Stdev.: 14.72; Median: 0). Videos exhibit in average a larger probability for consumption per user (Average: 19.15%; Stdev.: 28.5; Median: 0) in comparison to photos (Average: 7.75%; Stdev.: 14.53; Median: 0). This may have different reasons: First of all, the total number of videos is significantly lower in comparison to photos. In this study, nearly half of the posts belong to the category *photo*. The high quantity of photos may influence consumption negatively, whereas the exceptional occurrence of a video on the newsfeed could motivate users to watch a video. Additionally, photos, in comparison to videos, are easier to be consumed directly from the newsfeed without any interaction. Small versions of photos, so called thumbnail versions, are presented on the newsfeed as preview. Videos have a temporal dimension that can not be obtained by solely watching at the thumbnail.

The consumption of media is very diverse, as the median of 0% and the standard deviation (photo: 14.53; video: 28.5) for both video and photo consumption show. Especially for videos, a group of users within this study did not watch content at all, whereas some other participants consumed a large share of the videos presented. Users can be divided into non-media consuming and media-consuming users. Another idea focuses on the friend's interests. A comparison of the interests of users participating in this study and the interests of their friends is done. Due to the anonymization of data transferred to our servers, only the interest categories could be compared. To better understand the influence of common interests. Both the friends' interests who shared posts and the users' interests are compared for each individual post. The Euclidian distance between the number of interests in the categories, i.e. music, tv, movies, books, and games is calculated. For both photos as well as videos, a lower distance indicates that user and friend have more interests in common. For all posts, in mean a distance of 17.9 is calculated for watched, in comparison to 21.4 for non-watched content. This already indicates that interests could play a significant role for predicting multimedia consumption. This difference shall be calculated for the high consuming group of users. A reduced distance can be observed with an increased difference between the means of watched and non-watched videos. In contrast to our expectations, it is not a valid classifier for high consuming users, but still valid for those watching videos or photos only seldom.

As a result, the separation of users according to specific groups based on their behavior in the OSN is a promising approach that has to be evaluated in detail. It promises answers to the question on how to predict media consumption in OSNs.

### B. Origin of a post

The origin of posts is grouped into 'shared by a friend' and 'shared by other sources'. Shared in this context refers to any interaction of the friend that resulted in showing the post on the user's newsfeed. E.g. the creation of a post, commenting or liking a post may bring it to the newsfeed.

The majority of all posts are shared by friends and only 32.45% of all posts stem from other sources. Specifically, this rate was significant higher for videos with 44.56% and 40.9% for photos compared to other media. Thus, in comparison to other post types photos and videos are increasingly being distributed by non-friends. For videos, no difference is observable between posts sourced by friends (10.32%) and other sources (10.97%), such as pages or recommended posts on the Facebook newsfeed. There is no difference resulting from the source of a video or a photo on the consumption probability, if solely the average is evaluated. As the observed data suffers from significant variance, median offers a better insight. For videos a higher median for the proportion of watched videos can be observed for friend posts.

Content in our study is consumed shortly after publication and due to low numbers of likes and comments of most posts, we believe they are distributed in only small groups of friends. This can be supported by Table II, which investigates a subset of posts that has been shared by users within a friend's list. For both, videos and for photos, posts by close friends and family are preferred. The information on the social distance to a posting user could, thus, easily be identified and leveraged for prefetching mechanisms. The videos shared by close friends or family are predominantly those with low numbers of likes or comments. None of the videos re-shared by these friend list members has more than one thousand or more likes. This illustrates that videos watched from friends do not have the necessity to be very popular in an OSN.

TABLE II: Friend lists and their impact on consuming video and photos. The table shows the percentage of media shared by members of a friend list that is being consumed.

| | Photos | Videos |
|---|---|---|
| Close Friends | 70.1% | 85.7% |
| Family | 82.6% | 50% |
| Other lists | 9.2% | 8.3% |

### C. Global popularity

Previous studies [7] show a significant impact of popularity of content on its consumption. Content on video sharing sites follows a so called zipf-like distribution in which a large proportion of the videos is watched only by a very small amount of users, whereas the top percent has high access rates. Popularity in the OSN Facebook is represented by the total number of interactions with the post. Interactions evaluated are likes, comments, as well as shares of posts. The count of likes and comments is limited to 1,000 for presentation reasons. While this is a limitation that was removed for future studies, it only affects 1.1% of the comment and 6.1% of the like counts on posts in the dataset as they exceed this threshold. In the dataset, it can be observed that only a small share of posts reaches high numbers of likes, comments or shares.

Figure 2 shows the effect of likes and comments on the consumption of videos. No significant difference can be observed for the watched and non-watched videos influenced by likes and comments. Only 10% of the multimedia posts have high numbers of shares. This indicates that globally popular content which is being shared from user to user accounts only for a very small proportion of all posts.

Regarding photos, the median of the number of likes is 102 and 10 for comments. A difference between the values of watched and non-watched posts can be observed for photos. No similar pattern exists for the video posts. The median number of likes with 5 (Stdev.: 339) for watched videos and 53 (Stdev.: 402) for non-watched videos indicate that especially videos with a low number of interactions are of interest to a user. Figure 2 shows the small impact of popularity measured by likes, comments and shares for a large share (more than 80%) of the videos. For the remaining 20%, popularity has an impact, as the CDF for watched videos is significantly less increasing until around 900 comments. For this small proportion of posts a high number of shares gives additional potential for prediction. All video posts with more than 10,000 shares are watched, indicating that globally relevant content, even though seldom available, is consumed by users of our dataset. Interestingly, the CDF on the effect of increasing number of likes on the video consumption rate indicates that with increasing global popularity (likes) the consumption rate of a video decreases. Multiple explanations for the findings exist. First of all, videos being consumed are mainly shared within small groups and thus range. When creating a post on Facebook, a user has the opportunity to restrict the visibility of the post to himself/herself only, a specific group of friends, managed in friend's lists or public. Another explanation would be that the users very much consume videos shortly after being published. This theory can be supported by findings of Li et al. [7]. The researchers reported, that the average lifetime of a video in OSNs is very limited. As a result for this study, one finding is that for the majority of posts available on the newsfeed, the popularity expressed by like, comment, and share numbers is not affecting the consumption probability.

To better understand the results gathered, we compare those findings to the CDFs of photos shared on Facebook. The result can be obtained from Figure 3. It shows a stronger influence of high number of likes and comments on the probability of a photo being enlarged. Still, new content or available content not able to reach large numbers of shares and likes exhibit the strongest probability to be watched. In addition, the dataset shows that an approximation between likes, shares, and comments can be made based on having only a single information source. Especially for videos, the Pearson's correlation between shares and comments reaches 0.9263, whereas it drops for pictures to 0.17864. In contrast to that, the correlation between comments and likes for videos reaches 0.6287 and around 0.79986 for pictures. Thus, increasing numbers of comments allow to predict rising numbers of shares and comments.

Popularity is a highly accurate prediction mechanism for caching, but at least for this study only a limited impact on

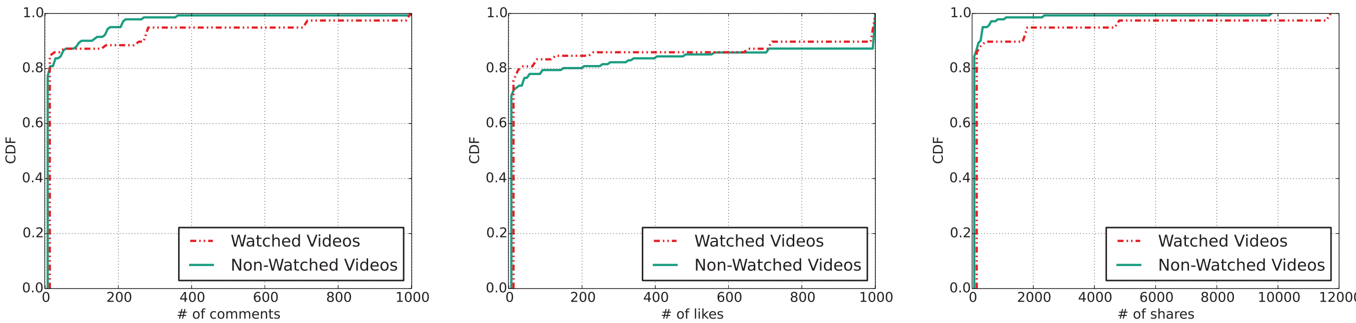


Fig. 2: Effect of number of (Left) comments, (Center) likes, and (Right) shares on the consumption rate of videos.

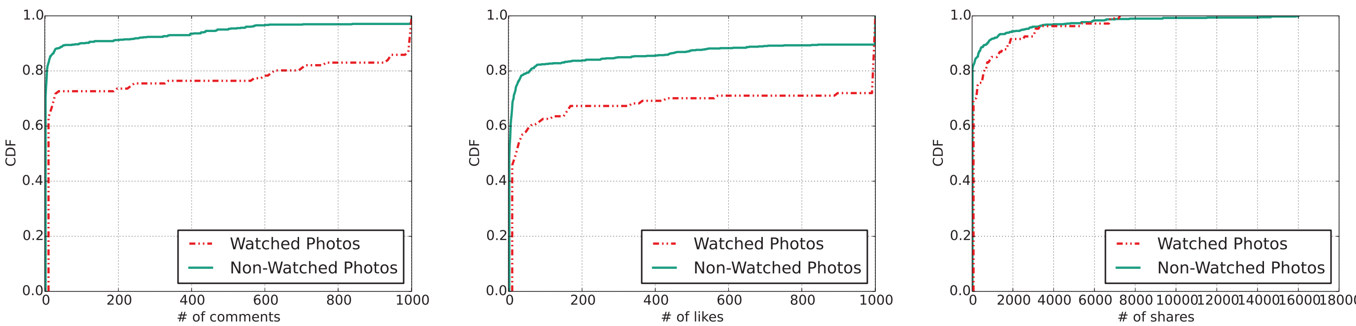


Fig. 3: Effect of number of (Left) comments, (Center) likes, and (Right) shares on the consumption rate of photos.

prefetching precision is observable. The consumption, especially of video, seems only to a very small portion being affected by the number of likes and comments.

## VI. Conclusion

In this work, an analysis on features available in the OSN Facebook has been conducted in order to test them according to their suitability to predict media consumption. The results indicate that multimedia content is increasing on Facebook, as more than half of all captured posts are images or videos. From this pilot study with only a limited number of users, no single predictor explaining media consumption can be retrieved. Prefetching mechanisms solely based on the pure number of likes and comments are not adequate for predicting whether a video is going to be watched or not. Likes and comments may predict quite well the popularity but most of the content is watched when it has only a local popularity. Possible prediction sources can be the social closeness. While social closeness is mostly measured by hops in the social graph, this study gives insights that being a friend alone is not enough. Friend lists seem to be, at least for those users who use them, one abstraction layer that increase prediction precision. The findings gained have to be validated in the upcoming large-scale study. The research challenges as well as the design of this large-scale user study is described in the work of Koch et al. [5].

## Acknowledgment

This work has been funded in parts by the European Union (FP7/#318398, eCOUSIN) and the German Research Foundation (DFG) as part of project C03 within the Collaborative Research Center (CRC) 1053 – MAKI.

## References

[1] X. Bai, F. Junqueira, and A. Silberstein, “Cache Refreshing for Online Social News Feeds,” in *ACM International Conference on Information and Knowledge Management (CIKM)*, 2013.

[2] N. Do, Y. Zhao, S.-T. Wang, C.-H. Hsu, and N. Venkatasubramanian, “Optimizing Offline Access to Social Network Content on Mobile Devices,” in *IEEE INFOCOM*, 2014.

[3] N. Gautam, H. Petander, and J. Noel, “A Comparison of the Cost and Energy Efficiency of Prefetching and Streaming of Mobile Video,” in *ACM Workshop on Mobile Video (MoVid)*, 2013.

[4] M. Kaafar, S. Berkovsky, and B. Donnet, “On the Potential of Recommendation Technologies for Efficient Content Delivery Networks,” *SIGCOMM Comput. Commun. Rev.*, vol. 43, no. 3, pp. 74–77, 2013.

[5] C. Koch and D. Hausheer, “Optimizing mobile prefetching by leveraging usage patterns and social information,” *IEEE International Conference on Network Protocols (ICNP 2014)*, 2014.

[6] S. S. Krishnan and R. K. Sitaraman, “Video Stream Quality Impacts Viewer Behavior: Inferring Causality Using Quasi-experimental Designs,” in *ACM Conference on Internet Measurement Conference (IMC)*, 2012.

[7] H. Li, H. Wang, J. Liu, and K. Xu, “Video Requests from Online Social Networks: Characterization, Analysis and Generation,” in *IEEE INFOCOM*, 2013.

[8] Z. Li, H. Shen, H. Wang, G. Liu, and J. Li, “SocialTube: P2P-assisted Video Sharing in Online Social Networks,” in *IEEE INFOCOM*, 2012.

[9] Sandvine, “Spring 2014 Global Internet Phenomena Report,” 2014.

[10] Z. Wang, L. Sun, and S. Yang, “Prefetching Strategy in Peer-Assisted Social Video Streaming,” in *ACM Conference on Multimedia (MM)*, 2011.