Systematic, Large-scale Analysis on the Feasibility of Media Prefetching in Online Social Networks

Thomas Paul*, Daniel Puscher*, Stefan Wilk†, and Thorsten Strufe‡

* P2P Networks, Technical University of Darmstadt, Germany
{thomas.paul@cs.|puscher@stud.} tu-darmstadt.de

† Distributed Multimedia Systems, Technical University of Darmstadt, Germany
swilk@cs.tu-darmstadt.de

‡ Datenschutz und Datensicherheit, Technical University of Dresden, Germany
thorsten.strufe@cs.tu-dresden.de

***Abstract*—Huge quantities of videos are shared via Online Social Networks (OSN) like Facebook and are watched on mobile devices. Internet connections via cellular networks (UMTS / LTE) require the scarce resources radio bandwidth and battery power. Prefetching of videos in areas of WLAN availability has the potential to reduce the power consumption in comparison to data transmission via cellular networks and prefetching can help to avoid users running into traffic caps of their network providers. Furthermore, startup delays can be reduced.**

Social networks offer contextual information such as likes and comments as well as social graph information which can potentially be used to predict which content will be consumed in the near future. In this paper, we elaborate possibilities to predict content consumption based on the number of likes, comments and the social graph distance. Our detailed analysis of the media access patterns of more than 700 users in Facebook shows that the media consumption does not solely depend on the number of likes or comments. Users tend to watch videos that are uploaded by close friends and family members. Furthermore, the time a video preview stays in the browser-viewport before being clicked (pre-click delay) can be exploited to decrease startup delays.

I. Introduction

Online Social Networks (OSN) are Internet-based communication and content sharing platforms. Users consume and produce content, like videos and photos, and share them amongst each other. The trend to access the Internet via mobile end-user devices shifts social network usage away from wired connections to wireless environments. The largest social network worldwide, Facebook, states that in March 2014 already 75.9% of all user accesses were made via mobile handhelds[1]. In general, the video consumption on mobile devices is expected to highly increase in the next years (Figure 1).

Typical mobile devices can either leverage cellular networks or infrastructure-based wireless LAN (WLAN). Cellular networks like UMTS or LTE are the first choice to access the Internet due to their ubiquitous availability. However, compared to WLAN, the client-side energy consumption is up to 32 times higher in cellular LTE networks [4].

[1] http://newsroom.fb.com/company-info/

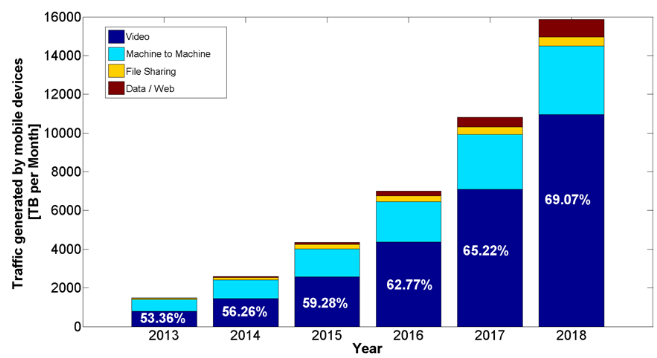

Fig. 1. Increase of mobile video traffic estimated by Cisco [5]

Furthermore, data caps in many contracts of cellular network providers limit the amount of high quality media content that can be up- and downloaded. Network providers charge extra fees at high prices for traffic extensions, since bandwidth is a scarce resource in cellular networks. This resource has to be shared with users within a certain local area: the bigger the coverage area of a single network cell, the more customers of network providers have to share bandwidth. The short range of WLAN allows spatial reuse of radio frequencies and causes less scarceness of wireless network bandwidth. Given a good estimate of which content is going to be accessed in the near future, it should be downloaded in advance, or prefetched, while the device has WLAN connectivity.

This good estimate can likely be derived from additional information which is linked to the video. Facebook offers a rich set of metadata such as likes, comments and social graph information for content which is uploaded to the social network. In this paper, we analyze this metadata and evaluate its feasibility to forecast future media consumption. We set up two user studies to observe video consumption habits of 34 users for 14 days in a mobile setting and 774 users for 34 days in average in a stationary setting. As data gathering for prefetching is a computationally intensive task which can be

978-1-4799-6390-4/15/$31.00 ©2015 IEEE

less attractive on mobile devices, our study is based firstly on stationary devices, using a browser plug-in, and secondly on 34 selected participants who allowed us to track their Facebook usage. We analyzed the newsfeeds with respect to the type of entries (text, pictures, videos, links), the pre-click delays, the relation to the author (friend or not) as well as the number of likes and comments.

In comparison to the previous work of Gautam et al. [3] and the Nettube system [2], we explicitly focus on the analysis of media access in the social network Facebook. We relate the video consumption of users in Facebook with contextual metadata, associated with the content, to find patterns that can be leveraged to predict future content retrieval. To the best of our knowledge, this is the first approach to investigate Facebook on such a scale to design prefetching methods.

Elaborating our measurements indicates that the social closeness of content producer and recipient helps in predicting media consumption for close friends and the family members, which are explicit subsets (groups) of what is called 'friend' in Facebook. We had hoped that prefetching based on affective expressions such as likes or comments can be effective. However, we did not find a strong correlation between the number of likes and comments and the probability for a video of being watched. As a further result, we argue that the time a user spends to evaluate a post before clicking on it can help to decrease startup delays. Our study participants tend to spend much more time to evaluate a wallpost before clicking on it as it would be the case without the intention to click on it. We thus suggest to start downloading videos before a click happens.

This work contributes to create a better understanding of predictions for prefetching videos based on likes, events, authorship and timings in OSNs. Applying our strategies can help to reduce network traffic in cellular networks for the mobile and to decrease startup times.

The remainder of this paper is structured as follows: In the next section, an overview on related approaches is given, showing the uniqueness of our approach regarding the number of participants and our method. The third section describes our experimental setup and shows which features are investigated to predict media consumption. With the fourth section, the gathered data is analyzed and discussed, elaborating on the impact of the features on the accuracy of media prefetching strategies. We finally conclude this work in section five.

II. Background and Related Work

Efficient mobile prefetching strategies are desirable for two reasons: 1) They allow to shift network traffic to the most cost and energy effective network interfaces of mobile devices, and 2) allow to reduce video startup delays. Huang et al. [4] show that generated LTE traffic is 32 times more energy consuming than WLAN.

Several studies focus on the impact of startup delays on the perceived QoE. In Krishnan and Sitaraman's work [8] the effect of different video startup times are evaluated. Thes demonstrate for short video clips such as those on YouTube that an increase of startup time increases the probability that a user cancels the streaming session. Prefetching may solve this issue as content is downloaded before the user requests it, resulting in a significantly reduced startup time.

An overview on video dissemination in OSNs is given by Li et al. [9]. They show that popularity of videos has a significant impact on their consumption via social networks. Unpopular video content disappears quickly from OSNs. This work concentrates on shared links of videos on an Chinese OSN that is claimed to be similar to Facebook. In contrast, we concentrate on videos that are directly included into the Facebook newsfeed. These 'internal' videos are more likely to be clicked and the social (graph) distances of authors and recipients can be evaluated. Our goal is to find the subset of videos which can be prefetched without wasting bandwidth and energy. In contrast, Li et al. focus on the popularity distribution of videos in general and thus neglect that unpopular videos still can be very interesting for a small subset of users e.g. because they are acquainted with people in the video.

Online Social Networks are increasingly leveraged by networking and caching researchers to predict media consumption. In [6], a recommendation-aware content placement strategy for Content Delivery Networks (CDN) is investigated. This work leverages similar information like we do in this paper, but focuses on storing content at different places within the network, whereas prefetching concentrates on downloading content to a device. Additionally Bai et al. [1] show social network related caching mechanisms for Facebook as well as in Yahoo News. Caching is powerful, but of course focuses on a large set of users, whereas prefetching through the downloading of content to an individual device can be tailored to each user individually. Prefetching, in contrast to caching in large-scale CDNs, has only a limited view on the OSN data - the view of the user. Thus, future video requests are potentially harder to predict.

Peer-to-Peer systems (P2P) are another research area which increasingly leverages social information gathered from OSNs. Wang et al. [11] investigate the most popular social network in China, called RenRen. They design a P2P-based prefetching algorithm that aims to reduce the video startup time. Similar to Wang's approach, we investigate with Facebook an OSN that has not been designed to share multimedia content by purpose. Their and our approach have in common that the social ties between users and the metadata can be used for prediction. In contrast, our work focus is on the suitability of prefetching on mobile devices with higher scarceness of resources. We aim to minimize the wasted bandwidth.

The SocialTube [10] system demonstrates the efficiency of a peer-to-peer-based social network. The purpose of the system is to reduce the playback startup time. The authors show that most of the video views are driven by social relationships and less by interests. We can confirm and refine the results. However, the P2P approach of the mentioned work seems not to be suitable to be migrated to mobile devices as it would generate much traffic on the client side - especially in cellular networks. Our aim is to leverage social network information to avoid this.

For scenarios in which video sharing sites are in focus, Khemmarat et al.[7] propose prefetching strategies considering related videos, and search query results. In contrast, Chen et al. [2] present the P2P streaming network NetTube which is customized for the video sharing site YouTube. In a long-term measurement, they show that 99.6% of all videos uploaded to YouTube have a playback time of less than 12 minutes. They give helpful insights on the structure of a video sharing site and how YouTube's infrastructure could be supported by the P2P paradigm. They propose to prefetch the first part of a video with a fixed length of ten seconds. Both approaches achieve high prediction rates but investigates YouTube a site which focuses on videos and neglects social ties. Their findings are closely related to caching experiments and recommender systems for video sharing sites.

None of the previous approaches concentrates on mobile devices and how users access OSNs from these devices. With the work of Gautam et al. [3] a mobile application is developed that allows to prefetch whole video clips based on arbitrary sources such as social networks or news feeds. The paper focuses on energy cost savings realized when applying prefetching on mobile devices. Zhao et al. [12] demonstrate a custom mobile Facebook application that integrates social network based algorithms for prefetching. It allows the conclusion that social prefetching is beneficial, but hard to conduct.

Our work enriches the field of research by adding a detailed large-scale analysis of OSN metadata with respect to its usefulness to predict videos being watched. We add new insights on the average pre-click viewing time of each post and its relation to the probability of a video to be watched in near future. The community additionally benefits from practical recommendations to identify the close friends in being a predominant factor of a video being watched.

III. Data Description

In this section, we describe the data collecting methods and the data gathered - both in the stationary as well as in the mobile setting.

A. Stationary Setting

Our approach in the stationary setting is to use browser plug-ins for Firefox and Chrome to collect data about user behavior in Facebook. We approached volunteers via newspaper articles to convince them to help us with our research. As a supportive incentive, we built statistic pages that help people to understand their own Facebook usage patterns.

The plug-ins read the Facebook wall of a user while the page was being rendered. Additionally, it collects usage patterns such as clicks. We saved and collected the meta information about which type of content was included in the walls (video, picture, link), the timing information (age of the entry, the time span between displaying and clicking / removing from the screen) as well as whether the author of a content item was part of the friendlist or not. The Facebook user ID, which was collected to distinguish the participants, was anonymized using hash functions to protect the privacy of users.

During our observation period of 123 days, the plug-ins (both versions for Firefox and Chrome together) has been installed by 2071 Facebook users. Since most people did not install the plug-in on the first day of our observation period and since some participants have left earlier, the average observation time was 34 days. We observed the phenomenon that many users were extremely passive and clicked only on very few content items. To ensure that we only use valid data, we excluded all cases where less than 100 wall entries were clicked. As a result, our analysis is based on 618,165 wall entries from 774 users.

B. Mobile Setting

To ensure that our findings from the stationary setting are also valid in mobile environments, we validated our findings by a small, prototypical evaluation on mobile devices. Data of 34 users over a time span of two weeks has been gathered and analyzed. Volunteering participants agreed on anonymously sharing their newsfeed's metadata (likes, comments etc.), information about their interactions with media posts as well as social graph information with us. We created an app for this purpose. In total, 8370 posts including 742 (8.9%) video posts and 3608 (43.1%) pictures have been analyzed.

IV. Analysis of the Collected Data

We analyzed our datasets to answer the following questions:

- What type of content can be found in the newsfeed?
- What type of content is consumed?
- Which fraction of the offered videos are watched?
- Does the number of its likes and comments change the probability of videos to be consumed?
- Does it depend on the author whether a video is being watched or not? - If yes, to which extent does the type of friendship or the membership in a global group play a role?

A. Stationary Setting

Our participants viewed an average number of 43 newsfeed entries per day (only days with activity are mentioned). Table I summarizes the newsfeed compositions with respect to the content type. It also shows what type of content was clicked by our participants.

	Photo	Video	Link	Total
Total	231,582	46,055	340,528	618,165
Clicked	13,342	5,259	21,636	40,237
% Clicked	5.76	11.42	6.35	6.51

TABLE I
SUMMARY OF THE NEWSFEED ENTRIES GATHERED IN OUR STUDY

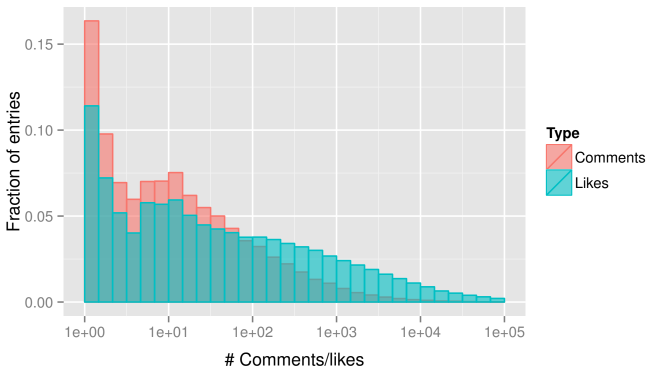

Fig. 2. Comparison of the number of comments and likes received by newsfeed entries

Interactions with the content are rare in general. Our participants clicked on 7%, liked 4% and commented less than 1% of the displayed content items. Figure 2 illustrates the distribution of clicks and likes. Newsfeed items usually have more likes than comments and items with multiple comments are very rare.

We are also interested in the relationship of the author and the user. This can be determined by looking at the friendlist. In average 49.5% of all newsfeed entries are authored by friends. The remaining posts are created by 'pages', (profiles maintained by companies to spread information) (41.4%) and content from strangers in case that friends liked or commented these items (9.1%).

Nearly half of the clicked newsfeed entries are those of friends (Figure 3). Pages are slightly less popular compared with their fraction in the newsfeed (41%). Nine percent of the viewed content was posted by strangers. Very interesting

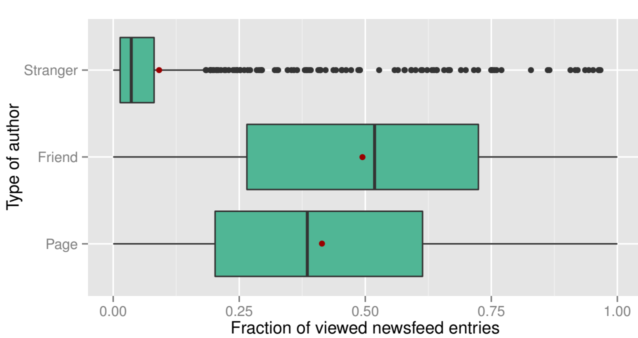

Fig. 3. Box-Whisker-Plot showing the distribution of clicked content items with respect to authorship; the red dots mark the averages

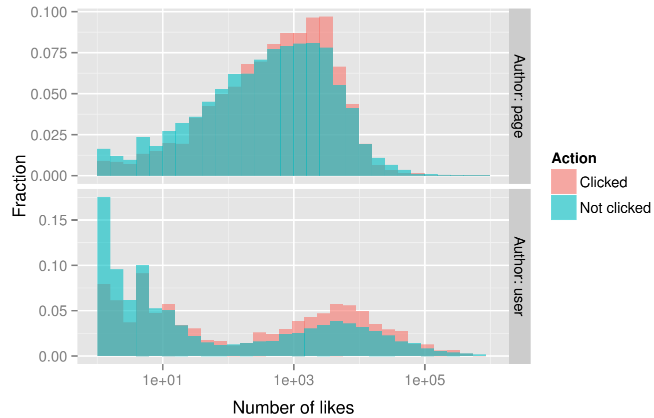

Fig. 4. Distribution of clicked and unclicked videos with respect to the number of likes; stacked plots indicate discrete numbers

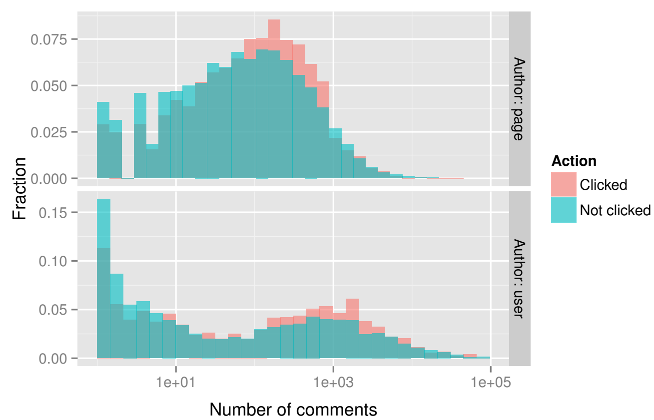

Fig. 5. Distribution of clicked and unclicked videos with respect to the number of comments; stacked plots indicate discrete numbers

at this point is that the fractions (with respect to authorship) of watched videos equals the fraction of displayed (videos shown on the wall) videos. As a result, authorship cannot be used as a predictor for prefetching content if it is considered at this granularity. We can later show in the mobile setting, that authorships of subsets of friends (close friends and family) can be used as predictors.

Figures 4 and 5 show the distribution of the number of likes and comments, which clicked and non-clicked videos have received before our study participants discovered the vidoes in their newsfeed. In both Figures, the distributions are very similar. They cover each other more than 90%, which means that the fraction of watched videos is nearly uncorrelated with the number of likes or comments. If the number of videos with a certain number of comments or likes increases, the number of watched videos with this certain number of comments or likes increases equally and the fraction of watched videos stays the same. That indicates that the number of likes and comments, attached to videos in Facebook, are no feasible prediction basis.

In Figures 4 and 5, we distinguished between professional pages, maintained by companies, and users since Facebook seems to use different algorithms for choosing newsfeed entries to display them on the user's wall. Our data indicates that content from professional pages needs to have far more likes and comments than items from users to be included into the newsfeeds of users.

However, the timing patterns (Figure 6) show a very clear indication that newsfeed entries that will be clicked, are watched for a longer time than those which will be removed without a click. This effect can be used to decrease the startup delay by starting to prefetch the item before it has been clicked.

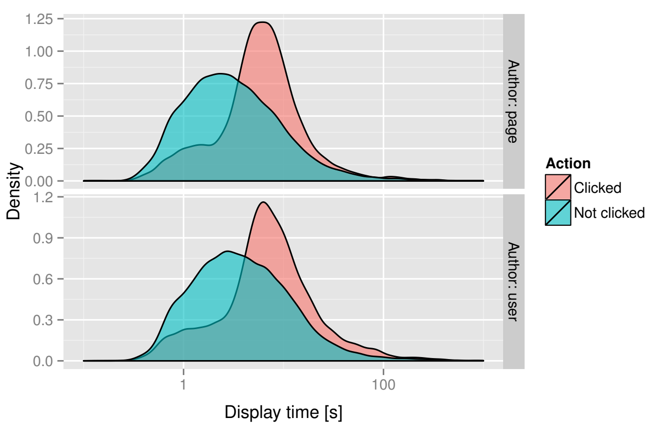

Fig. 6. Duration of newsfeed entries to stay in the browser viewport either before being clicked or removed

B. Mobile Setting

Our prototypical evaluation on mobile devices shows that the results obtained in the large-scale, stationary setting can be mapped to mobile devices. Figure 8 illustrates the influence of likes and comments on videos. Both categories, the one of clicked videos and the one of non-clicked videos, are normalized to one. A large proportion of the videos is neither commented nor liked at all.

It highlights two phenomena observable in social networks: 1) Many of the video posts, clicked as well as non-clicked, have no likes or comments and 2) a low number of likes and comments, does not mean that the videos are not watched. It shows that many videos are consumed shortly after publishing, or that they are distributed to only small groups of friends. This is supported by Table II which investigates friendlists.

The distance to a given user is approximated by her membership in a friendlist. We show in Table II only the subset of posts that has been shared by users within a friendlist. Despite the anonymization of traced data, standard Facebook groups such as 'family' and 'close friends' can still be identified. For both videos and photos, posts by close friends and family are preferred. The information on the social distance to a posting user can thus easily be identified and leveraged for prefetching

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

TABLE II
FRIENDLISTS AND THEIR IMPACT ON CONSUMING VIDEO AND PHOTOS - TABLE SHOWS PERCENTAGE OF MEDIA SHARED BY MEMBERS OF A FRIENDLIST THAT WAS CLICKED

mechanisms. The videos shared by close friends or family are predominantly those with low numbers of likes or comments. None of the videos reshared in our dataset has one thousand or more likes. It demonstrates that videos are either 1) consumed quickly after being shared, or 2) the videos watched from friends don't have the necessity to be very popular in the social network.

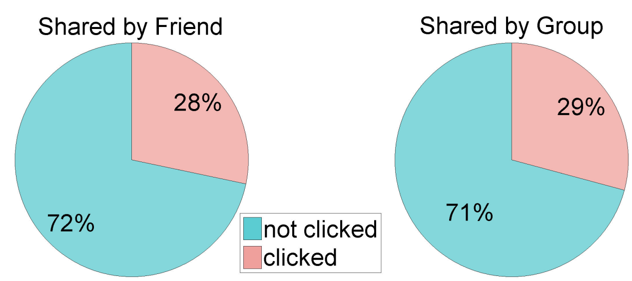

Fig. 7. Impact of photos and videos shared by friends versus global Facebook groups or pages

In total the proportion of unwatched videos is, as shown for stationary evaluation, nearly two times higher than those being watched. A significant outlier is located between 100 and 1000 likes. In this range, the percentage of clicked videos is nearly as high as for videos without likes.

Comparing content shared by friends with content posted by a Facebook group, it can be seen that video consumption patterns are very similar. This is true for both videos and photos. Figure 7 illustrates that for both video as well as photos the consumption behavior is the same as if the origin is a friend or a group.

V. CONCLUSION

In this work, a large-scale analysis on the impact of comments, likes and friends as originators on video consumption is evaluated. We discovered that short-time prefetching is an option to reduce the start-up delay of videos. The reason is that clicked posts remain longer in the browser viewport (before being clicked) than those that will not be clicked. This initial time can efficiently be leveraged to initiate download connections and stream the first chunks of a video. Since the exact time strongly depends on the user, we suggest to initiate download connections after two seconds as a rule of thumb.

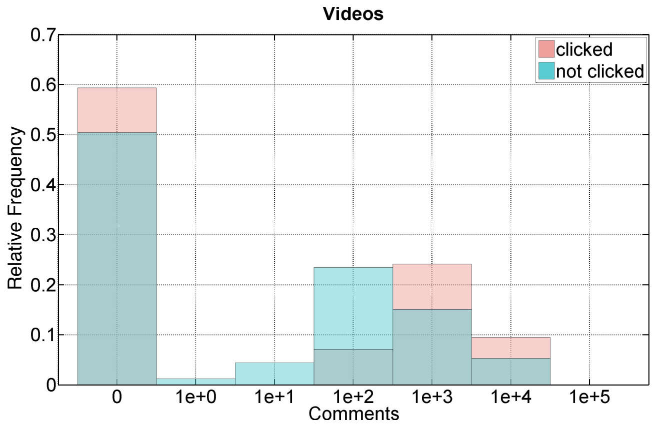

(a) Effect of number of comments

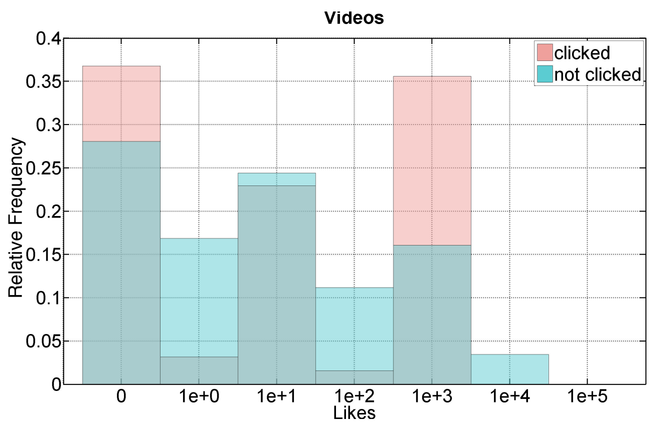

(b) Effect of number of likes

Fig. 8. Influence of the number of comments and likes on the consumption of videos on mobile devices

Furthermore, we did not observe videos to have a high likelihood of being watched in case that they count a high number of likes and comments. That indicates that prefetching mechanisms should not solely be based on the pure number of likes and comments. It is necessary to integrate measures such as closeness of a friend as originator.

We have also shown that 'being' a friend alone is not sufficient, but that the type of friendship is important. 85.7% of the videos, shared by users that are part of the group 'close friends', and 50% of the videos from family members are watched. This is independent from the number of likes, comments and the freshness of the videos (in case that they are still displayed in the feed). Since only 8.3% of the other friends' videos are watched, we suggest to prefetch the videos of authors which are labeled to be 'close friend' or family member. Relying on those both first major insights of this work, a next step is to design prefetching mechanisms for mobile devices that take the individual user characteristics into account.

VI. Acknowledgements

This work has been co-funded by the German Research Foundation (DFG) in the Collaborative Research Center (SFB) 1053 'MAKI – Multi-Mechanisms-Adaptation for the Future Internet.

We thank Wolfgang Effelsberg for his advices to improve this work.

References

[1] X. Bai, F. P. Junqueira, and A. Silberstein. Cache refreshing for online social news feeds. In *CIKM*. ACM, 2013.

[2] X. Cheng, S. Fraser, and J. Liu. Nettube: Exploring social networks for peer-to-peer short video sharing. In *IEEE INFOCOM*, 2009.

[3] N. Gautam, H. Petander, and J. Noel. A comparison of the cost and energy efficiency of prefetching and streaming of mobile video. In *MoVid*. ACM Press, 2013.

[4] J. Huang, F. Qian, A. Gerber, Z. M. Mao, S. Sen, and O. Spatscheck. A close examination of performance and power characteristics of 4g lte networks. In *MobiSys*, 2012.

[5] C. S. Inc. Cisco visual networking index: Forecast and methodology, 2013-2018, 2014.

[6] M. A. Kaafar, S. Berkovsky, and B. Donnet. On the potential of recommendation technologies for efficient content delivery networks. *SIGCOMM Comput. Commun. Rev.*, 43(3):74–77, July 2013.

[7] S. Khemmarat, R. Zhou, D. Krishnappa, and L. Gao. Watching user generated videos with prefetching. In *MMSys*, 2011.

[8] S. S. Krishnan and R. K. Sitaraman. Video stream quality impacts viewer behavior: Inferring causality using quasi-experimental designs. In *IMC*, pages 211–224. ACM, 2012.

[9] H. Li, H. Wang, J. Liu, and K. Xu. Video requests from online social networks: Characterization, analysis and generation. In *IEEE INFOCOM*, 2013.

[10] Z. Li, H. Shen, H. Wang, G. Liu, and J. Li. Socialtube: P2p-assisted video sharing in online social networks. In *INFOCOM*, 2012.

[11] Z. Wang, L. Sun, and S. Yang. Prefetching strategy in peer-assisted social video streaming. In *ACM Conference on Multimedia*, 2011.

[12] Y. Zhao, N. Do, S.-T. Wang, C.-H. Hsu, and N. Venkatasubramanian. O2sm: Enabling efficient offline access to online social media and social networks. In D. Eyers and K. Schwan, editors, *Middleware 2013*, LNCS, pages 445–465. Springer Berlin Heidelberg, 2013.