Understand Instant Video Clip Sharing on Mobile Platforms: Twitter's Vine as a Case Study

Lei Zhang
School of Computing Science
Simon Fraser University
Burnaby, BC, Canada
lza70@cs.sfu.ca

Feng Wang
Department of Computer and Information Science
The University of Mississippi
University, MS, USA
fwang@cs.olemiss.edu

Jiangchuan Liu
School of Computing Science
Simon Fraser University
Burnaby, BC, Canada
jcliu@cs.sfu.ca

ABSTRACT

With the rapidly development of mobile networking and end-terminals, anytime and anywhere data access become readily available nowadays. Given the crowdsourced content capturing and sharing, the preferred length becomes shorter and shorter, even for such multimedia content as video. A representative is Twitter's Vine service, which, available exclusively to mobile users, enables them to create ultra-short video clips, and instantly post and share them with their followers. In this paper, we present an initial study on this new generation of instant video clip sharing service over mobile platforms, taking Vine as a case. We closely investigate the architecture of Vine, and reveal how its service is empowered with a combination of advanced mobile and cloud computing platforms. Through a dataset of over $50,000$ video clips and over $1,000,000$ user profiles, which is available online for academic use, we examine the unique viewing behaviors of Vine uses, particularly batch viewing and passive viewing. We further analyze the video lifetime and propagation patterns in this new service, as well as the distinct social relations therein. Our study lead to critical observations that would help with improving the energy-efficiency and scalability of Vine-like services.

Categories and Subject Descriptors

H.5.1 [**Multimedia Information Systems**]: Video; C.4 [**Performance of Systems**]: Measurement techniques

General Terms

Measurement

Keywords

Mobile, Video Sharing, Social Network

NOSSDAV'14, March 19-21, 2014, Singapore, Singapore.

ACM 978-1-4503-2706-0/14/03 $15.00.
http://dx.doi.org/10.1145/2578260.2578278.

1. INTRODUCTION

In the past decade, the user-generated content available on the Internet (e.g., images, videos, micro-blogs, etc.) has presented an explosive growth trend, which provides significant opportunities for both understanding how users utilize the Internet and enhancing their experiences. Specifically, for the online video sharing service, we have seen a three-stage evolution, and the shared videos have shown different characteristics in the time dimension through these three stages. At the earliest stage, users upload and view videos directly on video sharing sites (VSSes) such as YouTube, and they can select the related videos recommended by VSSes [9] or search the interested videos via search engines. Usually, users browse VSSes and view videos in tens of minutes at this stage. The emergence of online social networks (OSNs) (e.g., Facebook, Twitter, Google+, etc.) however has greatly changed such access patterns through proactively and efficiently sharing among friends the video links from external VSSes [7]. The common interests of social groups direct users to the videos that can last for minutes. Recently, the new generation of social applications has been focused on media content sharing among mobile users such as photo sharing (Flickr) or short video sharing (Vine) or both (Instagram). The video clips consumed by these social applications are often as short as several seconds, and thus the file size of each clip is relatively small, which allows users to view multiple video clips within a very limited time from mobile devices.

To reveal the characteristics of video sharing services or social networking applications that allow video sharing, numerous measurement studies [1, 2, 3, 6] have been done. Yet very few have focused on instant video clips sharing over mobile platforms. To this end, we attempt to present an initial study on this new generation of instant video clip sharing service over mobile platforms, taking Vine as a case. In this paper, we closely examine the unique features of the Vine-like services. We first investigate the architecture of this mobile system in detail, which is empowered with a combination of advanced mobile and cloud computing platforms. Based on our observations, we reveal the unique viewing behaviors of Vine users, particularly batch view and passive view. Further, we implement customized crawlers to collect the repost traces of video clips and user profiles. Through the analysis of collected data, we identify the key differences between Vine-like services and traditional video sharing services, including highly skewed popularity, fast propagation, short lifetime, dense social network.

The rest of the paper is organized as follows. In Section 2, the architecture of Vine and the underlying cloud computing services are closely examined. We next present our measurement results and analysis in Section 3. A further discussion is provided in Section 4. Finally, Section 5 concludes the paper.

2. A CLOSER LOOK AT VINE

As a case study, this paper analyzes the characteristics of Vine, a representative of the new generation of instant video clip sharing services on mobile platforms. Vine is a mobile social networking application that enables its users to create ultra-short video clips, as well as post and share them with their followers. The service was introduced with a maximum clip length of six seconds and can be shared or embedded on social networking services such as Twitter (which acquired the application in October 2012) and Facebook. Unlike traditional social networking services, Vine exclusively focuses on mobile users, and does not have its counterpart on PCs. As claimed by Vine's official site, Vine has attracted over 40 million registered users since its initial release in January 2013 on iOS platforms. Though Vine was initially available only for iOS devices, it also has been appeared on other major platforms: Vine for Android was released on June 3, 2013 for devices with Android OS version 4.0 or higher; the Windows Phone 8 and Windows 8 version of the Vine app was also unveiled during the Microsoft 2013 Build Conference.

A typical mobile Vine client has four pages: Home/Feed, Explore, Activity, Profile. In the "Home/Feed" page, users can view, like, comment, and share (repost) the recent posts from others that they follow. The functionality of this page is very similar to the traditional social networking services such as Twitter. The only, but the significant difference is that, instead of texts, the media that people are viewing and sharing is ultra short video clip. Vine allows users to search the videos and the people they are interested in, and also provides the popular/trending posts in Vine, as well as many dedicated channels focused on the specific topics in the "Explore" page where it tries to offer services like traditional video sharing sites such as Youtube. The other two pages are for the normal usage: the "Activity" page acts as the notification center showing the recent events related to the current user; and the user can customize his personal settings in the "Profile" page. The social relationship in Vine is the follower-following relationship, which is similar to a number of other social network services, such as Twitter, Instagram, Flickr. Users of Vine can follow others whom they are interested in, and then receive updates from them. They can also hit the pubic channels to view promoted videos or videos on specific topics. If a user feels a video clip very interesting, s/he can comment, like or repost it. The similar social relationship and user actions make Vine representative for other social network services. On the other hand, the 6-second only instant video clip sharing in this mobile application makes Vine unique.

2.1 Service Framework

Although Vine has become extremely popular in such a short time, its underneath framework design is yet to be examined. To understand Vine's architecture as well as the potentials and challenges with this new generation of mobile social networking service, we have conducted a traffic measurement and analysis from our university campus. Besides capturing the traffic between the test device and the Vine server, we also intercepted the SSL connections between them to view the detailed requests from the application by using the mitmproxy tool. The traces show that Vine builds its system based on a cluster of cloud computing services including Amazon EC2, Amazon S3, Amazon CloudFront, and CDNs provided by Akamai and Fastly. We accordingly illustrate the Vine's service framework in Figure 1. The vine client on the mobile terminal initiates and maintains a HTTPS connection with the application server running on the EC2 instances with the domain name "api.vineapp.com". After the authorization process, users can then make requests to the Vine server, and the server in turn provides proper responses, which allow users to complete actions such as browsing, search, post, comment, like and so on. When a user logs into Vine (or returns to the "Home/Feed" page), the client makes a GET request for the timeline information which corresponds to the recent updates for this user. After receiving the response, the client can further make GET requests to CDNs with the domain name "v.cdn.vine.co" or "mtc.cdn.vine.co" to download the video clips and the corresponding thumbnails according to the URLs provided by the Vine server. The same thing happens when a user searches and views his interested videos in the "Explore" page. The slight difference is that the static web images in this page (such as the pictures for the channels) are distributed by Amazon CloudFront.

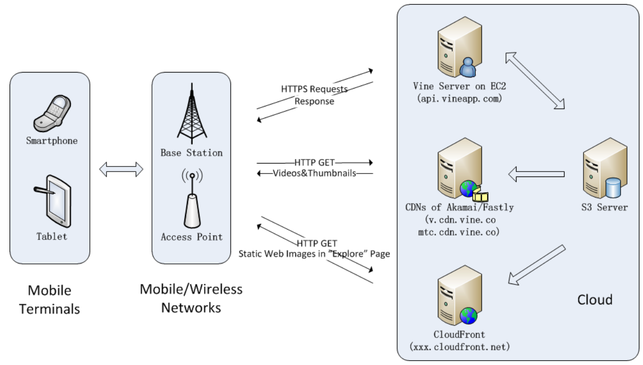

Figure 1: Vine's service framework

Figure 2 shows another important operation scenario where a mobile user makes a new post on Vine. First, the user uses the toolkit in the application to record a video clip, and then compresses it locally on the mobile device so that the compressed video clip satisfies the general standard for different platforms/devices (file format, file size, frame rate). Our measurement shows that Vine uploads the video clip (a .mp4 file and a .jpg file) to the load balancer on EC2 with the domain name "media.vineapp.com" once it completes recording and compressing. If the user is satisfied with the recorded video and confirms to post it, a "Post" request is sent to the Vine server, and then the server generates an unique ID for the post and updates the user's profile. While uploading the video clip, the user can also tag the post or assign it to a specific channel, so that other users can identify and classify the content of this post, and access it through searching. Now both the user and his followers can get the updated timeline information if they return to the "Home/Feed" page (other users may access it

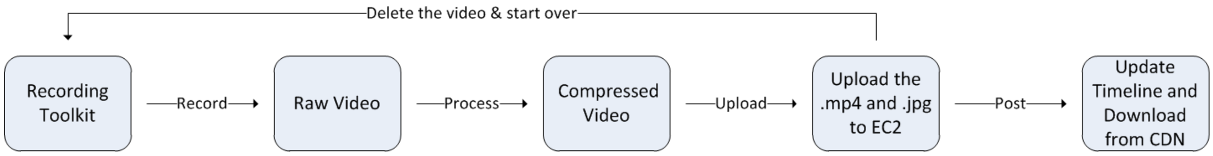

Figure 2: Vine post flow

in the "Explore" page if the tag or the assigned channel is provided), and can view the post after the video (.mp4 file) and the thumbnail (.jpg file) are downloaded from CDNs. There are a couple of findings worth noting here, which may be useful for the future work to improve Vine's performance. First, the recorded video is uploaded right after it is compressed, even before the user decides to post it. Second, the user himself also needs to download the post that he records from the CDNs to view it on the home page.

2.2 Key Behaviors: Batch View & Passive View

As Vine is dedicated to provide instant video clip sharing service exclusively for mobile users, its user experience is significantly different from traditional video sharing services and traditional social networking services that allow video sharing. In traditional VSSes and OSNs, users need to click on the interested item to view or link to one specific video, which only allows users to view one video each time/click. Vine, however, returns a playlist of video clips after the user decides to view the updates for certain users, tags, or channels. As the user scrolls the screen of his smartphone/tablet, a number of video clips from the generated list are played seamlessly. It is worth noting that scrolling the screen is the key and unique user action for mobile users of instant video clip sharing service, as the instant video clips from mobile users are often short and of small size, and thus can be watched in a very short time (while the user scrolls the screen). Therefore, we introduce a special term for this unique user behavior of viewing multiple video clips with the screen scrolling: Batch View. The batch view behavior implies that mobile users can watch a considerable amount of instant video clips within the playback time of one normal video such as a YouTube video.

Besides batch view, we also identify another important user behavior, Passive View, as the mobile users have to "passively" watch some of the video clips in the playlist. Since Vine has a Twitter-like user interface, the selected contents that will be viewed are arranged in order, and users have no control over the order of the playbacks. If the user is only interested in two specific video clips that are separated in the playlist, he may have to watch all the video clips between them. The only way to proactively skip a video clip is to scroll the screen so fast that there is not enough time to play it, which can hardly be controlled well in practice. If the user screen moves slowly, the following ultra short video clip may have been downloaded and started to play, and thus the user will probably watch it. As the 6-second playback time is so short that it is even not enough to finish reading the description of some video clips (users may have already watched the video clip before making the decision based on the description), the Vine users tend to passively view numerous video clips. The two key user behaviors abstract and emphasize the unique features of the new generation of mobile video clip sharing service. And the observation of these two key user behaviors can help us further understand the propagation of popular video clips among the mobile users. Detailed discussion are presented in the following section.

3. MEASUREMENT AND ANALYSIS

3.1 Methodology and Data Sets

One critical challenge to measure and analyze Vine is that Vine does not provide official application programming interfaces (APIs) for developers or researchers to collect the related data. To this end, we have to develop a series of mechanisms to crawl the data. In particular, we first capture the packets sent by the mobile device when using Vine service. Through the traffic analysis, we have identified that Vine uses HTTP and HTTPS for the communication. Besides capturing the traffic between the test device and the Vine server, we also intercept the SSL connections between them to learn the detailed requests from the application by using the mitmproxy tool. We then develop customized crawlers that imitate the behaviors of the mobile application by sending the exact same requests and substituting the header information with the headers of the mobile browser. By this means, our crawlers can get any data that a normal user can get when using Vine on a mobile device.

Our data set consists of two parts: repost traces of over 50,000 video clips, and over 1,000,000 user profiles. As the most important part, we collected the traces of the video clips that were posted on 16 user channels (47,794 posts) and 2 promotion channels (8,891 posts) from November 16 to December 14, 2013. After the collection, for each video clip, we accessed and recorded its repost history, including the exact time when it was shared and the user who made this repost. All the recorded video clips are from two types of public channels. One type of public channels are user channels focusing on dedicated topics, where each channel has two sections: recent and popular. Users can upload their video clips to any recent section of these 16 channels, and each user channel lists a small number of popular posts in the popular section. Unlike user channels, the other type of public channels, the promotion channels, do not accept the posts directly from the normal users. Instead, they choose most popular and most trending videos clips among all the recent posts on Vine. The second part is the user profile collection. For each user, we recorded the number of followers, the number of followings, and the number of posts. A more detailed description and the dataset can be found at http://netsg.cs.sfu.ca/vinedata.html.

3.2 Video Popularity

As we cannot access the actual number of views for each video clip, we use the number of reposts to evaluate the popularity of the video clips. We first plot the number of reposts as a function of the rank of the video clip by its

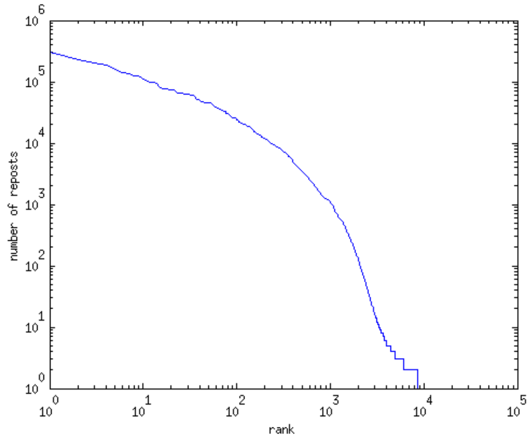

Figure 3: Video clips rank ordered by the number of repost

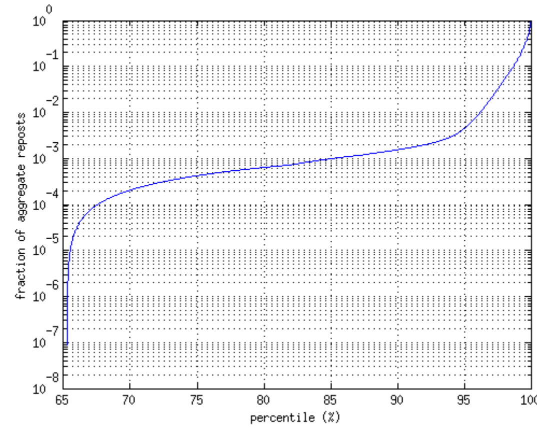

Figure 4: Skewness of popularity across video clips from the user channels

popularity in Figure 3. To reveal the property of general video clips, we use the data from all the 16 user channels. As we can see, the plot does not follow a Zipf distribution, which should be a straight line on a log-log scale. This is consistent with the results [3] for traditional video sharing services. Another observation is that the curve in the figure drops quickly. This implies that the number of reposts drops dramatically as the rank of the video clip getting lower and a large number of video clips get no repost at all.

To further understand how the popularity is distributed among the posts, we plot the cumulative proportion of the total number of reposts versus the percentile of the video clip in Figure 4. It shows that the popularity of video clips in the user channels is extremely skewed: the top 5% video clips accounts for more than 99% reposts. It heavily deviates from the Pareto Principle (or 80-20 rule), which is widely used to describe the skewness in distribution. This result is quite surprising, since other video sharing services [2] show much smaller skewness. We also check the popularity skewness of video clips in the promotion channels, and it fits the Pareto Principle well (the figure is not presented). These results imply that the promoted video clips are much more popular than the unpromoted ones. The rationale can be provided based on the organization of user channels, and the batch view and the passive view behavior. When a video clip is first unloaded to a user channel, it appears on the top of the playlist of the recent section. As time goes by, the unpopular video clips move towards the end of the playlist, and then can hardly be seen by other users again. On the contrary, the popular video clips will be moved to the popular section, and when the other users access the popular section they can easily reach these posts. As the video clips getting more and more popular, they keep on the top of the playlist of the popular section, and therefore can be accessed more frequently by other users. This leads to that the popular video clips become more popular, and the unpopular ones can hardly get users' attention.

3.3 Video Lifetime and Propagation

From the discussion above, we know that some video clips are very popular (the number of reposts grows very fast), while others are not. To investigate how the number of reposts changes with time, we plot Figure 5, which shows the average daily number of reposts after the video clips were created. As the popularity of the collected video clips is highly skewed, we only consider popular video clips in the following analysis, specifically, the top 5% reposted video clips from the user channels and all the video clips from the promotion channels. One thing worth to note is that, although our data collection period is slightly shorter than 1 month, many of the video clips that we explored may have been there for a while before we started crawling. We can see that the average number of reposts for the popular video clips monotonically decrease day by day. This result indicates that, even for many of the popular video clips, they are most popular during the first day after the initial posts and are getting less and less popular afterwards.

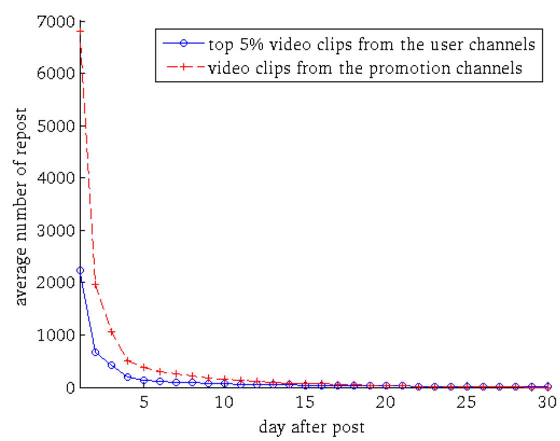

Figure 5: Daily number of reposts

We define the active lifetime of a video post as the duration from its initial post to the first day in which it gets no repost, and the CDF of active lifetime of the popular video clips is plotted in Figure 6. Here we use a real value (0) as the threshold to decide whether the video clip is active in propagation, instead of other metrics such as the changing rate and the moving average. The reason is two-fold: first, as shown in Figure 5, although the number of reposts for the popular video clips may change dramatically in the first few days, it still can be a large value; second, we can hardly know the impact of one repost, as the number of passive viewers after each repost varies significantly (if the user who shares the video clip has a large number of followers, this repost can have a potentially large impact on the propagation of the video clip). Although we give a loose definition of active lifetime, as we can see in Figure 6, more than half of the popular video reposts can only stay in active for less than 10 days. This result is quit different from the previous observations on tradition video sharing services [3]: a popular YouTube video can easily last for weeks, even months.

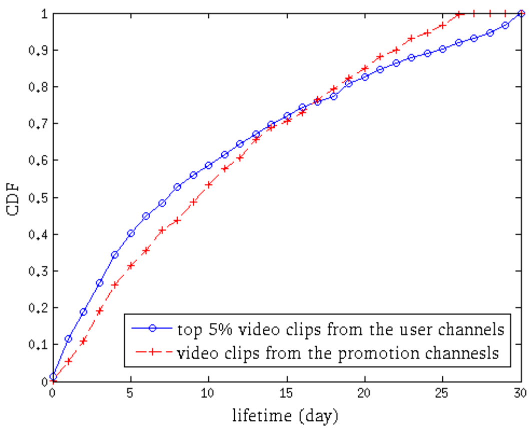

Figure 6: CDF of lifetime

We further examine the propagation speed of the most popular video clips (top 1% video clips from the user channels and video clips from the promoted channels) by tracking the average number of reposts with different granularities. Figure 7(a) tracks the average number of reposts of the popular video clips in the first 24 hours after their initial posts with the granularity of 1 hour; Figure 7(b) plots the average number of reposts in the first 2 hours with the granularity of 5 minutes. As shown in the two figures, the popular video clips propagate extremely fast, and the propagation speed reaches the maximum within a couple of hours. This is consistent with our previous findings: once a video clip cannot attract users' interests in the early stage after its initial post, it can hardly be watched by the massive users later; on the other hand, if a video clip can gain users' attention at the very beginning, it will probably propagate fast in the network and stay popular. Another observation is that the propagation speed drops to a minimum after about 15 hours in the left figure. It is reasonable because most of the video clips can only be shared locally around the creator's community. Therefore, we can expect to see a daily pattern in the video clips' propagation since the human activities exhibits a daily pattern.

3.4 Social Relationship

As Vine is a social networking application, we analyze its social relationship based on 1,151,938 users' profiles. Figure 8 indicates that the number of followers does not have a strong correlation with the number of followings. Our data further shows that the average number of followers and

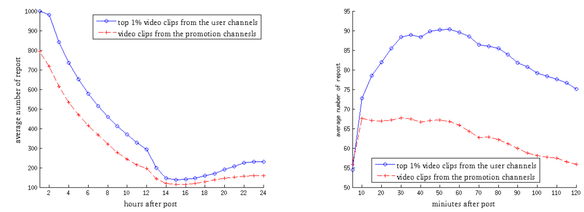

(a) the average number of reposts for the first 24 hours (b) the average number of reposts for the first 2 hours

Figure 7: Propagation of the popular video clips

followings in this Vine social network are 244.4 and 102.4 respectively, both of which are much higher than those of the Twitter social network [5]. These results imply that the social relationship in Vine does not rely on the real-world social relationship (such as friends, colleagues, employers, etc.). One user may follow a significant number of users that he does not know. Therefore, we can infer that the social relationship in Vine is a content-driven relationship: the reason for user A to follow user B is that user B has posted some video clips that user A is interested in. To support our argument, we further check the top 10 of the most followed usernames on Vine as shown in Table 1. In the traditional social networks, the most followed users can span a wide variety of public figures and news sources [1], which often include mass media, politicians, athletes, as well as other celebrities like actors, writers, musicians, etc. However, all the top 10 users in Table 1 are more famous for their large numbers of followers in Vine. Other users follow them because they post interesting and creative video clips, rather than they are real-world celebrities or news sources.

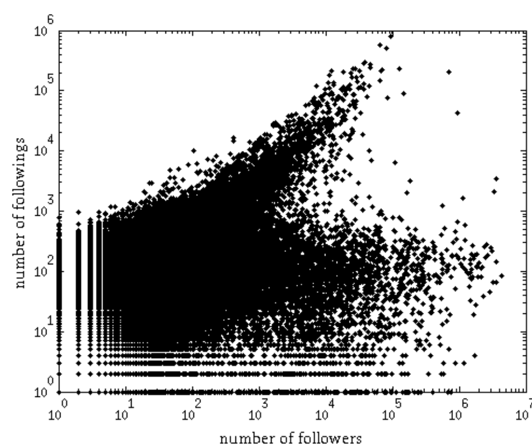

Figure 8: number of followers *vs* number of followings

4. POTENTIAL IMPROVEMENT AND APPLICATION

As we have shown in the previous section, Vine is significantly different from traditional video sharing services such as YouTube and other social applications that allow video sharing such as Twitter, in terms of popularity skewness, video lifetime and propagation, social relationship, etc.

username	number of followers
KingBach	4,327,113
JEROME JARRE	3,603,981
Marcus Johns	3,565,338
Josh Peck	3,547,111
Nicholas Megalis	3,398,624
DEM_WHITE_BOYZ	3,186,003
Curtis Lepore	3,067,646
Kc James	2,978,526
Alx James	2,859,786
Rudy Mancuso	2,686,820

Table 1: Top 10 of the most followed users

Therefore, previous systems that are designed based on the properties of tradition video sharing services may not work well in the mobile instant video clip sharing scenario. Our measurement results can shed the light on the approaches for performance improvement. For example, energy efficiency is always a key concern for mobile systems. For the uploading side, to fully utilize the cloud's capabilities and reduce the energy cost on mobile devices, we can offload the local computation-intensive tasks such as video compression to the cloud [8]. For the downloading side, some pre-fetching schemes can be applied to reduce the transmission cost based on the users interests prediction and current network conditions. Other issues such as replica management in CDNs can also be reconsidered and improved. The popular video clips propagate so fast in Vine that the peak of CDNs' workload can be reached within a couple of hours after the initial posts. Meanwhile, the lifetime of instant video clips is so short that most of them will be replaced in a few days. Therefore, there may be some opportunities to design effective replication strategies that suit the propagation of the instant video clips among mobile users. More interesting topics can be investigated if we consider the human interactions with mobile devices. As people take time to understand media content such as video and text, some smart downloading and playback strategies can be applied, which may reduce the frame rate of the video clips as users scroll the screen so fast that missing frames would not affect the users experience.

As Vine's users tend to have large number of followers and followings, Vine can be a potentially suitable platform for CrowdSourcing applications [4]. The fast propagation of the instant video clips can accelerate the information diffusion among the interested users. Name a simple example: if users have seen great deals during the Black Friday shopping time, they can make video clips, and upload them to a public channel or assign them to a specific tag, so that other users can easily access them. As more and more users contribute to this channel or tag, the information about great deals will aggregate and become more valuable. Users (consumers of the information) can use others' feedback to evaluate a certain piece of information: if this video clip has been reposted or liked many times, it would probably be a good deal; or, users can directly read the comments from other users to decide whether this information is useful.

5. CONCLUSIONS

In this paper, we have investigated the characteristics of the new generation of instant video clip sharing services on mobile platforms. Taking Vine as a case study, we have first closely examined the architecture of the system of Vine, and studied underlying services that enable this mobile social application. Further, based on the analysis of our crawled data set, we have identified the key difference between Vine and traditional video sharing services, including highly skewed popularity, fast propagation, short lifetime, dense social network. Finally, we have provided critical observations and discussions that would help with improving the energy-efficiency and scalability of Vine-like services and extending Vine-like services as approaches of crowdsourcing.

6. ACKNOWLEDGMENTS

Part of this work is supported by a Canada NSERC Discovery Grant, an NSERC Engage grant, and an NSERC Strategic Project Grant. Feng's work is supported by a Startup Grant from the University of Mississippi.

7. REFERENCES

[1] M. Cha, H. Haddadi, F. Benevenuto, and P. K. Gummadi. Measuring user influence in twitter: The million follower fallacy. *ICWSM*, 10:10–17, 2010.

[2] M. Cha, H. Kwak, P. Rodriguez, Y.-Y. Ahn, and S. Moon. I tube, you tube, everybody tubes: analyzing the world's largest user generated content video system. In *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, pages 1–14. ACM, 2007.

[3] X. Cheng, C. Dale, and J. Liu. Statistics and social network of youtube videos. In *Quality of Service, 2008. IWQoS 2008. 16th International Workshop on*, pages 229–238. IEEE, 2008.

[4] P. Jain, J. Manweiler, A. Achary, and K. Beaty. Focus: clustering crowdsourced videos by line-of-sight. In *Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems*, page 63. ACM, 2013.

[5] A. Java, X. Song, T. Finin, and B. Tseng. Why we twitter: understanding microblogging usage and communities. In *Proceedings of the 9th WebKDD and 1st SNA-KDD 2007 workshop on Web mining and social network analysis*, pages 56–65. ACM, 2007.

[6] H. Kwak, C. Lee, H. Park, and S. Moon. What is twitter, a social network or a news media? In *Proceedings of the 19th international conference on World wide web*, pages 591–600. ACM, 2010.

[7] T. Rodrigues, F. Benevenuto, M. Cha, K. Gummadi, and V. Almeida. On word-of-mouth based discovery of the web. In *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*, pages 381–396. ACM, 2011.

[8] Y. Zhao, L. Zhang, X. Ma, J. Liu, and H. Jiang. Came: cloud-assisted motion estimation for mobile video compression and transmission. In *Proceedings of the 22nd international workshop on Network and Operating System Support for Digital Audio and Video*, pages 95–100. ACM, 2012.

[9] R. Zhou, S. Khemmarat, and L. Gao. The impact of youtube recommendation system on video views. In *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*, pages 404–410. ACM, 2010.