Beyond the Watching: Understanding Viewer Interactions in Crowdsourced Live Video Broadcasting Services

Xiaodong Wang, *Student Member, IEEE,* Ye Tian, *Member, IEEE,*
Rongheng Lan, Wen Yang, and Xinming Zhang, *Senior Member, IEEE*

***Abstract*—Crowdsourced live video broadcasting services like Twitch and YouTube Live are becoming increasingly popular in recent years. In such a service, viewers are allowed to perform rich interactions, such as posting comments and donating monetary virtual gifts, when watching the videos. Understanding viewer interactions is essential for people to comprehend the productions and consumptions of the crowdsourced live video content, and improve the service. However, the basic characteristics of the viewer interactions are still unknown. In this paper, we present a comprehensive measurement study of the viewer interactions on Douyu, a popular crowdsourced live video broadcasting website in China. Our measurement spans four months, and contains comment posting and virtual gift donating interactions from tens of millions of viewers in hundreds of thousands of channels. Based on the measurement data, we carry out a content analysis on danmu comments, and characterize the patterns of the viewer interactions. We build a suite of models for capturing the gift donating process, viewer activity, and channel popularity. We further analyze the influences of the broadcaster's behavioral factors on a channel's popularity, and present methodologies for popularity predicting. Our measurement and analysis have important implications on the design and business policy of the crowdsourced live video broadcasting services.**

***Index Terms*—Crowdsourced live video broadcasting, viewer interactions, channel popularity, measurement, modeling.**

I. Introduction

CROWDSOURCED live video broadcasting (*livecast* for short) services, which allow anyone to broadcast live videos from anywhere on the Internet, have attracted millions of viewers and formed a big entertainment industry in recent years. For example, Twitch, a crowdsourced game video livecast website owned by Amazon, has become the fourth largest source of the Internet peak traffic in the US [1]. And Douyu, one of the prominent Chinese crowdsourced livecast websites, has attracted over two billion RMB of investment since its founding in 2014.

This work was supported in part by the National Natural Science Foundation of China under Grants 61672486 and 61672485, and in part by subtask of the National Major Science and Technology Project of China on New Generation Broadband Wireless Mobile Communication Networks under Grant 2017ZX03001019-004. *(Corresponding author: Ye Tian.)*

X. Wang, Y. Tian, W. Yang, and X. Zhang are with the Anhui Key Laboratory on High Performance Computing and Applications, School of Computer Science and Technology, University of Science and Technology of China, Hefei, Anhui 230026, P.R. China (e-mail: qsmywxd@mail.ustc.edu.cn; yetian@ustc.edu.cn; wenyang0@mail.ustc.edu.cn; xinming@ustc.edu.cn).

R. Lan is with the School of Information Science and Technology, University of Science and Technology of China, Hefei, Anhui 230026, P.R. China (e-mail: lan0729@mail.ustc.edu.cn).

Different from the conventional IPTV services (e.g., [2] and [3]) that broadcast TV channels over the Internet, crowdsourced livecast has some unique features. First, most broadcasters are amateurs, and they do not broadcast $7/24$, while on the other hand, there can be up to thousands of channels concurrently available on a livecast platform. Second, nearly all the crowdsourced livecast websites allow viewers to perform rich interactions such as posting comments and sending hearts, while watching a channel, and websites like Twitch and Douyu even allow viewers to purchase virtual goods/gifts from the website with real money, and donate to their preferred broadcasters for financially supporting them.

Understanding user behavioral patterns is essential for improving Internet video streaming services [4] [2] [5] [6] [7] [8] [9]. However, previous studies are focused on conventional IPTV systems, where a viewer can only watch channels passively [2] [5] [9]. On the other hand, the emerging crowdsourced livecast service differs substantially from the traditional IPTV, and the rich interactions of comment posting and monetary virtual gift donating provide ways for people to better understand how crowdsourced live videos are produced and consumed.

In this paper, we carry out a large-scale measurement study on Douyu, one of the most popular crowdsourced livecast websites in China, and present an in-depth analysis of the viewer interactions. Our data spans four months. We have collected over 7.5 million livecast sessions from about 0.7 million broadcasters, and more than one billion comments posted by over 14 million viewers. The data also contains nearly 65 million virtual gifts worth more than 6.8 million US dollars donated in four weeks.

With the measurement data, we study the viewer interactions of posting comments and donating monetary virtual gifts, and investigate the implicit popularities of livecast channels that are derived from the interactions. We develop models for three important aspects of the interactions, namely the gift donating process, the viewer activity, and the channel's popularity. Rather than simply relying on curve fitting, we further interpret the model parameters with real-world meanings. Our contributions in this work are summarized as follows:

- We carry out a content analysis on danmu comments, and find that the time-synchronized comments can be effectively exploited to capture the highlights in live videos.

- We find that livecast has become people's regular entertainment choice; we also find that livecast sessions experience "cold start", and viewers have strong preferences in their interactions.
- Big gift donations in a livecast session can be modeled as a discrete-time stochastic process, the process's first donating time and inter-donating interval can be well captured with the Weibull distributions, and the distributions' shape factors suggest a "Lindy effect" in gift donating.
- Viewer activities of posting comments and donating monetary virtual gifts are highly skewed, and fit well with the stretched exponential (SE) distributions, with the distributions' stretch factors reflecting the "effort" that a viewer takes to perform the interactions; however, the two viewer activities are only weakly correlated.
- Livecast channels' explicit and implicit popularities that are derived from the viewer interactions can be well captured with the SE distributions, with the distributions' stretched factors reflecting how much "effort" a viewer needs to take to consume the content via different interactions; however, the implicit and explicit channel popularities are only weakly to moderately correlated.
- A channel's popularity is influenced by a number of broadcaster behavioral factors, and it is feasible to accurately predict the implicit popularity of a livecast channel from the behavioral factors and its explicit popularity.

As far as we know, our work is the first comprehensive measurement on viewer interactions in a large-scale crowdsourced livecast service. Our analysis, based on the large volume of measurement data, has several implications. First, the highly skewed viewer activity suggests that differentiated video delivery services could be provided to viewers regarding their contributions, especially to the ones donating many monetary virtual gifts. Second, the SE distributions of implicit and explicit channel popularities and their correlations suggest that the conventional popularity in term of simultaneous viewers has its limitations, for example, for increasing the financial income, channels with higher gift-popularity should be given priority in allocating service resources. Finally, our analysis and proposed methodologies for predicting individual channel's popularity are valuable for the livecast website, individual broadcasters, as well as third parties such as advertisers, to assess their business policies.

The rest part of this paper is organized as follows: Section II discusses the related works. We introduce Douyu and our measurement data in Section III. We present a brief content analysis on danmu comments in Section IV. Section V analyzes the interaction patterns and models the big gift donating process. We investigate viewer activity and channel's popularity in Section VI. Section VII analyzes the factors that influence a channel's popularity, presents and evaluates the popularity predicting methodologies. Finally, we conclude in Section VIII.

II. Related Work

There is a rich literature on tracing and characterizing user behaviors in Internet video streaming services [4] [2] [5] [6] [7] [9]. Yu *et al.* [4] mine the video access logs of a large-scale video-on-demand (VoD) system, and find that video popularity matches a Zipf-like distribution. Li *et al.* [6] analyze the user accesses to a VoD service from WiFi and 3G connections, and find that individual users' video views follow a stretched exponential distribution, with the stretch factor correlating with user's connection type. Chen *et al.* [7] investigate a large VoD system, and reveal a number of relations between user watching time and various video features such as type, duration, and popularity.

For understanding user behaviors in live video streaming. Qiu *et al.* [2] investigate the channel access data from Set-Top Boxes (STBs), and find that TV channel's popularity can be well captured by a Zipf-like distribution. Li *et al.* [5] analyze the view patterns of mobile users in a large mobile TV service, and report that channel popularity in term of total segment downloads follows a Pareto distribution. Li *et al.* [9] investigate access logs of a large IPTV system, and develop a series of models for session duration, user activity, and time dynamics of user arrivals/departures. However, most previous works are focused on viewers' watching behaviors, as the systems under study do not allow other viewer interactions beyond the watching. We believe that by analyzing the rich interactions such as comment posting and virtual gift donating, people can better understand how videos are produced and consumed, and improve the services accordingly.

With crowdsourced livecast becoming popular in recent years, studies were carried out for investigating such a novel service. Zhang *et al.* [10] examine the message and streaming latencies in Twitch channels, and study the view patterns that are influenced by both events and video sources. Wang *et al.* [11] investigate the performances of two popular applications, namely Periscope and Meerkat, and analyze the causes of streaming latencies. Many works study livecast channel's popularity in term of simultaneous viewers. Kaytoue *et al.* [12] show that a video game tournament channel's popularity can be accurately predicted from its early stage. Deng *et al.* [13] investigate the high churn in channel popularity caused by game tournaments broadcasted on Twitch. Pires *et al.* [14] compare YouTube Live and Twitch, and find that their channel popularities are more heterogeneous than what have been observed on other user-generated contents. Jia *et al.* [15] also find that game broadcasters on Twitch have highly skewed popularity, with a significant heavy tail phenomenon. For improving crowdsourced livecast systems, He *et al.* [16] propose a cloud-assisted architecture and resource allocating algorithms for improving viewer satisfactions; Ma *et al.* [17] present a suite of distributed algorithms for synchronizing cross-viewer community interactions.

Only a few recent works were focused on the rich interactions in video streaming applications. He *et al.* [18] study the danmu comments in a VoD service, and use the comment volume to predict video popularity. Zhang *et al.* [19] analyze the touch operations in mobile gamecasting to predict viewer's gazing patterns for smartphone energy saving. Zhu *et al.* [20] present a statistical study on virtual gift values received by broadcasters on Douyu. Jia *et al.* [21] analyze the social features of the videos on a YouTube-like website, and use

Fig. 1. An example of a livecast channel on Douyu. Danmu messages are sliding from right to left; comments posted by viewers are displayed in colored texts, and the message starting with the "shark riding on a rocket" icon (in blue circle) is a system-wide message indicating a big gift donating event; virtual gifts (in red circle) under different prices are displayed at bottom of the video screen for viewers to purchase.

them to predict whether a video is popular or not. As far as we know, our work is the first comprehensive measurement study on viewer interactions in a large-scale crowdsourced livecast service. Based on our measurement data, we derive models and present in-depth analysis on several important issues in user behaviors and channel access patterns, and provide insights for both system design and business policy.

III. Background and Data

A. Background

`Douyu.com` is a popular crowdsourced livecast website in China. Similar to Twitch, on Douyu, any registered user can set up a *channel* and broadcast live videos on the Internet. Although initially focused on games, Douyu allows a wide range of contents to be broadcasted. For example, a *broadcaster* can sing, rap, tell stories and jokes, share experiences and opinions in her channel, or she can broadcast outdoor activities such as fishing, hiking, sightseeing to her audience. Fig. 1 demonstrates the web interface of a livecast channel[1], which is featured with good-looking broadcaster singing popular songs.

As many Chinese video websites, Douyu enhances video streaming with *danmu*. Danmu is a word originated from Japanese, literally meaning "bullet curtain", and it refers to a commentary sharing mechanism in which a *viewer* can paste *comments* directly on top of the video while watching it. As shown in Fig. 1, when the "view danmu" toggle is turned on, a viewer can see danmu comments from all the viewers watching a same channel sliding over his video screen in real time. Danmu is very popular in China, as by interacting with broadcaster and other viewers, a viewer may feel more engaged in the broadcasted activity [22].

Besides posting comments, Douyu also allows viewers to donate *monetary virtual gifts* to broadcasters through the danmu mechanism. As shown in Fig. 1, when watching a livecast channel, a viewer can purchase virtual gifts, whose prices range from 0.1 to thousands of RMB yuan, to donate to the broadcaster. In fact, virtual gift donations are an important financial income source for both individual broadcasters and the Douyu website.

Note that comment posting and virtual gift donating are not unique on Douyu. In fact, nearly all the crowdsourced livecast websites allow viewers to comment, and recently, many mainstream websites start to allow monetary donations. For example, since 2017, Twitch allows viewers to purchase "bits", a monetary virtual goods, and donate to broadcasters[2]; YouTube Live also allows viewers to purchase "super chats" with real money to highlight their comments[3].

TABLE I
Summaries of the Datasets

Dataset	*4MONTH*	*4WEEK*
Duration	124 days	28 days
Livecast channels	716, 425	242, 697
Livecast sessions	7, 554, 875	1, 789, 027
Viewers	14, 075, 214	7, 482, 937
Danmu messages	1, 425, 801, 906	257, 352, 841
Comments	1, 400, 128, 173	250, 291, 347
Gift donations		64, 894, 747
Total gift value		USD 6, 807, 524.94

B. Dataset

In this work, we select Douyu as an example, and study the viewer interactions of posting comments and donating virtual gifts with a measurement approach. In our measurement, we maintain a list of active livecast channels by repeatedly crawling Douyu's portal web page every five minutes. The page lists all the active channels that currently have on-going livecast *sessions*. For each channel in the list, we collect the metadata such as the channel ID, content type, and its current viewer number. We also run a probing process, which employs Douyu's RESTful APIs [23] to subscribe to its danmu server, receives and logs all the real-time danmu messages in the channel. Specifically, for each livecast channel, two kinds of messages were recorded:

- **Comment message**: this message contains the comment text posted by a viewer in the channel, as well as the timestamp and ID of the comment posting viewer.
- **Gift donating message**: this message indicates a gift donating event in the channel, and contains the information such as the timestamp, the donating viewer's ID, and the gift ID.

Our measurement started on Sep. 9, 2016, and lasted 124 days. From our measurement, two datasets are obtained:

- The first dataset covers all the 124 days, and records over 7.5 million livecast sessions and 1.4 billion danmu messages. By mining the data, we have identified more than 700 thousand unique broadcasters and over 14 million distinct viewers. We refer to this dataset as *4MONTH*.
- The second dataset covers 28 days from Nov. 22, 2016 to Dec. 19, 2016, which is a subset of the 4MONTH

[1]In this paper, we use the terms "broadcaster" and "channel" interchangeably.

[2]https://help.twitch.tv/customer/portal/articles/2449458
[3]https://support.google.com/youtube/answer/7277005

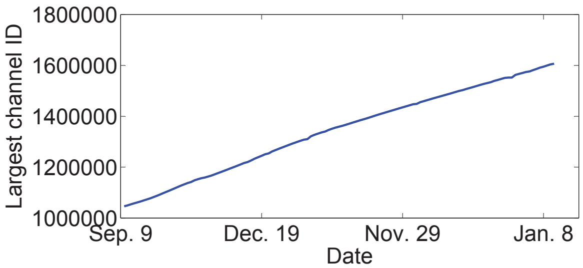

Fig. 2. Largest channel ID encountered by each day in the 4MONTH dataset.

dataset in time. However, in this dataset, we managed to trace monetary prices of all the donated virtual gifts. As a result, we have traced nearly 65 million pieces of virtual gifts, which are worth over 6.8 million US dollars[4], being donated. We refer to this dateset as *4WEEK*. Table I lists summaries of the two datasets.

We exploit a livecast channel's ID to estimate the scale of our measurement data and growth of Douyu's crowdsourced livecast service. We note that each livecast channel has a numerical ID, which we speculate is sequentially assigned by Douyu. To testify, we label each day in the 4MONTH dataset as day_1, day_2, $\cdots$, day_{124}, and for each day, say day_i, we find the largest channel ID that appears in $[day_1, \cdots day_i]$. Fig. 2 presents the largest channel ID encountered by each day in the dataset, from which we can see that it increases linearly with time. The figure confirms our conjecture that Douyu sequentially assigns numerical ID to channels; moreover, the largest channel ID in the dataset is 1606770, suggesting that 44.6% and 15.1% of the broadcasters on Douyu are included in our 4MONTH and 4WEEK datasets respectively. From Fig. 2, we also estimate that on average, $4,707$ broadcasters were newly registered every day, indicating that Douyu was expanding very fast during our measurement period. As far as we know, our measure data is the largest one on viewer interactions in crowdsourced livecast services.

IV. Content Analysis on Danmu Comments

A. Length and Sentiment

With the assistance of the measurement data, we first carry out a content analysis on danmu comments. We notice that danmu comments are very short: the averaged length of the comments in the 4MONTH dataset is 6.45 characters. We then employ the `Jieba` tokenizer [24], a powerful Chinese word segmentation package, to segment words from the comments. We find that averagely, a danmu comment contains only 3.8 words, which is also very short comparing with other text-based social medias (e.g., tweet).

We are wondering what the viewers are saying in such short danmu comments. By manually inspecting a number of samples, we find that unlike tweets and review comments, danmu comments are usually not long sentences, but only a few words expressing the sentiment feelings (like or dislike) of the viewers towards the broadcasters or broadcasted videos. To verify our point, for each danmu comment in the dataset, we compute its *sentiment value* as the following: We first segment words from the comment using the `Jieba` tokenizer, then we query each word in sentiment dictionaries to find out its *orientation*, i.e., positive (+1) or negative (-1). We employ two widely-used Chinese sentiment dictionaries, namely `BosonNLP` [25] and `NTUSD` [26], and extend them with the common slang terms on Douyu. Finally, we compute the comment's *sentiment value* as the sum of the orientations of all its words. Fig. 3 presents the distribution of the sentiment values of the comments in the 4WEEK dataset. From the figure we can see that most comments fall in the range of $[-2, +3]$, and about 40% comments have a sentiment value of 1. The observation confirms that most danmu comments are short sentences that express viewers' sentiment feelings.

B. Highlight Labeling

One salient feature of danmu comments is that they are synchronized with video contents. Recently, time-sync comments have been exploited to produce video summaries [27] and label highlight video shots [28] for stored videos in offline ways. In this subsection, we explore if the time-sync danmu comments can be explored to label highlights for live videos in real time.

We present a simple methodology for our purpose. For a live video, we divide its content into segments, with each segment lasting one-minute playback duration. For example, the video segments in the last hour can be denoted as $S = [-59, -58, \cdots, -, 0]$. For each segment, say segment i, we compute the segment values of all the danmu comments that were posted during its playback duration, using the method as previously discussed. We consider a danmu comment with a positive/negative sentiment value as a positive/negative vote, and estimate the viewers' overall sentiment feeling towards the video segment by computing the segment's *sentiment score* $s(i)$ with Wilson's Score Interval [29] as

$$s(i) = \frac{\hat{p} + \frac{1}{2n} z_{1-\alpha/2}^2 - z_{1-\alpha/2} \sqrt{\frac{\hat{p}(1-\hat{p})}{n} + \frac{z_{1-\alpha/2}^2}{4n^2}}}{1 + \frac{1}{n} z_{1-\alpha/2}^2} \qquad (1)$$

Note that the above formula not only takes $\hat{p}$, the faction of the positive votes out of the total votes, but also takes n, the total votes, into consideration, and we let $z_{1-\alpha/2} = 1.96$ be the variate value from the standard normal distribution for 0.95 percentile.

Fig. 4 presents an example of the sentiment scores of the segments in a livecast session. One can see that there are some peaks that are abnormally higher than other segments, we consider these segments as *candidate highlight segments*. More specifically, we employ Twitter's abnormal detecting package, which is based on the Seasonal Hybrid ESD (S-H-ESD) algorithm [30], to detect if the current segment (segment 0) is a highlight candidate: if the segment's sentiment score is abnormally high comparing with the ones in the last hour, we label it as a highlight candidate. For example, in Fig. 4, we have labeled 7 highlight candidates in the entire session.

We apply our method on 20 livecast sessions from 8 livecast channels, and have labeled 112 highlight candidates. We also

[4]Virtual gifts are priced with RMB, however in this paper, we fix the rate as 1 US dollar = 6.9 RMB yuan.

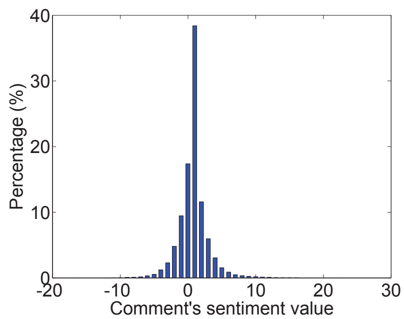

Fig. 3. Distribution of the sentiment values of the danmu comments.

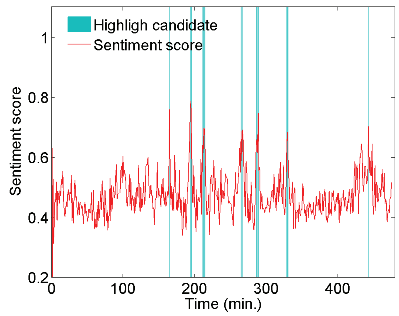

Fig. 4. An example of sentiment scores and highlight candidates in a livecast session.

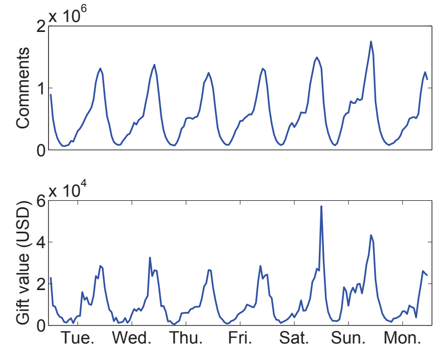

Fig. 5. Weekly patterns of viewer interactions.

record the live videos and ask three volunteers to watch the video to decide whether a labeled candidate is a true video highlight. Among the 112 labeled candidates, as many as 95 were considered as true highlights by all the three volunteers. The result suggests that our methodology is quite accurate in labeling highlights for crowdsourced live videos.

V. Interaction Patterns

A. Temporal patterns

In this section, we investigate the patterns of the viewer interactions, based on our measurement data from Douyu.

1) Daily and week patterns: We first look at the temporal patterns of the viewer interactions. In the top and bottom figures of Fig. 5, we present volume of the comments posted and value of the virtual gifts donated by viewers in each hour of a week in the 4WEEK dataset respectively. From the figures, we can see that Douyu is large in scale, as the website hosts thousands of concurrent livecast channels, and attracts virtual gifts worth tens of thousands of US dollars in peak hours. Moreover, the figures suggest a clear weekly pattern, as more interactions happen in weekends than weekdays.

We then investigate viewers' activities in each hour of a day. By bucketing comments in the 4MONTH dataset and gifts in the 4WEEK dataset in each hour of a day's 24 hours, we obtain the averaged comment volume and gift value in each hour, and plot them in the top figure of Fig. 6. For comparison, we also plot broadcasters' activities regarding the averaged ongoing sessions and the aggregated broadcasting time in each hour in Fig. 6's bottom figure. From the figures we can see that both viewers' and broadcasters' activities exhibit a strong diurnal pattern, and they are highly synchronized, with the Pearson's correlation coefficients between any two of the four activity data series higher than 0.85. However, we find that the peak hour of the viewer activity regarding both comment posting and gift donating appears in 23:00∼24:00, while the broadcasters are most active in 20:00∼21:00 (in term of ongoing sessions) and 21:00∼22:00 (in term of broadcasting time). In other words, viewers are staying up later at night than broadcasters. Such a mismatch can be explained with the fact that most viewers watch livecast at home, however, for some broadcasters, such as the outdoor broadcasters, they can not stay as late as the viewers.

We note that the daily and weekly patterns of the viewer and broadcaster activities on Douyu are similar to the conventional TV watching patterns [2], but are different from early studies on Twitch [10] [14]. For example, we do not observe that the overall viewer activity can be highly influenced by some events, such as e-sports competitions as in [14]. Our observation suggests that crowdsourced livecast websites such as Douyu have already become people's regular entertainment choice, rather than a platform for broadcasting special events.

2) Patterns within livecast sessions: In this subsection, we analyze how viewers' interaction patterns evolve within a livecast session. To this end, we divide all the sessions in our dataset into 2-minute intervals, and for each interval, we compute the averaged amount of comments posted and gift value donated during the interval from all the sessions that contain the interval, and normalize the values per hour. Fig. 7 presents the results, in which we only plot the interaction intensities within the first ten hours after a session starts, since very few sessions in our datasets last over ten hours.

From the figure we can see that there are relatively fewer comments posted and virtual gifts donated very soon after a session starts. The interaction intensities increase and reach to a stable level after several hours, then decline slowly over time. The observation suggests that in general, a broadcaster will experience a "cold start", as she can not attract many viewers right after starting to broadcast; a session needs to last long enough for accumulating viewers, who visit Douyu and the livecast channel sporadically from time to time.

B. Preference

As we have seen in Fig. 5, Douyu concurrently hosts a large number of livecast channels. One natural question is, how many broadcasters a viewer chooses to interact with? To answer this question, we rank all the viewers in the 4WEEK dataset in two descending orders: one is based on the comments they have posted, and the other is based on the value of their donated virtual gifts. In each order, we compute the averaged number of the broadcasters a viewer interacts with, for all the viewers ranked at $1 \sim 100$, $101 \sim 1000$, $1001 \sim 10000$, $10001 \sim 100001$, and the remaining positions in the order. Fig. 8 presents the results. From the figure we can see that the most active viewers post comments in dozens of channels, but the number decreases as the viewers

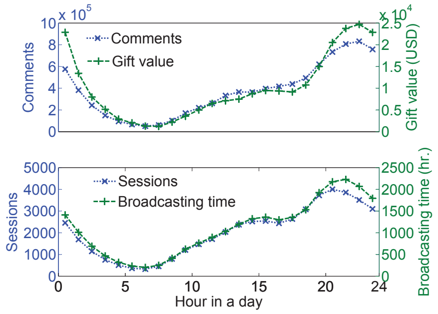

Fig. 6. Diurnal patterns of viewers' and broadcasters' activities.

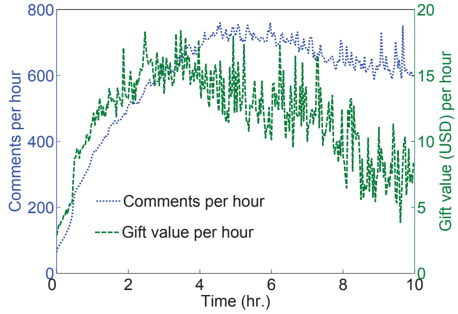

Fig. 7. Interaction patterns within a livecast session.

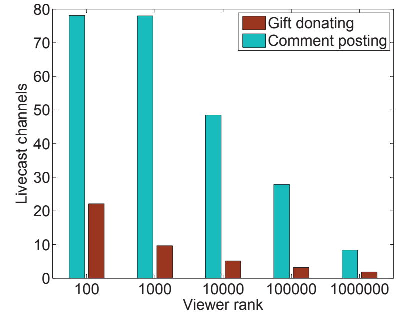

Fig. 8. Preferences of viewers in comment posting and gift donating interactions.

TABLE II
AMOUNT AND VALUE OF THE GIFTS UNDER DIFFERENT PRICES BEING DONATED IN THE 4WEEK DATASET.

Price (RMB)	< 1	$1 \sim 99$	≥ 100
Amount	$64,352,895$	$375,105$	$166,747$
Value (USD)	$1,164,896.3$	$291,422.3$	$5,351,206.3$

become less active. Comparing with comment posting, viewers donate virtual gifts much more conservatively: even the top-100 generous viewers donate to less than 25 broadcasters on average, and as many as 59.5% viewers donate in only one channel. The observation suggests that, viewers have strong preferences in their interactions, and they are much more selective in donating virtual gifts than posting comments, as the former requires more "effort" from the viewers by costing them real money.

C. Gift donating

Different from the subscription-based business model in the traditional TV industry, livecast services on Douyu are free, but viewers are encouraged to donate monetary virtual gifts to their preferred broadcasters. Douyu's business model is very successful, as the website can receive gifts worth up to tens of thousands US dollars in an hour. In this subsection, we present an in-depth analysis on gift donating.

1) Big gifts: Douyu provides virtual gifts under various prices for viewers to purchase and donate. In Table II, we list the numbers and total value of the virtual gifts within different price ranges in the 4WEEK dataset. From the table we can see that, 99.2% of the donated virtual gifts have prices less than 1 RMB yuan; on the other hand, 78.6% of the gift value is from the gifts with prices no less than 100 RMB yuan, although they constitute less than 0.3% of the gifts. Note that Douyu refer to the gifts with prices higher than 100 RMB yuan as *big gifts*, and when such a gift is donated, the donation is announced with a system-wide danmu message in all the active livecast channels. For example, in Fig. 1, the message starting with the "shark riding on a rocket" icon announces the event that a big gift named "rocket", whose price is 500 RMB yuan, is just donated by a viewer in a channel as indicated in the message text. Since big gifts constitute 78.6% of the total gift value, in the following part of this section, we focus on big gifts, and study the timing of their donations.

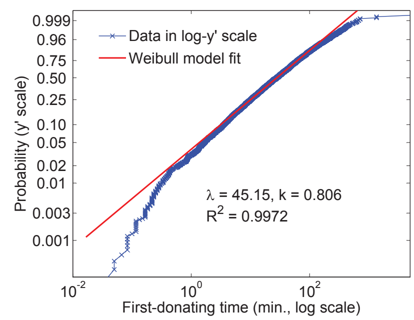

Fig. 9. Distribution of the first-donating time t_1 and its Weibull model fitting under the $log\text{-}y'$ scale.

2) Modeling big gift donating: On Douyu, virtual gifts can only be donated during on-going livecast sessions, thus we can model the big gift donations within a session as a discrete-time stochastic process. More specifically, for a session in which d big gifts have been donated, we denote the *big gift donating process* as $P = \{t_i | i = 1, 2, \cdots, d\}$, where t_i is the time of the i^{th} big gift donation since the session starts. We are interested in two properties of the process: 1) the *first-donating time* t_1, and 2) the *inter-donating interval* $\tau_i = t_{i+1} - t_i$ between two consecutive donations. The properties reflect how much effort a broadcaster needs to take (by broadcasting continuously) to receive the first and subsequent big gifts.

To investigate the big gift donating process, we select all the livecast sessions in the 4WEEK dataset that contain at least three big gift donations (i.e., $d \geq 3$). The selected sessions cover 80.4% of the big gift donations. By examining the sessions, we find that it is much more difficult for a broadcaster to attract the first big gift than the subsequent ones, as the mean first-donating time, which is $\overline{t_1} = 51.7$ minutes, is much longer than the mean inter-donating interval $\overline{\tau_i} = 19.3$ minutes. Note that our observation here conforms to the "cold start" phenomenon as observed in Fig. 7.

We then consider the distributions of the first-donating time t_1 and the inter-donating interval τ_i. We try to fit the empirical data with various models such as exponential, lognormal, and Weibull distributions using the maximum-likelihood estimation (MLE), and find that both data can be best fitted with the *Weibull distributions*, whose cumulative distribution function

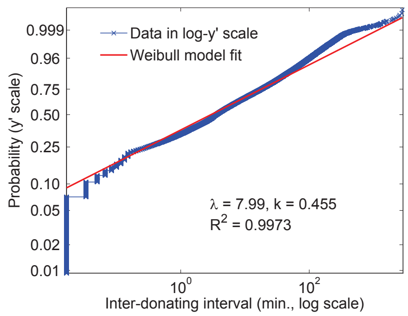

Fig. 10. Distribution of the inter-donating interval τ_i and its Weibull model fitting under the log-y' scale.

TABLE III
PARAMETERS OF THE FITTED WEIBULL DISTRIBUTIONS FOR BIG GIFT DONATING UNDER VARIOUS FILTERING CONDITIONS.

	First-donating time	Inter-donating interval
$d = 3$	$\lambda = 45.15, k = 0.806$	$\lambda = 7.99, k = 0.455$
$d = 4$	$\lambda = 42.11, k = 0.809$	$\lambda = 7.34, k = 0.458$
$d = 5$	$\lambda = 35.02, k = 0.811$	$\lambda = 6.26, k = 0.459$

(CDF) can be expressed as

$$F(t) = \Pr[X \le t] = 1 - e^{-\left(\frac{t}{\lambda}\right)^k} \tag{2}$$

where λ and k are the distribution's scale and shape parameters respectively. Note that by letting $y' = \ln(-\ln(1 - F(t)))$, Equation (2) can be re-written as

$$y' = k \ln t - k \ln \lambda \tag{3}$$

In other words, the distribution curve can be plotted as a straight line under the log-y' scale.

We plot the empirical distributions of the first-donating time t_1 and the inter-donating interval τ_i under the log-y' scale in Fig. 9 and Fig. 10 respectively. We also plot the fitted Weibull distribution models, and present their shape and scale parameter values on the figures. To evaluate the goodness of the fit, we compute the *coefficient of determination* R^2, which measures the proportion of total variation of the data explained by the model, and find that the R^2 values are close to one, suggesting that the Weibull distributions explain the data closely

We note that the two Weibull distributions have their shape factors k smaller than one, indicating that their hazard functions are monotonically increasing with the elapsed time t. In other words, for a broadcaster to receive the first or a subsequent big gift, there exists a "Lindy effect": the longer a broadcaster has waited for the gift, the longer she is expected to wait in future [31].

Finally, we examine the robustness of the Weibull distribution model. We change our filtering condition by considering the livecast sessions with no less than $d = 4$ or 5 big gift donations respectively, and fit the obtained first-donating times and inter-donating intervals to Weibull distributions. The model fitting results are given in Table III. From the table, we can see that the data can still be modeled with the Weibull distributions very well. Moreover, the fitted distributions only differ at the scale parameter λ, while the shape parameter k, which dictates how skewed the distribution is, doesn't change substantially. The observation suggests that, Weibull distribution is robust for modeling the big gift donating process; furthermore, the distribution's shape parameter, which does not change over samples, reflects the intrinsic hardness of a broadcaster to receive the first or a subsequent big gift donation within a livecast session.

VI. VIEWERS AND BROADCASTERS

In this section, we characterize viewers' activities in comment posting and monetary virtual gift donating interactions, and analyze channels' popularities, especially the implicit popularities derived from the viewer interactions.

A. Viewer activity

1) Modeling viewer activities: We first study viewers' interaction activities. Using the 4MONTH dataset, we compute the *skewness coefficient* of the comments posted by individual viewers, which is defined as

$$\gamma = \mathrm{E}\left[\left(\frac{X - \mu}{\sigma}\right)^3\right] \tag{4}$$

where μ and σ as the mean and standard deviation of the comments posted by a viewer. The coefficient value is 68.0. Such a large positive skewness coefficient suggests that the data is biased towards a few viewers, who have posted much more comments than average. For example, the top 10% viewers contribute as many as 75.9% comments in the 4MONTH dataset.

We also investigate gift donations using the 4WEEK dataset. We find that only 14.5% of the viewers have ever donated virtual gifts, and the gift donating activity is even more skewed than comment posting: the top 10% viewers contribute as much as 92.5% gift value, and the skewness coefficient is as high as 587.7.

For further understanding the viewers' comment posting acrivity, we rank all the viewers in the 4MONTH dataset in a descending order, and present the rank distribution in Fig. 11(a). We first seek to fit the distribution with a power-law model, which has been observed in users' lifetime tweets on Twitter [32]. However, the data fits the power-law model poorly, as we can see that the distribution curve under the log-log scale (the x- and the right y-axles) doesn't appear as a straight line.

On the other hand, a recent study [33] shows that on many social medias, users' behaviors of generating contents can be well captured with a stretched exponential (SE) distribution, whose definition is

$$\Pr[X \ge x] = e^{-\left(\frac{x}{x_0}\right)^c} \tag{5}$$

where x_0 and c are constants. In the comment posting case, if we rank the viewers in a descending order according to the comments they have posted, then the rank distribution function can be expressed as

$$y_i^c = -a \log i + b \tag{6}$$

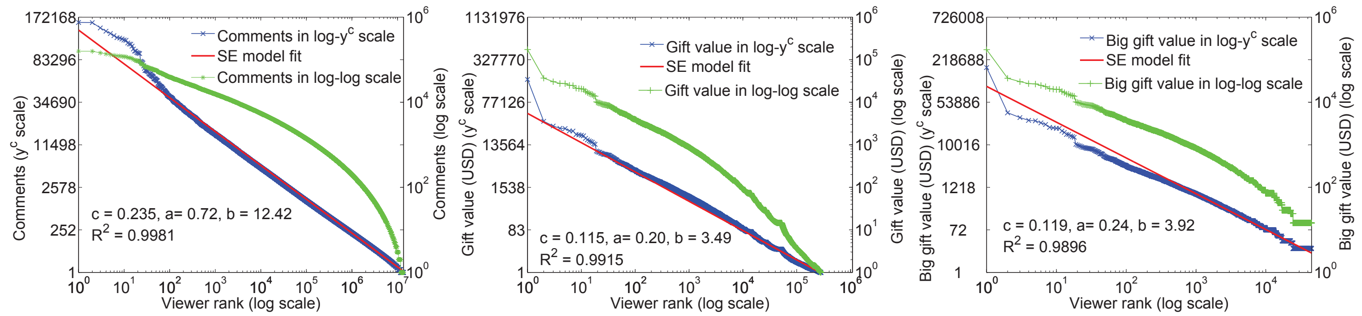

(a) Comment posting (b) Gift donating (c) Big gift donating

Fig. 11. Rank distributions and model fittings of (a) comments posted, (b) gift value, and (c) big gift value donated by viewers, under the log-y^c (x- and left y-axles) and log-log (x- and right y-axles) scales.

where y_i is the amount of the comments posted by the i^{th}-ranked viewer, $a = x_0^c$ and $b = y_1^c$. In other words, the curve of viewers' comment posting rank distribution should be a straight line under the log-y^c scale.

We use the MLE-based method in [33] to fit the comment posting rank distribution with the SE model, and present the result in Fig. 11(a). Note that the figure's left y-axle is under the y^c scale. From the figure we can see that the distribution curve appears as a straight line, suggesting that the SE model fits the viewers' comment posting data very well. We also compute the *coefficient of determination* R^2 of the fitting, and find that its value is close to 1.

We then examine viewers' gift donating activity. For the $1,087,452$ viewers who have donated virtual gifts in the 4WEEK dataset, we present their rank distribution regarding the donated gift value under the log-y^c and log-log scales in Fig. 11(b). From the figure, we can see that the data doesn't follow a power-law distribution, but can be captured with an SE distribution closely ($R^2 = 0.9915$). In Fig. 11(c), we present the rank distribution of the big gifts donated by viewers, and again we can fit the data with an SE distribution closely. We note that the parameters of the SE distributions for gift and big gift donations are very close, suggesting that for the viewers ranked high in gift donating, most of their donated gifts are big gifts, therefore the distribution in Fig. 11(c) overlaps with the head part of the distribution in Fig. 11(b), and they have close model parameter values.

By comparing the three viewer activity distributions, we can see that for comment posting, its SE distribution's stretch factor, i.e., parameter c in Equation (5), is much larger than the ones of the models for gift and big gift donating. Recall that [33] relates the stretched factor to the "effort" that is needed to generate a content: the more effort it requires, and smaller the factor is. We can see that such an interpretation also applies to our observation, as it takes more "effort" for a viewer to donate a monetary virtual gift, which cost him real money, than posting a short text comment.

2) Activity correlations: As a viewer can post comments and donate virtual gifts to interact with a broadcaster, one question is, are the two activities correlated? To answer this question, we examine the $958,909$ viewers in the 4WEEK dataset who have performed both interactions, and compute the *Pearson's correlation coefficient* between them. We also compute the *Spearman's rank correlation coefficient* between the ranks of the viewers in the two activities, which is defined as

$$\rho = 1 - \frac{6\sum_{i=1}^{n}(c_i - g_i)^2}{n(n^2-1)} \tag{7}$$

where n is total number of the viewers, c_i and g_i are the positions of viewer i in the comment posting and gift donating ranks respectively. We find that viewers' two interaction activities are only weakly correlated, as the Spearman's correlation is 0.381, and the Pearson's correlation, which is further influenced by the highly skewed data, is only 0.108. The low correlations suggest that in a crowdsourced livecast service like Douyu, viewers are highly diverse in their activities: an active commenter is not necessarily a generous gift donor, and vice versa.

B. Channel popularity

In tradition, popularity of a TV channel is measured with access frequency and dwell time [2], and crowdsourced livecast websites such as YouTube Live and Twitch use simultaneous viewers as a channel's popularity [10] [14]. All these measures depend only on viewers' watching activity. On the other hand, on a crowdsourced livecast website like Douyu, viewers can perform rich interactions rather than simply watching the channel. In the following, we analyze a livecast channel's implicit popularity based on how viewers interact with the broadcaster in posting comments and donating virtual gifts.

1) Implicit and explicit channel popularities: We first define the implicit popularity of a channel that is derived from the viewer interactions. Specifically, for a livecast channel, we define its *comment-popularity* as the comments posted by its audience in the channel during a unit broadcasting time (e.g., an hour). Similarly, we define a channel's *gift-popularity* as the value of the gifts donated by viewers in the channel during a unit broadcasting time. For clarity, we refer to the explicit popularity in term of simultaneous viewers as *viewer-popularity* in this work.

We employ our measurement data to investigate the implicit and explicit popularities. More specifically, among all the livecast channels in the 4WEEK dataset, we compute non-zero comment-popularities for $123,286$ channels, and non-zero gift-popularities for $43,385$ channels. We also exploit the

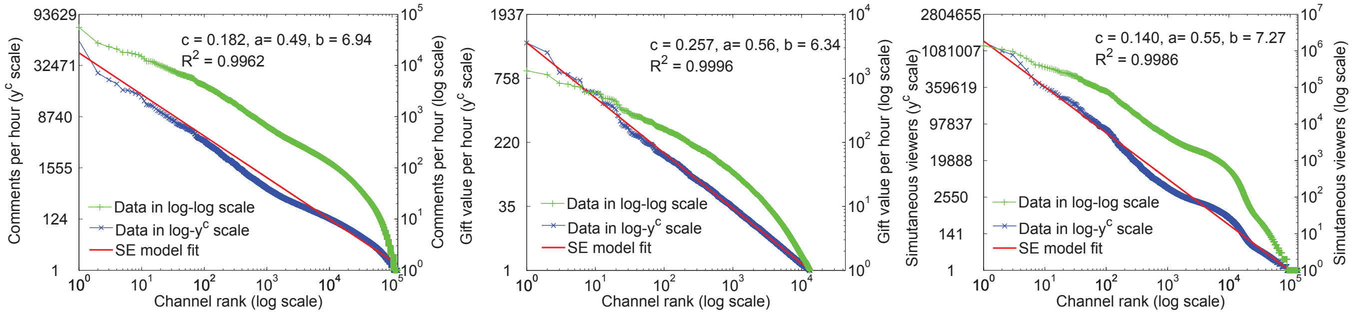

(a) Comment-popularity (b) Gift-popularity (c) Viewer-popularity

Fig. 12. Rank distributions and model fittings of (a) comment-popularity, (b) gift-popularity, and (c) viewer-popularity of livecast channels, under the log-y^c (x- and left y-axles) and log-log (x- and right y-axles) scales.

viewer numbers displayed on channels' web pages, and obtain non-zero viewer-popularities for $133,173$ channels. We find that the channels vary significantly in all the three popularity measures, as the skewness coefficients of the comment-, gift-, and viewer-popularities are as large as 62.5, 23.5, and 91.4 respectively.

2) Modeling channel popularity: Popularities of live and on-demand videos are usually observed to follow Zipf-like distributions [4] [2] [10] [14] [34], while a recent study [35] reports that the reference rank distributions of a number of Internet media objects should be modeled with stretched exponential (SE) distributions rather than Zipf-like distributions. To investigate which model fits the channel popularities, we plot the rank distributions of the comment-, gift-, and the conventional viewer-popularities of the channels in Fig. 12(a), (b), and (c) respectively. In each figure, we use the method in [33] to fit the data to the SE model, estimate the distribution parameters, and present the distribution curves under the log-y^c scale (left y-axle). We also fit the data to Zipf-like distributions by plotting the distributions under the log-log scale (right y-axle) for comparison.

The figures in Fig. 12 indicate that, the rank distributions of the implicit and explicit channel popularities do not follow Zipf-like models, as the curves are not straight lines under the log-log scale; however, the empirical data is well captured by the SE distributions, as the curves appear as straight lines under the log-y^c scale for all the three popularity measures, and the R^2 values, which measure goodness of the fits, are close to one.

We note that the stretch factor c of the SE model for the gift-popularity, which is 0.257, is the largest among the three distributions, then followed by the factor of the comment-popularity (0.182), while the viewer-popularity has the smallest stretch factor (0.140). The authors of [35] relate the stretch factor c with the sizes of the Internet media objects: the larger the object is, the greater the factor will be. However, [35] considers only the stored media objects (e.g., files, movies), but a livecast channel does not have a "content size". To explain the stretch factors of the SE popularity models, we extend the interpretation in [35] by relating the factor c with the amount of "effort" a user needs to take for consuming the content: the more "effort" is required, the greater the stretch factor will be. With the extended interpretation, we can explain the positive correlation between the media object size and the stretch factor value in [35], as it generally takes a user more "effort" (e.g., downloading and watching time) to consume a larger media object than a smaller one. The extended interpretation also explains our observations on channel popularities, as a viewer can post a comment or donate a virtual gift only when he has already taken the "effort" to watch the channel, so the viewer-popularity has the smallest stretch factor; in addition, it takes more "effort" for a viewer to donated a monetary virtual gift than posting a text comment, therefore the gift-popularity has a greater stretch factor than the comment-popularity.

TABLE IV
PEARSON'S AND SPEARMAN'S CORRELATION COEFFICIENTS AMONG THREE CHANNEL POPULARITIES.

Popularity1	Popularity2	Pearson	Spearman
comment	gift	0.429	0.641
comment	viewer	0.637	0.687
gift	viewer	0.222	0.540

3) Popularity correlations: Our analysis shows that crowdsourced livecast channels have implicit and explicit popularities. A question one may ask is: are these popularity metrics closely related, and can we use one popularity to replace the others? To answer this question, in Table IV we compute the Pearson's correlation and Spearman's rank correlation coefficients between any two of the three channel popularities. From the table we can see that, the three popularities are only weakly to moderately correlated, indicating that no single popularity measure can replace the others. We believe that the three popularities have their different usages. For example, when allocating streaming bandwidth to channels, the viewer-popularity should be considered; but the comment-popularity reflects the overall audience activeness of a channel, which is important for online advertisement; finally, the gift-popularity is closely related to the financial profitabilities of individual broadcasters and the livecast website.

C. Discussion

Our study in this section shows that viewers differ significantly in activities of posting comments and donating virtual gifts, which means that for a crowdsourced livecast website

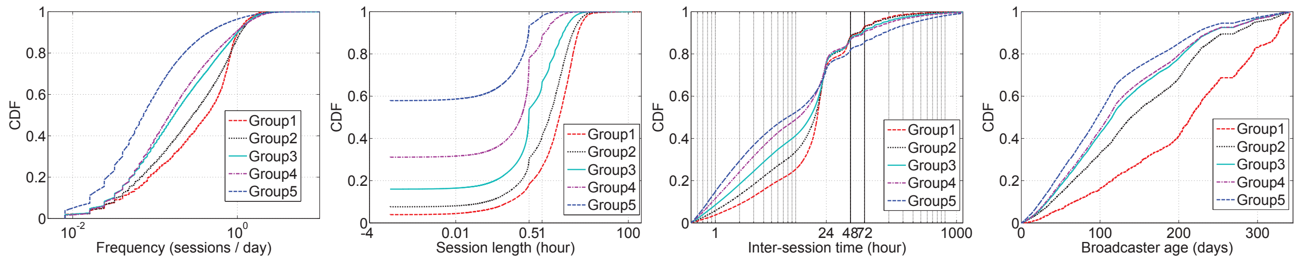

(a) Broadcasting frequency (b) Session length (c) Inter-session time (d) Broadcaster age

Fig. 13. Distributions of (a) broadcasting frequencies, (b) session lengths, (c) inter-session times, and (d) ages of the broadcasters in different groups.

like Douyu, some viewers, in particular, those donating lots of virtual gifts, are much more valuable than others. On the other hand, existing crowdsourced livecast systems treat all the viewers indifferently when delivering live videos to them [13].

One potential improvement for the crowdsourced livecast system is to provide those valuable viewers a differentiated video delivery service. For example, the system may schedule better (but may be more expensive) network paths for streaming live videos to valuable viewers. It is reasonable to expect that when being provided with a better quality of experience (QoE), the valuable viewers will better enjoy the service and reciprocate with more monetary virtual gift donations.

Our analysis also reveals that livecast channels have implicit as well as explicit popularities, and they are only weakly to moderately correlated. On the other hand, most existing crowdsourced livecast systems allocate service resources among the channels based on the conventional viewer-popularity. For example, it is observed that Twitch only transcodes videos to multiple representations for the channels that have most simultaneous viewers [16].

We suggest that the interaction-derived implicit channel popularities should be considered when allocating system's service resources. For example, when allocating the transcoding servers, the channels with higher gift-popularity should be given priority. This is because by providing more video representation options to viewers, especially those valuable viewers, channels with higher gift-popularity are likely to receive more monetary virtual gift donations from the viewers with improved satisfaction levels [36].

VII. Predicting Channel Popularity

Our previous analysis and discussion show that the implicit and explicit channel popularities have different implications. In this section, we investigate the factors that influence the popularities of livecast channels, and explore the feasibility of accurately predicting them.

A. *Factors influencing channel popularity*

We first investigate the factors that influence a livecast channel's popularity. We recognize that the main reason for a channel to be welcomed by viewers is the attractiveness of its video content. However, it is difficulty to objectively assess a channel's content attractiveness, as different viewers may perceive a same content differently. In this work, we focus on the factors of the broadcaster's behaviors. More specifically, for each channel, we consider the following factors: 1) *broadcasting frequency*; 2) *session length*; 3) *inter-session time*; and 4) *age of the broadcaster* on the website.

We employ the 4MONTH dataset to study the above four behavioral factors. For each channel, the second and third factors can be easily obtained; the broadcaster's age can be estimated from the channel's numerical ID, as it is estimated that averagely, $4,707$ broadcasters were newly registered on Douyu each day in Section III-B; for estimating a broadcaster's broadcasting frequency, we divide the number of her sessions with the days that her lifetime, which is estimated from the channel ID, overlaps with our measurement period.

For differentiating the broadcasters, we rank all the channels with non-zero comment-popularity in a descending order, and cluster them into five groups such that group1 represents the most popular broadcasters and group5 represents the least popular ones, and each group has approximately equal aggregated comment-popularity. Table V lists the size as well as the mean values of the four behavioral factors for each group.

From Table V, we can see that broadcasters of different comment-popularities differ significantly in their broadcasting behaviors. To further understand the differences, we plot the distributions of the broadcasting frequencies, session lengths, inter-session times, and broadcaster ages of the five groups in Fig. 13. Fig. 13(a) shows that popular broadcasters tend to broadcast more frequently than unpopular ones, for example, 10.9% of the broadcasters in group1 broadcast at least once every day on average, but this ratio drops to 3.6% for group5.

Fig. 13(b) presents the session length distributions of the broadcaster groups, from which we can see that popular broadcasters tend to have longer sessions, as 73.5% of the sessions in group1 last over one hour, but this ratio drops to less than 2.0% for group5. We also find that except for group1, there are considerable sessions in the other four groups lasting exactly 30 minutes. Further investigation shows that it is because Douyu rewards a broadcaster if she continuously broadcasts over 30 minutes, therefore many unpopular broadcasters stop their sessions immediately after the threshold. Obviously, the existence of so many 30-minute sessions is a good indication of lack of popularity for the channels in these groups.

In Fig. 13(c), we plot the distributions of the inter-session times of the broadcasters in the five groups. We make two observations from the figure: First, for all the five groups, the percentages of the inter-session times within 24 hours are

TABLE V
SIZES AND MEAN BEHAVIORAL FACTOR VALUES OF THE FIVE BROADCASTER GROUPS.

	Group1	Group2	Group3	Group4	Group5
Num. of broadcasters	$1,379$	$8,879$	$20,685$	$46,257$	$286,793$
Mean broadcasting frequency (sessions / day)	0.498	0.456	0.369	0.344	0.193
Mean session length (hr.)	3.38	2.23	1.16	0.47	0.16
Mean inter-session time (hr., within 24 hours)	13.10	11.57	9.96	8.31	6.80
Mean broadcaster age (days)	205.1	148.8	129.4	126.0	111.0

approximately same, which are about 70%, regardless of the popularity. However, the more popular a group is, the larger proportion of its inter-session times are close to 24 hours. We also observe similar (but much weaker) phenomena of the inter-session times around 48 and 72 hours on the figure. The observation suggests that popular broadcasters tend to schedule their broadcasting activities regularly on daily basis; on the other hand, unpopular broadcasters tend to broadcast more casually.

The second observation is that the inter-session times of different broadcaster groups have diverse characteristics within and beyond 24 hours. When only considering inter-session times within 24 hours, we are surprised to find that popular broadcasters tend to have longer inter-session times than unpopular ones. For example, as indicated in Table V, broadcasters in group1 have the longest mean inter-session time of 13.1 hours, while it is only 6.8 hours for group5, which is the shortest among all the five groups. Note that such an observation is counter-intuitive and seems to contradict with the fact that popular broadcasters broadcast more frequently than unpopular ones. However, after further examination, we find that many unpopular broadcasters make a few sessions within a short time, then never broadcast again; on the other hand, as we have observed, popular broadcasters have regular schedules for their broadcasting activities. For example, we find that 24.0% of the broadcasters in group5 have all their sessions in only one day, while this ratio drops to 5.3% in group1.

Fig. 13(d) presents the age distributions of the broadcasters in the five groups. We can see that popular broadcasters are generally older than unpopular ones. For example, we find 64.7% of the broadcasters in group5 are registered within our 124-day measurement period, while it is only 18.6% in group1.

Note that although Fig. 13 only illustrate the influences of the broadcaster's behavioral factors on comment-popularity, however, we have made similar observations on the other two channel popularity metrics, that is, the viewer-popularity and gift-popularity. We omit the results due to space limitation.

We conclude our findings as the following: popular broadcasters are *diligent* broadcasters, who broadcast longer and more frequently; popular broadcasters are also *regular* broadcasters, who have well-planned schedules for their broadcasting activities, while unpopular broadcasters tend to broadcast in an impromptu way; finally, popular broadcasters are *persistent* broadcasters, who accumulate audiences and even fanbases on Douyu for a longer time.

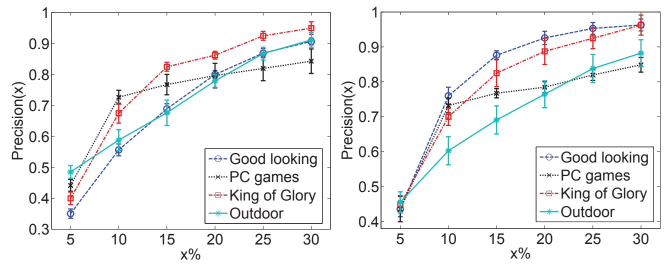

(a) Comment-popularity predicting (b) Gift-popularity predicting

Fig. 14. $Precision(x)$ of (a) comment-popularity predicting and (b) gift-popularity predicting for channels of different types under various $x\%$.

B. Predicting danmu-popularity

In this and next subsections, we explore the feasibility of accurately predicting channel's implicit popularities. We first seek to predict a channel's comment-popularity based on its conventional viewer-popularity and the broadcaster's behavioral factors, that is, the broadcasting frequency, mean session length, mean inter-session time, and the broadcaster's age. Note that for inter-session time, we only consider the ones within 24 hours, due to the reason as previously discussed. We believe that the predicting problem is practical, as some crowdsourced livecast websites do not provide APIs for collecting comment posting interactions, thus it is infeasible for any third parties to obtain a livecast channel's comment-popularity directly.

Instead of using one predictor for all the channels, we consider channel's content type, and develop one predictor for each type. In our experiment, we select four representative channel types, namely "good looking", "PC games", "King of Glory" (a popular mobile game), and "outdoor" from the 4WEEK dataset, and use channels of these types as the samples for training and evaluating the predicting algorithms. We list the number of the channels of each type in Table VI.

The popularity predicting problem is a regression problem, and we explore three representative regression algorithms, namely linear regression, regression tree, and random forrests [37] for solving it. For training and evaluating the algorithms, we randomly select 80% of the channel samples as the training set, and use the remaining 20% as the test set. Each experiment is repeated 10 times. For measuring the predicting accuracy, we use the *normalized mean squared error* (*NMSE*), which compares the predicting error with the error of the baseline approach that uses the global mean of the training set as the predictions, that is,

$$NMSE = \frac{\sum_i (p_i - a_i)^2}{\sum_i (g - a_i)^2} \tag{8}$$

TABLE VI
PREDICTING ACCURACIES OF REGRESSION ALGORITHMS, FOR DIFFERENT TYPES OF LIVECAST CHANNELS.

Channel type			Good looking	PC games	King of Glory	Outdoor
Num. of channels			$5,150$	$2,199$	997	904
Comment-popularity predicting	Linear regression	NMSE	0.504	0.481	0.316	0.412
		Correlation	0.707	0.691	0.785	0.749
	Regression tree	NMSE	0.787	1.070	0.879	0.532
		Correlation	0.628	0.538	0.716	0.691
	Random forrests	NMSE	**0.432**	**0.459**	**0.254**	**0.334**
		Correlation	**0.759**	**0.749**	**0.871**	**0.819**
Gift-popularity predicting	Linear regression	NMSE	0.452	0.480	0.360	0.522
		Correlation	0.748	0.779	0.800	0.702
	Regression tree	NMSE	0.657	0.758	1.772	1.371
		Correlation	0.663	0.712	0.586	0.440
	Random forrests	NMSE	**0.361**	**0.368**	**0.332**	**0.450**
		Correlation	**0.799**	**0.795**	**0.821**	**0.745**

TABLE VII
IMPORTANCE SCORES OF THE FEATURES USING RANDOM FORRESTS, FOR DIFFERENT TYPES OF LIVECAST CHANNELS.

Channel type		Good looking	PC games	King of Glory	Outdoor
Comment-popularity predicting feature importance	Broadcasting frequency	0.058	0.050	0.031	0.059
	Inter-session time	0.091	0.053	0.034	0.047
	Session length	**0.503**	**0.556**	0.264	0.156
	Broadcaster age	0.092	0.056	0.049	0.081
	Viewer-popularity	0.255	0.285	**0.622**	**0.657**
Gift-popularity predicting feature importance	Broadcasting frequency	0.049	0.041	0.053	0.040
	Inter-session time	0.061	0.059	0.117	0.067
	Session length	0.087	0.077	0.101	0.056
	Comment-popularity	**0.598**	**0.646**	**0.569**	**0.613**
	Broadcaster age	0.074	0.067	0.081	0.068
	Viewer-popularity	0.131	0.109	0.079	0.156

where a_i is the actual comment-popularity of channel i in the test set, p_i is channel i's comment-popularity predicted by the algorithm, and g is the global mean of the samples in the training set. We also compute the Pearson's correlation coefficient between the predicted and the ground-truth comment-popularities for evaluating the predicting accuracy.

We implement the three predicting algorithms using the `scikit-learn` Python package [38], and evaluate with the samples from our measurement data. The comment-popularity predicting results are given in Table VI. From the table we can see that the random forests (RF) algorithm outperforms the other two algorithms, as it constantly has the lowest NMSE and the highest correlation coefficients among the three. More specifically, with RF, the NMSEs for all the channel types are lower than 0.5, suggesting that the algorithm significantly improves the predicting accuracies over the baseline approach. In addition, the correlations between the predictions and the ground truths are $0.749 \sim 0.871$, which are much higher than the ones in Table IV.

The better accuracy of RF over regression tree is easy to understand, as RF employs a collection of regression trees rather than a single tree for learning, thus can significantly reduce the errors caused by bias and variance in the training samples [37]. The better accuracy of RF over linear regression can be explained with the fact that the relations between the comment-popularity and the behavioral factors are not strictly linear, as we have seen in Section VII-A.

We further examine the RF-predictors, as they make the most accurate predictions. Table VII presents the importance scores of the features in predicting. We can see that viewer-popularity is an important feature, but session length also has high scores, and for the types of "good looking" and "PC games", it is indeed the most useful feature.

In some cases, people may only want to identify the most popular broadcasters instead of predicting their exact popularity values. To this end, we consider the problem of how well the RF-predictors can identify the channels with the highest comment-popularity. More specifically, we label the top 10% channels regarding their comment-popularity in the test set as *popular channels*, and compute the percentage of the popular channels that fall in the top $x\%$ of our predictions. We refer to such a percentage as $Precision(x)$. Formally,

$$Precision(x) = \frac{\text{Popular channels in top } x\% \text{ of the predictions}}{\text{All the popular channels}} \tag{9}$$

Fig. 14(a) presents $Precision(x)$ of the predictions for the four types of channels under various $x\%$. From the figure we can see that $Precision(x)$ increases rapidly with $x\%$, and for all the content types, as many as $84 \sim 95\%$ popular channels fall in the top 30% of our predictions, indicating that the RF-predictors can precisely identify the most popular channels.

C. Predicting gift-popularity

In this subsection, we consider the problem of predicting a channel's gift-popularity based on its viewer- and comment-popularities, as well as the broadcaster's behavioral factors. Our motivation is that some crowdsourced livecast websites do not allow or fully allow virtual gift donations, for example, both Twitch and YouTube Live do not allow monetary donations until 2017, and YouTube Live doesn't allow monetary donations in channels with less than $10K$ subscribers. For such a website, an interesting question is: given a broadcaster's broadcasting behaviors and her viewer- and comment-popularities, how much money the website and the individual broadcaster can make if virtual gift donating is fully allowed.

The gift-popularity predicting problem is also a regression problem, and we explore the algorithms of linear regression, regression tree, and RF for solving it. We consider all the behavioral factors, i.e., broadcasting frequency, mean session length, mean inter-session time (within 24 hours), and broadcaster's age, as well as the channel's viewer- and comment-popularities as the features, and use the channels of the four types in Table VI for training and evaluating the predictors.

As in the previous subsection, we use the metrics of NMSE and the Pearson's correlation coefficient for evaluating the predicting accuracies, and present them in Table VI. From the table we can see that the predictions made by the RF-predictors are much more accurate than the baseline approach, and the correlations between the predictions and the ground truths are much higher than the ones in Table IV. Moreover, as in the comment-popularity predicting, RF outperforms the other two algorithms. We also list the features' importance scores with the RF-predictors in Table VII, and find that comment-popularity, which is also derived from viewer interactions, is the most important feature, while the conventional viewer-popularity is also very useful. Finally, we present $Precision(x)$ of the predictions for the four channel types in Fig. 14(b), from which we can see that the RF-predictors can precisely identify the most profitable broadcasters, as $84 \sim 96\%$ of the top 10% most profitable broadcasters fall in the top 30% of our predictions.

In summary, in this section we analyze a number of broadcaster's behavioral factors that influence a channel's popularity. Based on our analysis, we present machine-learning based methodologies for predicting a channel's comment- and gift-popularities. Using the measurement data, we show that our proposed approach can predict channels' popularities accurately.

VIII. Conclusion

Rich interactions in crowdsourced live video broadcasting services, such as posting comments and donating monetary virtual gifts, have greatly improved viewer experiences and brought huge incomes to both website and individual broadcasters. In this paper, we tracked, analyzed, and modeled viewer interactions in a large scale crowdsourced livecast website. Our results showed that danmu comments help to capture video highlights, livecast sessions experience "cold start", and viewers have strong preferences in their interactions. We found that timing of the big gift donations can be modeled with Weibull distributions, and both the viewer activity and the channel popularity can be captured with stretched exponential distributions closely. We also analyzed the influence of a broadcaster's behavioral factors on channel popularity, and present machine-learning based methodologies for popularity prediction.

Our analysis provides important implications for improving crowdsourced livecast systems, and we also provide insights for different players in the crowdsourced livecast ecosystem, including the website, individual broadcasters, and third-party advertisers, to assess their business policies. For future works, we seek to extend our measurement study on other mainstream platforms like YouTube Live, and improve the content delivery and resource allocation mechanisms in livecast services based on viewers' interaction characteristics.

References

[1] M. DiPietro, "Twitch is 4th in peak us internet traffic," Feb. 2014, https://blog.twitch.tv/twitch-is-4th-in-peak-us-internet-traffic-90b1295af358.

[2] T. Qiu, Z. Ge, S. Lee, J. Wang, Q. Zhao, and J. Xu, "Modeling channel popularity dynamics in a large IPTV system," in *Proc. of SIGMETRICS'09*, Seattle, WA, USA, June 2009.

[3] S. Zhou, Q. Wang, J. Ge, and Y. Tian, "Peer-assisted video streaming with RTMFP flash player: A measurement study on PPTV," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 28, no. 1, pp. 158–170, 2018.

[4] H. Yu, D. Zheng, B. Y. Zhao, and W. Zheng, "Understanding user behavior in large scale video-on-demand systems," in *Proc. of EuroSys'06*, Leuven, Belgium, Apr. 2006.

[5] Y. Li, Y. Zhang, and R. Yuan, "Measurement and analysis of a large scale commercial mobile internet TV system," in *Proc. of IMC'11*, Berlin, Germany, Nov. 2011.

[6] Z. Li, J. Lin, M.-I. Akodjenou, G. Xie, M. A. Kaafar, Y. Jin, and G. Peng, "Watching videos from everywhere: a study of the PPTV mobile VoD system," in *Proc. of IMC'12*, Boston, MA, USA, Nov. 2012.

[7] Y. Chen, B. Zhang, Y. Liu, and W. Zhu, "Measurement and modeling of video watching time in a large-scale Internet video-on-demand system," *IEEE Trans. Multimedia*, vol. 15, no. 8, pp. 2087–2098, 2013.

[8] H. He, Y. Zhao, J. Wu, and Y. Tian, "Cost-aware capacity provisioning for internet video streaming CDNs," *The Computer Journal*, vol. 58, no. 12, pp. 3255–3270, 2015.

[9] Z. Li, M. A. Kaafar, K. Salamatian, and G. Xie, "Characterizing and modeling user behavior in a large-scale mobile live streaming system," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 27, no. 12, pp. 2675–2686, 2017.

[10] C. Zhang and J. Liu, "On crowdsourced interactive live streaming: A Twitch.TV-based measurement study," in *Proc. of NOSSDAV'15*, Portland, OR, USA, Mar. 2015.

[11] B. Wang, X. Zhang, G. Wang, H. Zheng, and B. Y. Zhao, "Anatomy of a personalized livestreaming system," in *Proc. of IMC'16*, Santa Monica, CA, USA, Nov. 2016.

[12] M. Kaytoue, A. Silva, L. Cerf, W. M. Jr., and C. Raïssi, "Watch me playing, I am a professional: A first study on video game live streaming," in *Proc. of WWW'12*, Lyon, France, Apr. 2012, pp. 1181–1188.

[13] J. Deng, G. Tyson, F. Cuadrado, and S. Uhlig, "Internet scale user-generated live video streaming: The Twitch case," in *Proc. of PAM'17*, Sydney, Australia, Mar. 2017.

[14] K. Pires and G. Simon, "YouTube Live and Twitch: A tour of user-generated live streaming systems," in *Proc. of MMSys'15*, Portland, OR, USA, Mar. 2015.

[15] A. L. Jia, S. Shen, D. Epema, and A. Iosup, "When game becomes life: The creators and spectators of online game replays and live streaming," *ACM Trans. Multimedia Comput. Commun. Appl.*, vol. 12, no. 4, 2016.

[16] Q. He, J. Liu, C. Wang, and B. Li, "Coping with heterogeneous video contributors and viewers in crowdsourced live streaming: A cloud-based approach," *IEEE Trans. Multimedia*, vol. 18, no. 5, pp. 916–928, 2016.

[17] X. Ma, C. Zhang, J. Liu, R. Shea, and D. Fu, "Live broadcast with community interactions: Bottlenecks and optimizations," *IEEE Trans. Multimedia*, vol. 19, no. 6, pp. 1184–1194, 2017.

[18] M. He, Y. Ge, L. Wu, E. Chen, and C. Tan, "Predicting the popularity of danmu-enabled videos: A multi-factor view," in *Proc. of International Conference on Database Systems for Advanced Applications (DASFAA'16), LNCS 9643*, Dallas, TX, USA, Apr. 2016, pp. 351–366.

[19] C. Zhang, Q. He, J. Liu, and Z. Wang, "Beyond the touch: Interaction-aware mobile gamecasting with gazing pattern prediction," in *Proc. of IEEE INFOCOM'17*, Atlanta, GA, USA, May 2017, pp. 1–9.

[20] Z. Zhu, Z. Yang, and Y. Dai, "Understanding the gift-sending interaction on live-streaming video websites," in *Proc. of International Conference on Social Computing and Social Media*, Vancouver, BC, Canada, Jul. 2017, pp. 274–285.

[21] A. L. Jia, S. Shen, D. Li, and S. Chen, "Predicting the implicit and the explicit video popularity in a user generated content site with enhanced social features," *Elsevier Computer Networks*, vol. 140, pp. 112–125, 2018.

[22] O. L. Haimson and J. C. Tang, "What makes live events engaging on Facebook Live, Periscope, and Snapchat," in *Proc. of CHI'17*, Denver, CO, USA, May 2017.

[23] "Douyu third-party developper api document v2.1," Apr. 2016, http://dev-bbs.douyutv.com/forum.php?mod=viewthread&tid=113.

[24] J. Sun, "Jieba tokenizer v0.33," Sep. 2018, https://github.com/fxsjy/jieba.

[25] "Bosonnlp: Open platform for chinese semantic," Sep. 2018, https://bosonnlp.com/.

[26] NLPSA, "NTUSD - NTU sentiment dictionary," Sep. 2018, http://academiasinicanlplab.github.io/.

[27] L. Xu and C. Zhang, "Bridging video content and comments: Synchronized video description with temporal summarization of crowdsourced time-sync comments," in *Proc. of AAAI'17*, San Francisco, CA, USA, Feb. 2017, pp. 1611–1617.

[28] G. Lv, T. Xu, E. Chen, Q. Liu, and Y. Zheng, "Reading the videos: Temporal labeling for crowdsourced time-sync videos based on semantic embedding," in *Proc. of AAAI'16*, Phoenix, AZ, USA, Feb. 2016, pp. 3000–3006.

[29] S. Wallis, "Binomial confidence intervals and contingency tests: Mathematical fundamentals and the evaluation of alternative methods," *Journal of Quantitative Linguistics*, vol. 20, no. 3, pp. 178–208, 2013.

[30] J. Hochenbaum, O. S. Vallis, and A. Kejariwal, "Automatic anomaly detection in the cloud via statistical learning," *arXiv preprint arXiv:1704.07706*, 2017.

[31] C.-D. Lai, *Generalized Weibull Distributions*. Springer, Berlin, Heidelberg, 2014.

[32] D. R. Bild, Y. Liu, R. P. Dick, Z. M. Mao, and D. S. Wallach, "Aggregate characterization of user behavior in Twitter and analysis of the retweet graph," *ACM Trans. Internet Technol.*, vol. 15, no. 1, pp. 1–24, 2015.

[33] L. Guo, E. Tan, S. Chen, X. Zhang, and Y. E. Zhao, "Analyzing patterns of user content generation in online social networks," in *Proc. of KDD'09*, Paris, France, Jun. 2009.

[34] K. Sripanidkulchai, B. Maggs, and H. Zhang, "An analysis of live streaming workloads on the Internet," in *Proc. of IMC'04*, Chicago, IL, USA, Nov. 2004.

[35] L. Guo, E. Tan, S. Chen, Z. Xiao, and X. Zhang, "The stretched exponential distribution of internet media access patterns," in *Proc. of PODC'08*, Toronto, Canada, Aug. 2008.

[36] L. Toni, R. Aparicio-Pardo, K. Pires, G. Simon, A. Blanc, and P. Frossard, "Optimal selection of adaptive streaming representations," *ACM Trans. Multimedia Comput. Commun. Appl.*, vol. 11, no. 2s, 2015.

[37] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed. Springer, 2009.

[38] "scikit-learn: Machine learning in Python," Jul. 2018, http://scikit-learn.org/.

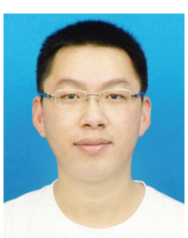

Xiaodong Wang received the bachelor's degree in computer science from the Nanjing University of Aeronautics and Astronautics in 2015. He is currently a Ph.D. candidate at the School of Computer Science and Technology, University of Science and Technology of China, Hefei, China. His research is focused on multimedia networks and Internet measurement. He is a student member of the IEEE.

Ye Tian received the Ph.D. degree from the Department of Computer Science and Engineering, The Chinese University of Hong Kong, Hong Kong, in 2007. He joined the University of Science and Technology of China (USTC) in 2008 and is currently an Associate Professor with the School of Computer Science and Technology, USTC. His research interests include multimedia network, Internet measurement, software-defined networking and future Internet architecture. He has published over 70 papers and co-authored a research monograph published by Springer. He is the winner of the Wilkes Best Paper Award of the Oxford Computer Journal in 2016. He is currently serving as an Associate Editor of the Springer Frontiers of Computer Science Journal. He is a member of the IEEE.

Rongheng Lan received the bachelor's degree in statistics from the University of Science and Technology of China (USTC), Hefei, China, in 2016. He is currently working toward the master's degree at the School of Information Science and Technology, USTC, Hefei, China. His research interest is focused on Internet big data.

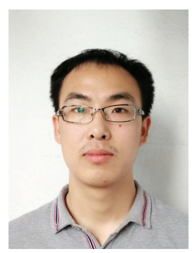

Wen Yang received the bachelor's degree in computer science from Liaoning University. He is currently pursuing the mater's degree with the School of Computer Science and Technology, University of Science and Technology of China, Hefei, China. His research interest is focused on Internet measurement.

Xinming Zhang received the B.E. and M.E. degrees in electrical engineering from China University of Mining and Technology, Xuzhou, China, in 1985 and 1988, respectively, and the Ph.D. degree in computer science and technology from the University of Science and Technology of China, Hefei, China, in 2001. Since 2002, he has been with the faculty of the University of Science and Technology of China, where he is currently a Professor with the School of Computer Science and Technology. From September 2005 to August 2006, he was a visiting Professor with the Department of Electrical Engineering and Computer Science, Korea Advanced Institute of Science and Technology, Daejeon, Korea. His research interest includes wireless networks. He has published more than 100 papers in networking. He won the second prize of Science and Technology Award of Anhui Province of China in Natural Sciences in 2017. He is a senior member of the IEEE.