When Game Becomes Life: The Creators and Spectators of Online Game Replays and Live Streaming

ADELE LU JIA, College of Information and Electrical Engineering, China Agricultural University, China
SIQI SHEN, Parallel and Distributed Processing Laboratory, National University of Defense Technology, China; School of Computer, National University of Defense Technology, China; Software and Computer Technology Department, Delft University of Technology, the Netherlands
DICK H. J. EPEMA and ALEXANDRU IOSUP, Software and Computer Technology Department, Delft University of Technology, the Netherlands

Online gaming franchises such as World of Tanks, Defense of the Ancients, and StarCraft have attracted hundreds of millions of users who, apart from playing the game, also socialize with each other through gaming and viewing gamecasts. As a form of User Generated Content (UGC), gamecasts play an important role in user entertainment and gamer education. They deserve the attention of both industrial partners and the academic communities, corresponding to the large amount of revenue involved and the interesting research problems associated with UGC sites and social networks. Although previous work has put much effort into analyzing general UGC sites such as YouTube, relatively little is known about the gamecast sharing sites. In this work, we provide the first comprehensive study of gamecast sharing sites, including commercial streaming-based sites such as Amazon's Twitch.tv and community-maintained replay-based sites such as WoTreplays. We collect and share a novel dataset on WoTreplays that includes more than 380,000 game replays, shared by more than 60,000 creators with more than 1.9 million gamers. Together with an earlier published dataset on Twitch.tv, we investigate basic characteristics of gamecast sharing sites, and we analyze the activities of their creators and spectators. Among our results, we find that (i) WoTreplays and Twitch.tv are both fast-consumed repositories, with millions of gamecasts being uploaded, viewed, and soon forgotten; (ii) both the gamecasts and the creators exhibit highly skewed popularity, with a significant *heavy tail* phenomenon; and (iii) the upload and download preferences of creators and spectators are different: while the creators emphasize their individual skills, the spectators appreciate team-wise tactics. Our findings provide important knowledge for infrastructure and service improvement, for example, in the design of proper resource allocation mechanisms that consider future gamecasting and in the tuning of incentive policies that further help player retention.

CCS Concepts: • **Networks** → **Network measurement**; • **Social and professional topics** → *User characteristics;*

Additional Key Words and Phrases: Online game communities, gamecast sharing sites, repository characteristics, popularity dynamics, user behaviors

---

This work was partially supported by the National Science Foundation for Young Scholars of China (NSFYSC) No. 61502500, by the National Basic Research Program of China (Grant No. 2014CB340303), NSFYSC No. 61303064, by TU Delft, by the NWO/STW Veni grant large (11881), and by the COMMIT NL project P20.
Authors' addresses: S. Shen (corresponding author), College of Computer, National University of Defense Technology, China; email: shensiqi@nudt.edu.cn; A. L. Jia, College of Information and Electrical Engineering, China Agricultural University, China; email: ljia@cau.edu.cn; D. H. J. Epema and A. Iosup, Department of Software and Computer Technology, Delft University of Technology, the Netherlands; emails: {D.H.J.Epema, A.Iosup}@tudelft.nl.

DOI: http://dx.doi.org/10.1145/2957750

ACM Reference Format:
Adele Lu Jia, Siqi Shen, Dick H. J. Epema, and Alexandru Iosup. 2016. When game becomes life: The creators and spectators of online game replays and live streaming. ACM Trans. Multimedia Comput. Commun. Appl. 12, 4, Article 47 (August 2016), 24 pages.
DOI: http://dx.doi.org/10.1145/2957750

1. INTRODUCTION

Online games are today entertaining hundreds of million of users and form a multi-billion-dollar global industry [McGonigal 2011]. Similar to professional sports such as football, the community involved in the activity includes not only amateur and professional players, but also a large group of spectators. Together, players and spectators watch and even study gamecasts (the equivalent of recorded or streamed broadcasts in sports) of talented teams or individual players, sometimes repeatedly, for entertainment or educational purpose. Numerous online communities provide the opportunity for players to share their gamecasts. These communities often archive millions of gamecasts that are watched and commented upon by millions of users. With the huge user base and the large revenue involved, gaming communities have attracted the attention of many industrial magnates. Among the leading communities, Amazon acquired Twitch.tv in 2014 and YouTube recently launched its own game streaming site.

In this work, we investigate two online gaming communities, WoTreplays [2013] and Twitch [2011], which are leading representatives for the two most popular classes of gamecast sharing sites. WoTreplays is community-maintained and *replay*-based. In just 2 years of operation, it has archived more than 380,000 replays of games played by nearly 2 million players and has attracted more than 5 million downloads. Twitch is a leading commercial community for game *live streaming*, wherein players broadcast and comment on their gamecasts in video streams to spectators who, if they enjoy the gamecast, can follow the channel, give a heart, or further make a donation. This new form of communication introduces new relationships between Twitch players and their spectators. In February 2014, Twitch became the fourth largest source of US peak Internet traffic [Fitzgerald and Wakabayashi 2014].

Our analysis of WoTreplays and Twitch mainly focuses on characterizing the repository and on analyzing user activities. We do this for two reasons:

First, repository characteristics play an important role in infrastructure and service improvement; for example, in designing proper resource allocation mechanisms based on content popularity. Although much research has already focused on characterizing general content sharing sites [Cha et al. 2009; Yu et al. 2006; Pouwelse et al. 2005; Abrahamsson and Nordmark 2012], little is known about the gamecast sharing sites' user behaviors. Given the rapid increase in the popularity and market size of online games, we believe it is important and timely to characterize these sites.

Second, user activities provide important knowledge for maintaining community prosperity, for example, through customizing incentive policies based on user preference. Compared to general UGC sharing sites that often cover a wide range of topics and that lack explicit information on contents, gamecast sharing sites have the advantages that they are exclusively for gaming content and that many communities archive detailed game statistics such as the winning team and the in-game scores of each player. These game statistics provide fine-grained information for inferring user activities, including their upload and download preferences.

In this work, we conduct an in-depth analysis of gamecast sharing sites starting from two real-world, large-scale datasets. Our analysis reveals the basic characteristics of the repositories and user behaviors. We summarize our main contributions as follows:

(1) We collect, use, and offer public access[1] to the dataset that contains the full history of WoTreplays, with detailed statistics for 1,956,256 gamers (including their download and upload behaviors) and 382,760 games (including team formation, game result, and the reward each player obtained). We further include in our analysis an earlier published dataset on Twitch [Pires and Simon 2015], which archives more than 5 million game streaming sessions (Section 3).
(2) We investigate and compare the basic characteristics of WoTreplays and Twitch. Our analysis includes (i) the repository scale, including the gamecast injection rate and duration, and (ii) the statistical properties of the gamecast popularity, including the *skewness*, *heavy tail*, and *long tail* phenomena (Section 4).
(3) We provide a detailed analysis of user activities, including the activity level, upload delay, creator popularity, and interactions between users and the repository. With the detailed in-game statistics archived by WoTreplays, we further investigate the upload and download preferences of its users (Sections 5 and 6).

2. RELATED WORK

In this section, we compare our work with the characterization research on gamecast sharing sites, game workload, general content sharing sites, and content popularity.

Gamecast sharing site characterization. Gamecast sharing is a relatively unexplored area. Kaytoue et al. [2012] and Pires and Simon [2015] provide preliminary characterizations on Twitch.tv. They analyze the dynamics of game spectators and propose models for predicting video popularity. Our previous work [Shen and Iosup 2011] analyzes XFire, a social network of games and players. The analysis focuses on the global network, gaming activity, and the social structure in XFire, with preliminary results on UGC. In this work, we discover more than triple the amount of players and multiple orders of magnitude more replays than in Shen and Iosup [2011] and Kaytoue et al. [2012]. Most importantly, we complement these studies with a finer-grained dataset: The game statistics included in our WoTreplays dataset provide a better reference for understanding user activities, including their upload and download preferences.

Although not much quantitative analysis has been performed, gamecast sharing has been studied qualitatively. Cheung and Huang [2011] provide a qualitative account of the experiences of StarCraft II spectators and find nine personas in the data that demonstrate who the spectators are and why they watch. Downs et al. [2013] study the core aspects of audience experience in social video gaming. Hamilton et al. [2014] find that informal communities emerge from Twitch streams where users socialize with each other.

Game workload characterization. The popularity of the spectator activity of a game largely depends on the game's popularity. Understanding the workload of games and players can help us better understand the workload of gamecasts and develop better systems supporting gamecast sharing sites. A significant body of work studies the workload of First Person Shooter (FPS) games and Massively Multiplayer Online Games (MMOGs) [Suznjevic and Matijasevic 2013]: Henderson and Bhatti [2001] analyze the player population and network packets of FPS games. Armitage et al. [2006] study the FPS games' client round-trip time and hop-count distributions. Chambers et al. [2010] characterize player and session distribution of FPS games and MMOGs. Recently, cloud gaming systems such as OnLive, Gaikai, and StreamMyGame, have received much attention: Claypool et al. [2014] analyze the workload of OnLive. Chen et al. [2014] study the quality of service of OnLive and StreamMyGame. Differing from the research focusing on game workloads, we focus on the workloads of gamecast sharing sites.

[1]The dataset is released in the Game Trace Archive [Guo and Iosup 2012] http://gta.st.ewi.tudelft.nl/.

General content sharing site characterization. General content sharing sites, wherein the shared content covers a wide range of topics, have been extensively studied. Cheng et al. [2008] find YouTube videos have noticeably different statistics compared to traditional streaming videos. Gill et al. [2007] analyze YouTube traffic generated by a collection of clients and provide a detailed view of the local UGC service usage. Cha et al. [2007, 2009] provide a complementary global view by crawling data of complete sets of video categories. Qu et al. [2008] and Huang et al. [2008] provide a survey on peer-to-peer live streaming systems.

Content popularity characterization. One major contribution of this work is a detailed statistical analysis of gamecast popularity. Little related work has dealt with this topic. Kaytoue et al. [2012] provide a preliminary study, and they identify skewed gamecast popularity in Twitch. Pires and Simon [2015] further provide a Zipf fitting for gamecast popularity based on Normalized Rooted-Mean-Square Deviation (NRMSD). Our work adopts a more sophisticated method for quantifying power-law behaviors as proposed in Clauset et al. [2009]. Most importantly, we distinguish popular and unpopular gamecasts that often have different statistical properties, and we study the long and heavy tail phenomena that have been observed in many general content sharing sites but have not yet been explored for gamecast sharing sites.

In general content sharing sites, Cha et al. [2009] analyze content popularity in two popular UGC sites; their analysis shows power-law-like characteristics. Figueiredo et al. [2011] study content popularity in YouTube and find that copyright protected videos tend to get most views much earlier in their lifetimes. Pinto et al. [2013] present two simple models for predicting future popularity.

The skewed content popularity observed in many content sharing sites indicates that, instead of treating contents indiscriminately, users have certain preferences for sharing and consuming. However, because general content sharing sites often lack explicit and well-structured information on the contents (often providing only titles and descriptions), little related work has provided a detailed analysis of user preference. Our WoTreplays dataset contains detailed game statistics that provide important knowledge for analyzing user preference for sharing and consuming. To the best of our knowledge, this is the first work that sheds light on this topic.

3. METHOD FOR CHARACTERIZING GAMECAST SHARING SITES

In this section, we present an empirical method for characterizing gamecast sharing sites that consists of four main stages: (i) understanding the basic operations of these sites, (ii) identifying interesting and important characteristics and metrics, (iii) selecting and collecting datasets with representativeness and coverage, and (iv) analyzing and presenting the results. We introduce them in turn in the following sections.

3.1. Basic Operations of Gamecast Sharing Sites

A gamecast sharing site keeps a *gamecast repository* into which creators inject the gamecasts they generate and from which spectators watch the gamecasts they are interested in. During the viewing process, users (spectators and creators) may interact with each other via a number of methods like comment and chat. In the following sections, we give a more detailed introduction to the two examples we considered in this article, namely WoTreplays and Twitch.

3.1.1. WoTreplays. WoTreplays is a community-maintained replay-based gamecast sharing site for World of Tanks (WoT). WoT was developed and initially released in 2010 August by Wargaming and has more than 60 million registered players.[2] WoT is

[2]http://worldoftanks.com/en/news/19/world-tanks-hits-mark-60-million-registered-users/.

a typical MMOG in which teams of players, with a maximum of 15 players per team, confront each other in a battle. During a game, each player can gain some *credits*, with the actual amount depending on the player's actions, such as how many tanks have been killed by the player. Players' credits reflect the levels of their gaming skills.

WoTreplays maintains a repository of replays shared by its users, which, upon downloaded, can be viewed with the game engines. The replays capture all player actions, including actions from the keyboard, and they are useful for studying player techniques. In WoTreplays, uploaded replays are displayed by their upload time, and the latest ones are displayed first. To locate and download a replay, users can browse the website or search using various keywords (e.g., tank types). In addition to downloading, users can interact with other users through comments and giving hearts to express their appreciation.

3.1.2. Twitch.tv. Twitch.tv is a leading commercial community for game live streaming. It has more than100 million monthly unique viewers.[3] Twitch.tv contains multiple game genres, such as League of Legends, FanDuel, WoT, and StarCraft II. Twitch players maintain channels, wherein players broadcast gamecasts, chats, and explainations of thei game styles to spectators. Channels are grouped by games and sorted according to their number of concurrent views. Besides browsing channels, users can use some keywords to search channels. A Twitch user can watch the game stream, chat with the player and other spectators, and, if the user enjoys the gamecast, he can further follow the channel, give a heart, or make a donation. Further, Twitch adopts a partnership program that allows streamers to earn revenue by running advertisements and a subscription program that enables a viewer to subscribe to a channel and pay a monthly subscription fee [Kang 2014].

3.1.3. Terminology.

Gamecast (replays and streaming sessions): *Gamecast* refers to the record of a game being played. In WoTreplays, gamecasts are shared after the games are finished, and we call them *replays*. In Twitch, gamecasts are broadcast via live streaming, and we call them game streaming *sessions*.

User classification: For WoTreplays, there are four types of users: *creators*, *gamers*, *uploaders*, and *players*. Their meanings are defined over the entirety of the WoTreplay dataset and are listed as follows.

(1) Creators are users who have uploaded at least one replay.
(2) Gamers are users who have played at least one game.
(3) Uploaders are users who have uploaded at least one replay and have played at least one game.
(4) Players are users who have played at least one game but have not uploaded any replay.

For Twitch, we call players who create channels and stream their gamecasts *streamers*. Depending on the schedule of the streamers, the live streaming of a channel contains a series of *sessions*. The Twitch dataset does not contain detailed game information, so we indiscriminately analyzed the Twitch sessions. For WoTreplays, there are three major types of games: *Winning games*, *Losing games*, and *Survived games*, which are defined as follows.

(1) A winning/losing game is a game wherein its uploader's team has won/lost.
(2) A survived game is a game wherein its uploader's in-game representation (i.e., tank) stayed alive at the end.

[3] http://www.twitch.tv/year/2014.

3.2. Characteristics and Metrics

To characterize gamecast sharing sites and users, we identify the following three important aspects that make up the basic operations of these sites:

Repository characteristics. We consider in this article (i) *gamecast injection* and *duration* that measure the scale of the repository in terms of the number and workload of its contents and (ii) *gamecast popularity* that measures the preference of the spectators in the gamecast level, which is defined as the number of views collected by the gamecasts. We report the injection rate and the duration for gamecasts in the entire repository, analyze the statistical properties of the gamecast popularity, and study its correlation with other features including the age of the gamecasts.

Creator characteristics. We identify four aspects that cover most creator activities, including (i) *creator-level gamecast injection* and *duration* that measure the activity level of the creators; (ii) *creator popularity* that measures the preference of spectators in the creator level, which is defined as the total number of views collected by gamecasts shared by a creator; (iii) *upload delay* that measures the eagerness of the creators for sharing gamecasts, which is defined as the time lag between the finish time of a game and the upload time of its replay; and (iv) *upload preference* that measures the preference of creators for sharing gamecasts.

In this article, we report the injection rate and the duration of gamecasts shared by each creator, analyze the pattern of the creator popularity, and study its correlations with other features including the activity level of the creators. Further, we investigate four features that potentially influence the creators' upload preference, including (i) *game count*, which is the number of games a gamer has played, and similarly, *winning/losing/survived game count*, which is the number of games a gamer has won/lost/survived; (ii) *win ratio*, which is the fraction of games a gamer has won; (iii) *upload count*, which is the number of replays a creator has uploaded; and (iv) *upload ratio*, which is the upload count divided by the game count of an uploader.

Spectator characteristics. We identify two important aspects for the spectator activity specifically (i) *download preference* that measures the preference of spectators for gamecasts and (ii) *interactions with gamecasts* that demonstrates the explicit interactions between spectators and gamecasts, for example, through leaving comments and giving hearts to gamecasts they like.

3.3. Dataset

For each replay, WoTreplays achieves two types of metadata: (i) **gamecast statistics**, including the name of the creator, upload time, replay duration, and the number of downloads/comments/hearts; and (2) **in-game statistics**, including, for the game represented in the replay, the start time, team formation, winning team, and, for each player in the game, the damage and the number of kills he has made, the credits he has gained, and his status at the end of the game (survived or dead).

Twitch maintains basic **gamecast statistics**. For each streaming session in each channel, it contains the information on start time, the name of the player, the number of views per 5 minutes, and the number of hearts. Unlike WoTreplays, Twitch does not maintain metadata on in-game statistics.

To collect the metadata in WoTreplays, we performed a number of separate crawls (fetching webpages through web links), and we obtained the *full history* of WoTreplays since it was first launched in March 2013 until March 2014. In total, we obtained 382,760 replays including 1,956,256 players, uploaded by 63,308 creators. We counted the gamers and creators based on their unique in-game ids; it is possible that a user may have multiple accounts, but we just count those accounts that are from different users. In total, these replays received 5,818,625 downloads, 16,641 comments, and 53,485 hearts from spectators. For Twitch, we use an earlier published dataset [Pires and

Table I. Overview of the Datasets. (Y Means that the Dataset Contains Related Information)

Dataset	# gamecasts	# creators	# players	upload time	gamecast info.	in-game info.
WoTreplays	382,760	63,308	1,956,256	Mar 2013 - Mar 2014	Y	Y
Twitch	7,492,008	1,068,001	N/A	Jan 2014 - Apr 2014	Y	
DotA	991,720	N/A	82,876	Apri 2010 - Sep 2012		Y
StarCraft II	85,532	N/A	83,119	Mar 2012 - Aug 2013		Y

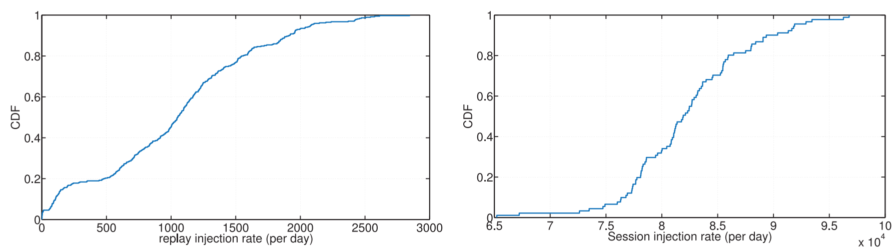

(a) WoTreplays (b) Twitch

Fig. 1. CDF of the gamecast injection rate. Note the difference in the scale of the horizontal axis.

Simon 2015] that contains detailed information on 1,068,001 streamers and 7,492,008 game streaming sessions.

In addition, we include two datasets (DotA and StarCraft II) from our previous work [Iosup et al. 2014]. Although gamecasts in these two datasets are recorded directly by game servers instead of shared individually, they help in comparing game features in different game genres, such as team size and gamecast duration. Particularly for the DotA dataset, we performed a second crawl in September 2012 to include more games. The basic statistics of the four datasets are summarized in Table I.

4. REPOSITORY CHARACTERISTICS

In this section, we provide a basic characterization on the gamecast repository of WoTreplays and Twitch. We mainly focus on three aspects, namely the system-level gamecast injection, duration, and the gamecast popularity.

4.1. Gamecast Injection

Figure 1 shows the Cumulative Distribution Function (CDF) of the gamecast injection rate (i.e., the number of new gamecasts per day). We find that, on average, 1,034 replays and 82,330 sessions are injected each day in WoTreplays and Twitch, respectively. Moreover, we find that the gamecast injection rate in WoTreplays is highly skewed: During two periods that both represent 20% of its history, it achieves fewer than 500 and more than 1,500 daily injections, respectively. To identify this difference, we show in Figure 2(a) the chronological game replay injection rate over the full history of WoTreplays. We see that *as WoTreplays evolves, it attracts more daily injections, indicating that this one-year-old game replay sharing site is gradually expanding*. Though Twitch was first introduced in June 2011, we were only able to obtain its history from January to April 2014, and we observe a relatively stable daily session injection rate, as shown in Figure 2(b).

We also observe from Figure 2, for both WoTreplays and Twitch, a clear *weekend pattern*, with higher daily injection rates on weekends than on work days. Table II shows the statistics, including the mean and standard deviation, of the daily injection rate on different days of a week. The weekend pattern has been observed in other UGC sharing sites, such as YouTube-Live [Pires and Simon 2015], Daum (a Koren UGC site) [Cha et al. 2009], and a DotA gaming site [van de Bovenkamp et al. 2013].

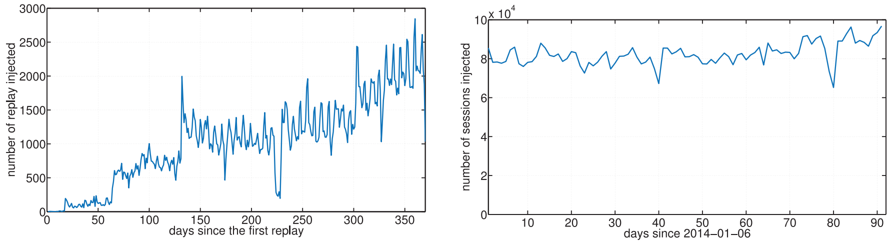

(a) WoTreplays (b) Twitch

Fig. 2. Chronological gamecast injection rate.

Table II. Daily Injection Rate (Mean and Standard Deviation) on Different Days of a Week

mean (std)	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday
WoTreplays	1,026 (606.2)	984 (577.7)	948 (549.7)	929 (557.5)	956 (576.4)	1,160 (735.2)	**1,239** (767.7)
Twitch	80,745 (3,622.9)	78,668 (9,433.6)	82,671 (4,651.0)	80,669 (7,830.3)	83,481 (3,458.6)	**85,366** (5,009.1)	82,235 (6,380.0)

4.2. Gamecast Duration

Gamecast duration directly measures the scale of the repository in terms of the workload of its contents, especially for game live streaming sites like Twitch. As shown in Figure 3, in general, sessions in Twitch have longer durations than replays in WoTreplays: Although 80% of replays in WoTreplays are within 10 minutes, 80% of sessions in Twitch are longer than 20 minutes (see Section 5 for further discussion).

This difference is possibly due to the fact that Twitch users can continue to stream commentary and interviews after the games are finished and that Twitch covers a wide range of game genres, including WoT, DotA, and the StarCraft series, whereas WoTreplays specializes in WoT. Consistently, our DotA and StarCraft II datasets (as introduced in Section 3.3) show that DotA and StarCraft II games are much longer than WoT, achieving an average duration of 36.3 minutes and 18.3 minutes, respectively.

More specifically, the average and the median gamecast duration for WoTreplays is 8.20 and 7.73 minutes, respectively, with a low standard deviation of 2.87 minutes. The average and the median session duration for Twitch is 96.37 and 50 minutes, respectively, with a high standard deviation of 234.09 minutes. Compared with previously published results [Cha et al. 2007, 2009], we find that the median gamecast duration in WoTreplays and Twitch is longer than the median content duration in general UGC sites (e.g., 3 minutes for YouTube), but shorter than the median content duration of non-UGC sites (e.g., 94 minutes for LoveFilm, one of Europe's largest online DVD rental stores [Cha et al. 2009]).

4.3. Gamecast Popularity

In any UGC sharing site, content popularity provides important knowledge for the activity level of users and the potential workload for maintaining the site. As a form of UGC, here we measure the popularity of a gamecast in terms of the number of downloads (views) it collected. Particularly because Twitch provides a gamecast streaming service, we consider both the number of total views a gamecast collected and the peak number of its concurrent views, denoted as *cumulative popularity* and *peak popularity*, respectively. In this section, we conduct a set of analyses that provide a holistic view of gamecast popularity. We first study the skewness of user requests across gamecasts. Then, we analyze how user requests are distributed across popular and unpopular gamecasts by examining the heavy and long tail phenomena [Mahanti et al. 2013].

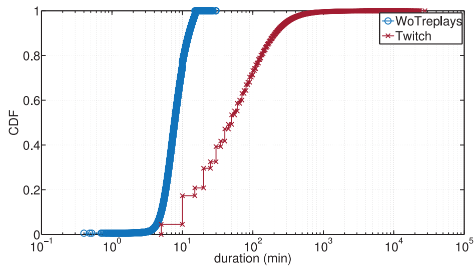

Fig. 3. CDF of the gamecast duration.

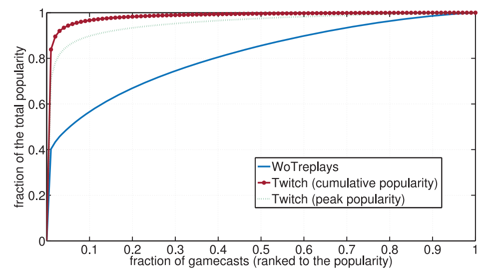

Fig. 4. Skewness of gamecast popularity.

4.3.1. The Skewness of Gamecast Popularity. To study skewness of popularity in gamecast sharing sites, we calculate the fraction of the total popularity aggregated by the rth most popular gamecasts in WoTreplays and Twitch, respectively. Results are shown in Figure 4. The horizontal axis represents the fraction of gamecasts ranked from the most popular to the least popular. For Twitch, 10% of the gamecasts account for more than 90% of the total popularity in the repository in terms of both cumulative and peak popularity. Compared to Twitch, gamecast popularity in WoTreplays is less skewed, with 10% of gamecasts representing less than 60% of the total popularity.

In general UGC sites like YouTube, the skewness of popularity has been observed as well, with roughly 10% of content representing 80% of the total popularity [Cha et al. 2007]. It is interesting to notice that the skewness of gamecast popularity in Twitch and WoTreplays, lies above and under the one for general UGC sites, respectively.

We conjecture that the difference in popularity skewness is a consequence of the recommendation algorithms used in these sites. By default, Twitch sorts its gamecasts by the number of concurrent views in its browse list, whereas WoTreplays merely displays gamecasts by date. It is likely that *the recommendation algorithm used in Twitch helps promote the dominance of popular gamecasts and therefore induces a highly skewed gamecast popularity.* The recommendation algorithm used in YouTube is more complicated, with considerations of both content popularity and many other aspects [Davidson et al. 2010], resulting in a content popularity that is less skewed than Twitch, wherein the default promotion is solely for popular gamecasts, and that is more skewed than WoTreplays, wherein the default promotion is irrelevant to the current gamecast popularity.

4.3.2. Statistical Properties. In this section, we delve further into the statistical properties of the gamecast popularity and examine whether power-law characteristics apply. A distribution is considered to follow a power-law relationship if its probability density function takes the form $f(x) = Cx^{-\alpha}$. The constant α is called the scale, and once α is fixed, the constant C is determined by the requirement that the distribution $f(x)$ sum to 1. Taking logarithms on both sides produces $log(f(x)) = -\alpha log(x) + log(C)$. This expression exhibits a linear relationship with a slope of $-\alpha$ when plotted on a log-log scale.

Two frequently occurring (and confusing) terms associated with power-law distributions are *heavy tails* and *long tails*. A distribution is considered to have a heavy tail if its tail is not exponentially bounded. A power-law distributed tail is one example for a heavy tail. In the context of content sharing, a heavy tail represents a small number of popular contents accounting for a large fraction of the total popularity. The long tail, on the other hand, is a manifestation of power-law relationships. The term became popular when researchers showed that online purchasing sites like Amazon benefit from the long tail: A large number of items, each attracting only a few customers, but altogether accounting for a significant part of the total sale [Anderson 2004].

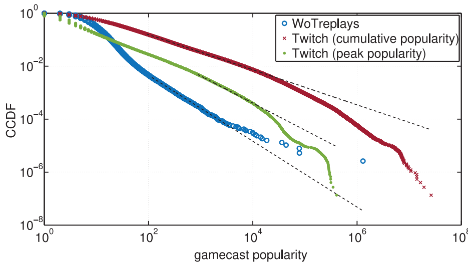

Fig. 5. CCDF of the gamecast popularity.

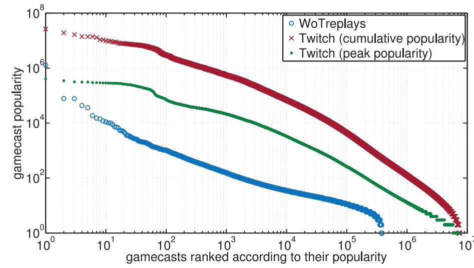

Fig. 6. Gamecast popularity, in log-log scale.

Table III. Power-Law Fitting Results for the Gamecast Popularity (Number of Downloads/Views)

	x_{min}	α	p-value	D	n_{tail}	n_{tail}/n	p_{tail}/p_{total}
WoTreplays	116	2.2437	0.513	0.0143	1424	0.37%	36.98%
Twitch (cumulative popularity)	92	1.6228	0.483	0.0069	1,218,805	16.29%	97.85%
Twitch (peak popularity)	880	2.0229	0.496	0.0187	36,134	0.48%	62.11%

In the context of gamecast sharing sites, the heavy and the long tail represent the popular and unpopular gamecasts respectively, which we will use for our analysis in this section. To study the heavy tail phenomenon, we consider a Complementary Cumulative Distribution Function (CCDF) graph that shows the fraction of gamecasts with popularity that is higher than a variety of values, as shown in the horizontal axis. To study the long tail phenomenon, we use a plot of gamecasts ranked in the decreasing order of their popularity. The tails of the CCDF graph and the plot represent the popular and unpopular gamecasts, respectively.

Heavy tail: The popular gamecasts. Figure 5 shows the CCDFs of gamecast popularity in WoTreplays and Twitch, respectively. The dashed lines represent the fitted power-law distributions. When plotted on a log-log scale, all exhibit a straight line, especially on the tails, indicating a power-law distributed characteristic. To be more rigorous, we further perform power-law fittings to test whether the heavy tail phenomenon occurs.

We use the tool proposed in Clauset et al. [2009] for discerning and quantifying power-law behaviors. The tool combines maximum-likelihood fitting methods with Goodness-of-Fit (GoF) tests. In practice, few empirical phenomena obey power laws for all values of the data. More often, the power law applies only for values greater than some minimum value denoted by x_{min}. Therefore, instead of fitting the whole data, this approach focuses only on $x \geq x_{min}$ (i.e., the tail)—this is exactly what we need for testing whether the heavy tail phenomenon occurs.

In addition to the start of the fitting, x_{min}, this method also returns the scale α and the p-value for the fitted distribution. The scale measures the heterogeneity of the data. In our analysis, a larger value of α indicates more skewed popularity. For the p-value, we use 0.05 as the significance level, below which the null hypothesis that the fitted distribution represents the empirical data is rejected. In addition, we also calculate the largest gap between the empirical Cumulative Distribution Function (CDF) and the fitted CDF, denoted by D, the number and the fraction of data contained in the tail (i.e., $x \geq x_{min}$), denoted by n_{tail} and n_{tail}/n, and the fraction of total popularity achieved by the tail, denoted by p_{tail}/p_{total}. The fitting results are shown in Table III.

We find that, for both WoTreplays and Twitch (and the two types of popularity we considered in Twitch), power-law distributed popularity holds for the very popular gamecasts (with p-values significantly larger than 0.05 for $x \geq x_{min}$), indicating that *the heavy tail phenomenon occurs in these two gamecast sharing sites*. Nevertheless, the size of the heavy tails and the fraction of total popularity achieved by them are

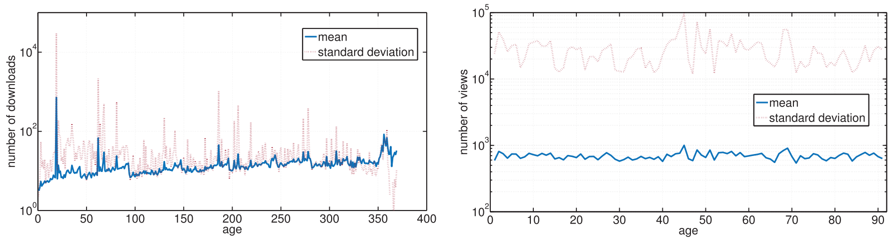

(a) WoTreplays (b) Twitch

Fig. 7. Gamecast popularity versus gamecast age; the vertical axis is in a logarithmic scale.

different across WoTreplays and Twitch. The tail for the cumulative popularity in Twitch contains the most gamecasts, which is more than 30 times more than the tail for the peak popularity in Twitch, and is more than 850 times more than the tail for the popularity in WoTreplays. Furthermore, we observe that the tails in these sites all account for a significant part of the total popularity and, again, the tail for the cumulative popularity achieves a higher fraction of its total popularity than the tail for the peak popularity in Twitch, which in turn is higher than the tail for popularity in WoTreplays. This is consistent with our previous result in Figure 4 that shows the same pattern of difference in the level of skewness observed in these two sites.

Long tail: The unpopular gamecasts. The long tail phenomenon is often analyzed based on the *rank/frequency* plot [Newman 2006], which we show in Figure 6. Here, the vertical axis shows the popularity for each gamecast, and the horizontal axis shows the ranking of gamecasts based on the decreasing order of their popularity. For each dot on this plot, its pair of values, say (x, y), represents that there are x gamecasts with popularity higher than y. The *rank/frequency* plot in fact is another form of the CCDF plot.

For both WoTreplays and Twitch, we observe roughly straight lines for the tails in this log-log plot, representing a power-law characteristic for the unpopular gamecasts. This result indicates that, in both sites, there exist long tails that contain a large number of gamecasts with low popularity. Together with the results shown in Figures 4 and 5, we find that 90% of the replays in WoTreplays and 60% of the sessions in Twitch have been watched in total less than 20 times, accounting for 40% and 2% of the total popularity, respectively. We also find that 84% of the sessions in Twitch achieve fewer than 10 peak concurrent views, accounting for less than 10% of the total peak popularity. These results may indicate that, *in both WoTreplays and Twitch, the long tail phenomenon occurs, and the long tail in WoTreplays plays a more important part for total popularity than that in Twitch.*

4.3.3. Popularity and Age. In this section, we study the cumulative effect of time by analyzing the popularity of gamecasts. We count the age of a gamecast as the total number of days since the gamecast was published online. Figure 7 shows the popularity (number of downloads/views) achieved by gamecasts in different age groups. For both WoTreplays and Twitch, we have two interesting observations. First, we did not find a strong correlation between gamecast popularity and age, which only achieves a Spearman Ranking Correlation Coefficient (SRCC)[4] of 0.0007 and 0.0004, respectively. Second, we observe very high deviations for gamecasts in the same age group: For WoTreplays, it is on the same level as the mean value, and for Twitch, it is two orders of magnitude higher than the mean value.

[4] In brief, SRCC assesses how well the relationship between two variables can be described using a monotonic function [Spearman 1904].

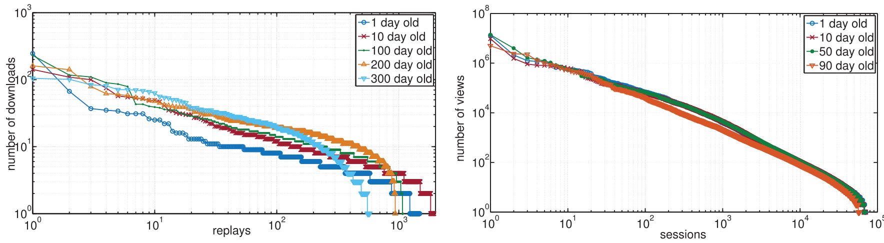

(a) WoTreplays (b) Twitch

Fig. 8. Gamecast popularity in different age groups, in log-log scale.

These results indicate that, statistically, in both WoTreplays and Twitch, *gamecasts do not necessarily accumulate more downloads and views over time*. To further demonstrate this effect, we show in Figure 8 the popularity of each gamecast in five different age groups, with the gamecasts ranked in decreasing order of their popularity. For WoTreplays, we only observe a slight increase from 1-day-old to 10-day-old gamecasts of the same rank, whereas for Twitch, gamecasts of the same rank but in different age groups achieve very similar popularity. In general UGC sites like YouTube, though contents often accumulate a large amount of views at their early ages, the cumulative effect of time is still observable—roughly, the content popularity for videos 3 months old is two orders of magnitude higher than that of videos 1 day old in YouTube [Cha et al. 2009].

Together with the comparatively large gamecast injection rate, we believe that *WoTreplays and Twitch are repositories of items that have short-term values: A considerable amount of gamecasts are injected, viewed, and soon forgotten*. We conjecture that this is due to the nature of gaming: As game-playing techniques evolve quickly over time, stale gamecasts no longer provide enough information for education or entertainment.

5. CREATOR CHARACTERISTICS

We have shown that WoTreplays and Twitch are constantly expanding repositories, with a large number of new gamecasts being injected on a daily basis. In this section, we investigate the basic characteristics of the creators. We first show their activity level and upload delay. Then, we conduct an analysis of creator popularity, with a focus on its statistical properties. Finally, with the detailed in-game statistics archived by WoTreplays, we analyze the upload preference of creators.

5.1. Activity Level of the Creators

We use the number and duration of gamecasts shared by creators to measure their activity level. As shown in Figure 9, for creators in WoTreplays and Twitch, their activity levels are highly skewed: While 70% of creators have shared less than 5 gamecasts, 3% of creators have shared more than 40 gamecasts. We observe similar skewness for the total duration of gamecasts shared by creators as well.

Meanwhile, although the number of gamecasts shared by creators in these two communities achieve similar statistical patterns, creators in Twitch in general accumulate one order of magnitude longer gamecast duration than do creators in WoTreplays. Together with the fact that our WoTreplays and Twitch datasets contain histories of 1 year and 3 months, respectively, we conclude that creators in Twitch are more active than those in WoTreplays. The detailed statistics are shown in Table IV.

Interestingly, we identify a considerable number of creators in Twitch who have streamed almost continuously. We find that 500 creators have streamed on average

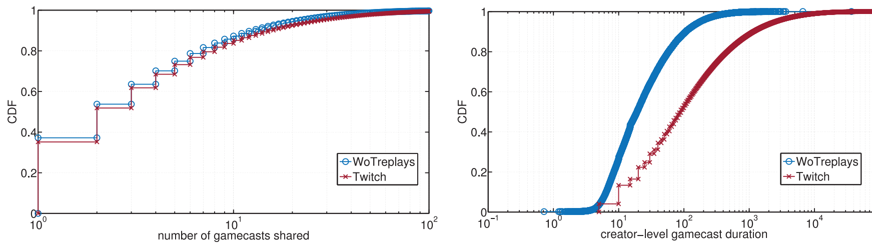

(a) The number of the gamecasts (b) The duration of the gamecasts

Fig. 9. CDF of the number and the duration of the gamecasts shared by creators.

Table IV. Statistics on Creator Activity Level, Including Mean, Median, Maximum, Minimum, and Standard Deviation (STD)

	Number of Gamecasts					Duration (Minute)				
	mean	median	max	min	std	mean	median	max	min	std
WoTreplays	6.04	2	4683	1	23.64	49.67	19.11	36,978	0.72	185.02
Twitch	7.01	2	5388	1	20.59	676.18	90	130,410	5	3320.6

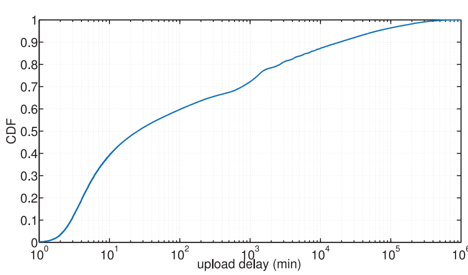

Fig. 10. CDF of the upload delay in WoTreplays.

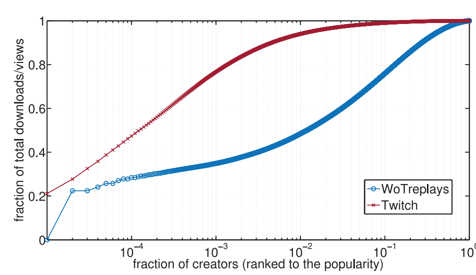

Fig. 11. Skewness of creator popularity.

more than 12 hours a day, and 100 creators have streamed in total more than 80 days during the whole data collection period (90 days). As Twitch's partnership program allows streamers to earn revenue through streaming, it is possible that some of these creators are teams of professionals in online e-sports streaming rather than single players.

In regular content sharing sites, the upload time often reflects the timeliness of the content. For example, contents related to breaking news are normally immediately uploaded when they are available. Here, we use the *upload delay* to measure how quickly creators upload their replays in WoTreplays, which is defined as the time difference between the finish time of a game and the upload time of its replay. Figure 10 shows the CDF of the upload delay for each replay in WoTreplays. We see that 50% of the replays are uploaded within 20 minute after the games are finished, indicating that creators are often very eager to share their games. In Twitch, the gamecasts are broadcast via live streaming (i.e., no upload delay). In the meantime, the broadcasters chat and explain their game styles to their spectators. This new form of communication helps broadcasters to build relationships with their spectators, which partially sparks interest for users in online e-sports.

5.2. Creator Popularity

In Section 4.3, we conducted a detailed analysis of gamecast popularity. In the context of UGC sites, content popularity has received great attention for scientific research

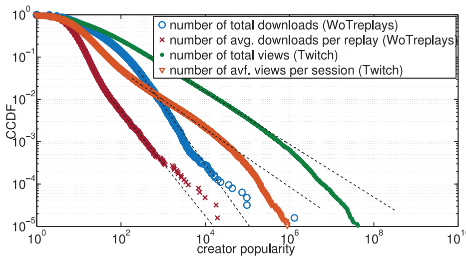

Fig. 12. CCDF of the creator popularity.

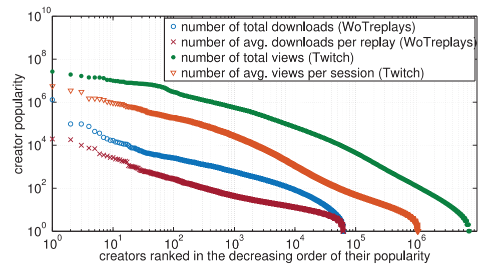

Fig. 13. The creator popularity, in log-log scale.

Table V. Results of the Power-Law Fitting for the Creator Popularity

	x_{min}	α	p-value	D	n_{tail}	n_{tail}/n	p_{tail}/p_{total}
WoTreplays (total downloads)	261	2.3961	0.494	0.0179	2954	4.67%	65.25%
WoTreplays (avg. downloads per replay)	56	2.2374	0.497	0.0224	670	1.06%	27.41%
Twitch (total views)	477	1.6230	0.502	0.0044	107,010	10.02%	99.01%
Twitch (avg. views per session)	181	1.6847	0.482	0.0155	32,894	3.08%	90.96%

because of the importance of understanding users' preferences. However, to the best of our knowledge, the popularity patterns for their creators are still unexplored. Nevertheless, it is equally (if not more) important to understand the popularity of content sharers—those who directly generate the contents and maintain high content popularity—by following the activity level of popular content sharers. In this section, we analyze creator popularity in gamecast sharing sites.

We use two metrics to measure the popularity of creators, namely the total number of views and the average number of views for the gamecasts they share. Following the methods used in Section 4.3, our analysis focuses on examining the skewness of popularity and studying the statistical properties of popular creators (i.e., the heavy tail) and unpopular creators (i.e., the long tail).

5.2.1. Skewness. In Figure 11, we show the fraction of the total popularity aggregated by the rth most popular creators in WoTreplays and Twitch, respectively. In general, we observe a very skewed creator popularity in both communities: For WoTreplays, 10% of creators account for 75% of the total downloads, and for Twitch, 1% of the creators have accumulated more than 90% of the total views. Together with the results of the skewness analysis for gamecast popularity (as shown in Figure 4), we find that, for the same community, creator popularity is more skewed than gamecast popularity. This result indicates that, *to attract more users, maintaining popular creators is potentially more effective than maintaining popular gamecasts*.

5.2.2. Heavy Tail: The Popular Creators. In Figure 12, we show the CCDF of creator popularity in WoTreplays and Twitch, respectively. The dashed lines represent the fitted power-law distributions. We used the same method for the power-law fitting as described in Section 4.3.2, and the fitting results are shown in Table V.

We find that, for both WoTreplays and Twitch, power-law distributed popularity holds for very popular creators (with p-values significantly larger than 0.05 for $x \geq x_{min}$), indicating that *the heavy tail phenomenon applies to creator popularity in both of these two sites*. Nevertheless, the size and the fraction of total popularity achieved by their heavy tails are different. Possible reasons are that, as shown in Figure 11, creator popularity in Twitch is more skewed than that in WoTreplays, and therefore, its heavy

Table VI. Correlations Between Creator Popularity, p_c, Number of Gamecasts, n_g, and Average Popularity of Gamecasts Shared by Creators, Avg. p_g

SRCC	p_c vs. n_g	p_c vs. avg. p_g	n_g vs. avg. p_g
WoTreplays	0.8446	0.7361	0.2906
Twitch	0.8323	0.8766	0.4906

tail contains a larger fraction of creators and accounts for a higher fraction of the total popularity—in other words, *the tail of the creator popularity is heavier in Twitch than in WoTreplays*.

5.2.3. Long Tail: The Unpopular Creators. To examine the long tail phenomenon, we show the *rank/frequency* plot of creator popularity in Figure 13. Here, the vertical axis shows the popularity for each creator and the horizontal axis shows the ranking of creators based on the decreasing order of their popularity. As discussed in Section 4.3, the *rank/frequency* plot is another form of the CCDF plot.

For both WoTreplays and Twitch, we observe almost straight lines for tails in this log-log plot, representing a power-law characteristic for the unpopular creators. This result indicates that *both WoTreplays and Twitch contain a large number of creators with low popularity*. Together with the results shown in Figures 11 and 12, we find that 90% of creators in WoTreplays have attracted in total less than 100 downloads for the replays they share, roughly accounting for 25% of the total downloads; in Twitch, more than 70% of creators have collected less than 100 views for their streaming sessions; however, these unpopular creators only account for 0.025% of the total views.

5.2.4. Discussion. Combining the preceding results on creator popularity with the results on gamecast popularity (as shown in Section 4.3), we conclude that *both the gamecast and the creator popularity in WoTreplays and Twitch are skewed*, with creator popularity more skewed than gamecast popularity. We also confirm that *the heavy tail phenomenon occurs in these two sites*, indicating that a small number of popular gamecasts and creators account for a large amount of the total popularity. Nevertheless, we observe *different long tail phenomena in these two sites*. While they both contain a large number of gamecasts and creators with low popularity, in WoTreplays, these gamecasts and creators accumulate a considerable amount of the total popularity; but in Twitch, their share in the total popularity is neglectable.

One immediate application of these findings is to maintain user activity level in these sites. Understanding users' preferences for gamecasts and creators, the administrators of these and similar sites can identify popular gamecasts and creators, customize incentive policies for them to share more, and hence attract more users to the sites.[5] Furthermore, for sites with significant long tail phenomenon (e.g., WoTreplays) wherein unpopular contents accumulate a considerable amount of the total popularity, better recommendation algorithms can be applied to promote unpopular content to improve their potential for attracting more users.

5.3. Building up Creator Popularity

We have shown that popularity distributions of creators in WoTreplays and Twitch are skewed. In this section, we explore possible reasons for this skewness. Table VI shows the SRCCs between creator popularity, the number of gamecasts, and the average popularity of gamecasts shared by creators. We observe clear correlations between either of the latter two metrics and creator popularity, but not between these two metrics. This

[5]Twitch's partnership program encourages popular creators to share more high-quality gamecasts. However, the details of the program are not specified to an extent that enables service-level analysis.

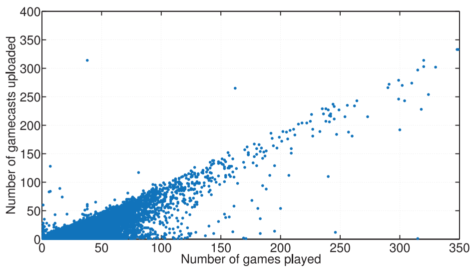

Fig. 14. Upload count versus game count.

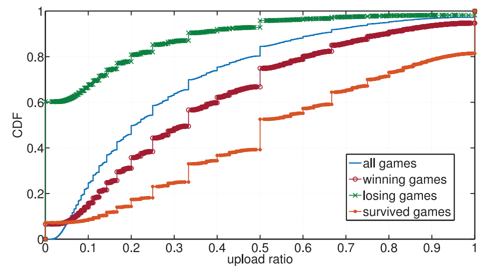

Fig. 15. CDF of the upload ratio.

result suggests that *a popular creator tends to have many popular gamecasts*. Nevertheless, we can still identify a small number of creators who have achieved extremely high popularity with only a few gamecasts. For example, in WoTreplays, 30 creators in the top 100 most popular creators have shared in total less than 20 gamecasts. And in Twitch, 37 creators in the top-100 list have shared no more than 50 gamecasts. We conjecture that their popularity is due to their gaming reputation, whereas, in general, creator popularity is built up gradually through constant sharing.

A creator's gaming reputation may consist of various aspects, including skill and the style of his gameplay. A more detailed analysis of the relationship between the reputation and the popularity of creators requires quantifying their gaming reputation, which potentially needs more data on the in-game performance and possibly a user survey. We leave this to future work.

5.4. Upload Preference

In this section, we analyze the upload preference of creators (i.e., whether and how creators are selective when choosing gamecasts to share). This analysis requires detailed game statistics, including team formation, the winning and losing teams, in-game skill of each player, and the reward the player gets at the end of the game. Since only the WoTreplays dataset contains such information, our analysis in this section focuses on uploaders in WoTreplays (i.e, creators who have played at least one game). The behavioral metrics we considered are as introduced earlier in Section 3.2.

5.4.1. Upload Everything?. We first show in Figure 14 the scatter plot of the upload count and game count for each uploader. We observe a modest correlation between these two metrics, achieving a SRCC of 0.5414. This result indicates that *uploaders do not necessarily upload more games when they have played more.*

To further illustrate this phenomenon, Figure 15 shows the CDF of the upload ratio (defined as the upload count divided by the game count of an uploader). We find that it takes a wide range of values, and only 20% of uploaders have uploaded more than half of all the games they have played. This result indicates that *the uploaders do not upload indiscriminately*. In the following sections, we investigate this selective upload behavior.

5.4.2. Upload the Winning and the Survived Games?. We conjecture that two possible reasons for an uploader to share a replay are the victory of his team and/or of himself. Following this intuition, we test whether an uploader is prone to upload those games he has won and those games he has survived (see Section 3.1.3 for definitions).

Figure 15 shows the CDF of the upload ratio for winning, losing, and survived games, respectively. We find that, in general, the upload ratio for winning games is much higher than that for losing games: While only 3% of uploaders have uploaded at least half of

Table VII. Statistics for the Upload Ratio and the Fraction of Different Types of Games in Played and Uploaded Games of Each Uploader, with Mean and Standard Deviation (STD)

	Upload Ratio mean (std)	Fraction in Played Games mean (std)	Fraction in Uploaded Games mean (std)
winning games	0.3717 (0.2736)	0.6088 (0.1758)	**0.8379** (0.2817)
losing games	0.1120 (0.1989)	**0.3912** (0.1758)	0.1475 (0.2686)
survived games	0.5526 (0.3185)	0.3738 (0.2017)	**0.7802** (0.3139)
losing and survived games	**0.4982** (0.4760)	0.0177 (0.0486)	**0.0338** (0.1347)
losing but not survived games	**0.0852** (0.1672)	**0.3736** (0.1744)	0.1173 (0.2402)

their losing games, 30% of uploaders have done so for their winning games. As it turns out, uploaders are even more prone to share survived games: 60% of uploaders have shared at least half of their survived games, and 20% of uploaders have shared all of them. We have also calculated the SRCCs between the upload count and the winning, survived, and losing game count, which are 0.6217, 0.7201, and 0.3067, respectively. Clearly, *uploaders are more prone to share winning and survived games than losing games.*

With this upload preference, it is natural that uploaders achieve higher win ratios in the archived games than players who did not upload any replays. We find that more than 20% of players have not won any games, whereas more than 50% of uploaders have won at least half of their games. We also observe that around 20% of players and around 5% of uploaders have achieved a win ratio of 100%. We conjecture that the complete victory is due to the small game count of these players (on average, fewer than 2 and 4 games for players and uploaders, respectively).

5.4.3. Upload the Losing but Survived Games?. Alhough uploaders prefer to upload winning and survived games, from Figure 15, we still observe a considerable amount of losing games being uploaded, and we can identify 40% of uploaders who have uploaded at least one of their losing games. In this section, we test what kind of losing games that uploaders are inclined to share. Following the results shown in the last section, we focus on those games an uploader has lost but survived in the end.

Table VII shows the statistics for the upload ratio of different types of games. Clearly, the upload ratio for losing and survived games is much higher than that for losing but not survived games. We conjecture that when an uploader decides to share a losing game, the uploader mainly does so to show his individual skills. To better illustrate this phenomenon, we further show in Table VII the statistics on the fraction of different types of games in the played and uploaded games of each uploader. We see that the winning games, survived games, and losing but survived games all take higher fractions in the uploaded games than in the played games, suggesting the uploader's preference for them. We conjecture that, *for uploading game replays, the uploaders in general focus more on their individual performance than on the outcome of their teams.*

5.4.4. Show off Skills?. In this section, we study the in-game skills of uploaders. The in-game skill of a WoT gamer can be measured via the number of tanks he has killed, the damage he has done, and the credit he has gained. In our analysis, we focus on the number of kills, since we can use the maximum team size (i.e., 15 players) as the baseline to infer a gamer's skill.

We find that, while in WoT each team has maximum 15 players, more than 40% of uploaders have killed at least 5 tanks (i.e., one third of the opponents). More specifically, the average number of kills takes a mean value of 4.64, a median value of 4.5, and a maximum value of 15 respectively, with a standard deviation of 2.19. It is very likely that uploaders are showing off their skills through sharing the replays.

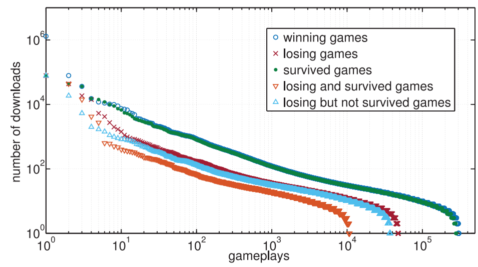

Fig. 16. Number of downloads for different types of games.

Table VIII. Statistics for the Size and Number of Downloads of Replays of Different Game Types

	Fraction of Games	Fraction of Total Downloads	mean	Number of Downloads median	max	std
winning games	79.65%	88.13%	15.74	7	1,299,174	2359.00
losing games	12.57%	11.87%	13.43	7	77,869	422.32
survived games	73.76%	61.09%	11.78	7	78,378	193.75
losing and survived games	2.8%	3.24%	16.38	7	43,546	444.39
losing but not survived games	9.8%	8.63%	12.58	7	77,869	415.75

6. SPECTATOR CHARACTERISTICS

Spectator activity captures the most basic characteristics of gamecast sharing sites. In the previous sections, we indirectly investigated their download activity through an extensive analysis of gamecast and creator popularity, which, as it turns out, are both highly skewed. In this section, we further explore the download preferences of spectators, and we conduct an analysis of the interactions between spectators and the WoTreplays repository in terms of the comments and the hearts they give.

6.1. Download Preference

To understand the download preferences of spectators, we first study the spectator's potential preference for replays of different game types (i.e., winning, losing, and survived games). Then, we analyze the possible influence of uploader gaming skills.

6.1.1. Preference for Game Types. We first show in Figure 16 the number of downloads for replays of different game types, where the replays of each game type are ranked in the decreasing order of the number of downloads they collected. In general, spectators prefer to download winning and survived games over losing games, and, within the losing games, they prefer survived over not survived losing games. The detailed statistics are shown in Table VIII. We have two interesting observations, as follows.

First, compared to the losing and the survived games, the winning games have the largest fraction (79.65%), account for the largest fraction of the total downloads (88.13%), and achieve the highest average number of downloads (15.74, the mean value). Second, losing games achieve a higher mean value for the number of downloads than survived games. The 5-times difference in the number of losing games and the number of survived games suggests that the survived games contain a longer tail than the losing games, and this long tail (meaning a large number of games with a small number of downloads) induces a relatively small mean value. If we compare the top-n survived and the top-n losing games ($n \leq 48{,}113$, which is the number of the losing games), as shown in Figure 16, the survived games always achieve a larger number of downloads. The same argument applies to the losing and survived games that achieve a large mean value of the number of downloads as well.

These results indicate that, when making download decisions, *the spectators appreciate team performance more than the individual performance of the uploaders*. Note

Table IX. Correlation Between the Number of Downloads (DLS) and Uploader Skills

SRCC	Kills	Stars	Damage	Credit	Win Ratio
dls (per replay)	0.1361	0.2076	0.2315	0.1572	N/A
dls (per uploader)	0.1702	0.2702	0.3232	0.2234	0.1032

Table X. Statistics on the Comments at the Replay Level

#Comments	Mean	Median	Std	Max	Frac. of Comments	Mean	Median	Std	Max
all	1.3195	1	1.2286	43					
from the uploaders	0.5951	1	0.5882	11	from the uploaders	**0.4994**	0.5	0.4733	1
from the gamers	0.1034	0	0.3307	5	from the gamers	0.0863	0	0.2723	1
from the spectators	0.6209	0	1.2184	41	from the spectators	**0.4143**	0	0.4676	1

that earlier, in Section 5.4, we showed an opposite tendency for the upload preference of uploaders: *they focus more on the individual than on the team performance*.

6.1.2. Preference for Uploaders' Skills. Uploaders' gaming skills are reflected by a number of metrics. Our WoTreplays dataset includes (i) for each game, the *number of kills*, the *experience* (number of stars), the *damage*, and the *credit* that the uploader has earned at the end of the game; and (ii) in the long run, his *win ratio* (the fraction of winning games in all the games he has played). We calculate the SRCCs between the number of downloads and these five metrics, at both the replay and the uploader levels. Taking the number of kills as the example, at the replay level, we calculate the SRCC between the number of downloads of each replay and the number of kills made by its uploader; at the uploader level, we calculate the SRCC between the average number of downloads for replays uploaded by each uploader and the average number of kills he has made in those replays. The results are shown in Table IX.

In brief, we did not observe any strong correlations between the number of downloads and the number of kills the uploaders have made, nor for their experience, damage, credit, or win ratios. Nevertheless, we do find that the SRCCs are in general larger when the metrics are considered at the uploader level than at the replay level.

6.2. Interactions with Replays: Comments and Hearts

As in a general UGC sharing site, users in WoTreplays can comment and give hearts to the replays they are interested in. In total, 12,377 (42,964) replays have received at least one comment (one heart), with a maximum number of comments and hearts of 43 and 80, respectively. Among these replays, 80% of them have received only one comment (one heart). Meanwhile, we find that the number of comments correlates with the number of hearts received by each replay, achieving an SRCC of 0.5407.

6.2.1. Comments at the Replay Level. While WoTreplays only archives the number of hearts a replay has received, it archives more detailed information on the comments, including the time, contents, and the names of those users who have left the comments. Table X shows the basic statistics on the comments at the replay level. We see that, for each replay, the comments it received are in general equally split among uploaders and spectators, achieving a mean value for the fraction of comments of 0.4994 and 0.4143, respectively. This result indicates that *the comment function maintained by WoTreplays provides the opportunity for uploaders and spectators to interact with each other explicitly*.

6.2.2. Comments at the User Level. We further conduct an analysis of the statistical properties of comments at the user level. We find that, in total, 8,875 users (<0.5% of total users) have left at least one comment, resulting in a total of 16,331 comments for the replays achieved by WoTreplays. Among these users, 5,319 users, 983 users, and

Table XI. Statistics on Comments at the User Level

#Comments	Mean	Std	Median	Max
all	1.8401	3.7891	1	205
as uploaders	**0.8300**	1.3782	1	74
as gamers	0.1442	0.4796	0	8
as spectators	**0.8659**	3.4956	0	**205**
#uploaded replays	12.8397	28.7852	5	1000
#played games	14.6173	12.1819	12	180

Table XII. Correlation Between Number of Downloads (DLS) and Number of Comments (c) from Uploaders (c_u), Gamers (c_g), and Spectators (c_a)

SRCC	c	c_u	c_g	c_a
dls (per replay)	**0.3060**	0.0798	0.0515	**0.3716**
dls (per uploader)	0.1712	**0.3783**	0.0720	0.1845

3,755 users have left at least one comment as an uploader, a gamer, or a spectator, respectively. The detailed statistics are shown in Table XI.

In general, users are more likely to leave comments as uploaders and as spectators, achieving an average number of 0.8300 and 0.8659 comments, respectively. This is consistent with our previous results, which show that, at the replay level, comments are mainly left by uploaders and spectators. Particularly, we find the user with the most comments left 205 comments in 181 replays, as a spectator. By manually checking the contents of these comments, we find that he is mostly asking for permission to share the replays on YouTube. Although these users have on average uploaded 12.8397 replays and played 14.6173 games, they have commented in less than 10% of them.

These results indicate that *many users in WoTreplays are silent spectators*: They mostly download replays without engaging in any explicit interactions. Nevertheless, through delicate modeling, rich (social) relationships can be detected in such communities and used to improve gaming services, such as match recommendation and game population retention, which are important for the user experience and the commercial value of those companies that run these communities. We refer interested readers to our previous work [Jia et al. 2015] for details.

6.2.3. Comments and Popularity. The number of comments received by a replay reflects its prosperity, and the number of comments left by an uploader reflects the uploader's activity level. Intuitively, replays and uploaders with many comments are potentially popular among spectators, and therefore they may collect a considerable amount of downloads. To study the influence of the number of comments on the replay and uploader popularity, we calculate the SRCCs between these two metrics at both replay and uploader levels. More specifically, for each replay, we consider its total number of comments and its number of comments from the uploader, gamer, and spectators. For each uploader, we consider his total number of comments and his number of comments as an uploader, as a gamer, and as an spectator. The results are shown in Table XII.

We do not observe any strong correlations between the number of comments and the replay and uploader popularity. Nevertheless, we do find that, at the replay level, the SRCC between replay popularity and the number of comments *from spectators* is one order of magnitude higher than SRCCs between replay popularity and the number of comments from uploaders or from gamers. And, at the uploader level, the SRCC between uploader popularity and their number of comments *as uploaders* is 2–5 times higher than SRCCs between uploader popularity and their number of comments as gamers and as spectators. These results suggest that, for the replays, comments from their spectators, and, for the uploaders, their comments as uploaders, are more important for boosting their popularity.

7. DISCUSSION

In previous sections, we investigated the workload and user behaviors in WoTreplays and in Twitch. In this section, we discuss possible future work that is worth exploring.

Gamecast platform : With the workload fluctuation we observed in WoTreplays and in Twitch, smart resource scheduling policies can be developed to reduce the operational cost of these two sites while still satisfying the Quality of Service (QoS). For example, Aparicio-Pardo et al. [2015] propose a method based on Adaptive Bitrate Streaming. Nevertheless, their method only deals with fixed instead of dynamic number of computing resources. Nae et al. [2011], Lu et al. [2013], and Li et al. [2015] further propose using resources in an on-demand manner. We believe that a combination of their methods could be a good direction to solve this problem.

Another way to reduce the operational cost is by taking into account the skewness of gamecast popularity. For example, to balance the workload, gamecasts could be allocated to servers based on their popularity instead of their pure number. It should be noted that, as the gamecast popularity in Twitch is more skewed than that in WoTreplays, a platform optimized for Twitch may be suboptimal for WoTreplays: A method dedicated to Twitch may place a number of the most popular gamecasts on one server, but for WoTreplays, placing the same number of top popular gamecasts on one server will make it underloaded. Furthermore, because gamecasts often attract fewer downloads over time, especially in Twitch-like platforms with streaming servers, gamecast platforms should take the evolution of gamecast popularity into account and allocate the latest gamecasts to fast cache/storage and the old gamecasts to cheaper but slower one.

Gamecast recommendation : Although recommendation mechanisms on general UGC sites such as YouTube have been widely analyzed [Davidson et al. 2010], gamecast recommendation is relatively less explored. Kim and Kim [2014] propose a method based on in-game features and user rating. Nevertheless, their method requires users to rate replays with nontrivial efforts. Furthermore, in addition to in-game statistics, we believe that the popularity and the co-watched counts of gamecasts[6] can also be leveraged to design better gamecast recommendation algorithms.

On the other hand, because we observed that a number of WoTreplays users have uploaded replays of games wherein they did not play, it would be interesting to understand the reason behind this behavior: Was it because these replays are well-played? An in-depth analysis of these games may provide insight into the design of gamecast recommendation algorithms.

Pre- and Post-game periods : In Twitch-like platforms, a gamecast records not only the game but possibly also what happened before and after the game, such as creators chatting with spectators. The pre- and post-game periods provide good opportunities for creators to socialize with their spectators and to increase their popularity. Because it is difficult to extract the pre- and post-game periods from a gamecast, the analysis of these periods is still missing. In this work, we provide a preliminary analysis of this topic by examining upload delay (i.e., the time between the end of a game and the upload time of the gamecast), which provides a hint for the post-game periods. However, a further analysis is needed to better understand the pre- and post-game periods.

Methedology: Our study, albeit conducted in good faith and with due care for scientific norms, suffers from the same limitations as the experimental, data-driven studies conducted by others. First, the experimental nature of this work means that an inability to reproduce the data collection process is among the core threats to validity. We tried

[6]The number of times a pair of gamecasts was watched in the same period.

to alleviate this problem by using rich, long-term, large-scale datasets in this study, which offset to some extent the biases of sparser, shorter, and smaller datasets. Future work should perhaps also focus on other games and game genres than this study. Second, much of the work in this study focuses on identifying meaningful correlations, but, due to methodological limitations such as an uncontrolled environment, it cannot establish causal influences between the studied variables. For example, although the number of gamecasts is highly correlated with the popularity of a gamecast creator, we cannot simply infer that building up a large number of gamecasts is guaranteed to cause high popularity for the creator. Future work could focus on new methodologies that allow us to prove causality, for example, by involving small-scale lab experiments in combination with large-scale uncontrolled data collection.

8. CONCLUSION

The growth in popularity and market size of gamecast sharing sites has not been matched by a comprehensive scientific analysis of their characteristics. To address this issue, in this article, we presented a large-scale measurement study on gamecast sharing sites exemplified by the leading and representative sites, WoTreplays and Twitch. We presented the first publicly accessible dataset with both *gamecast statistics* and *in-game statistics*. Based on statistics for more than 7 million gamecasts, more than 1 million gamecast creators, and nearly 2 million players, we investigated the repository, the creators, and the spectators of gamecast sharing sites.

Among our results, we found that both WoTreplays and Twitch are fast-consumed repositories, with a large number of gamecasts being shared, viewed, and soon forgotten. And, in both types of communities we studied, gamecast and creator popularity are highly skewed, with a significant *heavy tail* phenomenon that represents a small number of gamecasts and creators accounting for a large fraction of the total popularity. Furthermore, we observed a noticeable *long tail* phenomenon in WoTreplays, wherein a large number of gamecasts and creators with low popularity accumulate a considerable amount of the total popularity. Last, we also found that creators and spectators have markedly different preferences for sharing and consuming the replays: Whereas creators often emphasize their individual skills, spectators appreciate teamwise tactics. Our findings provide important reference for gamecast sites to improve their services, for example, through designing better recommendation algorithms for unpopular contents to increase their potential of attracting more users.

REFERENCES

H. Abrahamsson and M. Nordmark. 2012. Program popularity and viewer behaviour in a large TV-on-demand. In *Proceedings of the 6th Internet Measurement Conference (IMC'12)*. ACM, Boston, USA, 199–210.

C. Anderson. 2004. The long tail. *WIRED Magazine*. www.wired.com/2004/10/tail. (2004).

R. Aparicio-Pardo, K. Pires, A. Blanc, and G. Simon. 2015. Transcoding live adaptive video streams at a massive scale in the cloud. In *Proceedings of the 6th ACM Multimedia Systems Conference (MMSys'15)*. ACM, New York, USA, 49–60.

G. Armitage, C. Javier, and S. Zander. 2006. Post-game estimation of game client RTT and hop count distributions. In *Proceedings of 5th ACM SIGCOMM Workshop on Network and System Support for Games (NetGames'06)*. Singapore, 1–6.

M. Cha, H. Kwak, P. Rodriguez, Y. Ahn, and S. Moon. 2007. 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 (IMC'07)*. ACM, San Diego, USA, 1–14.

M. Cha, H. Kwak, P. Rodriguez, Y. Ahn, and S. Moon. 2009. Analyzing the video popularity characteristics of large-scale user generated content systems. *Transactions on Networking* 17, 5 (2009), 1357–1370.

C. Chambers, W. Feng, S. Sahu, D. Saha, and D. Brandt. 2010. Characterizing online games. *Transactions on Networking* 18, 3 (2010), 899–910.

K.-T. Chen, Y.-C. Chang, H.-J. Hsu, D.-Y. Chen, C.-Y. Huang, and C.-H. Hsu. 2014. On the quality of service of cloud gaming systems. *Transactions on Multimedia* 16, 2 (2014), 480–495.

X. Cheng, B. C. Burnaby, C. Dale, and J. Liu. 2008. Statistics and social network of YouTube videos. In *16th International Workshop on Quality of Service (WQoS'08)*. IEEE, Enschede, the Netherlands, 229–238.

G. Cheung and J. Huang. 2011. Starcraft from the stands: Understanding the game spectator. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'11)*. ACM, Vancouver, Canada, 763–772.

A. Clauset, C. R. Shalizi, and M. E. J. Newman. 2009. Power-law distributions in empirical data. *SIAM Reviews* 51, 4 (2009), 661–703.

M. Claypool, D. Finkel, A. Grant, and M. Solano. 2014. On the performance of OnLive thin client games. *Multimedia Systems* 20, 5 (2014), 471–484.

J. Davidson, B. Liebald, J. Liu, P. Nandy, and T. van Vleet. 2010. The YouTube video recommendation system. In *Recommender Systems (RecSys'10)*. ACM, Barcelona, Spain, 293–296.

J. Downs, F. Vetere, S.Howard, and S. Loughnan. 2013. Measuring audience experience in social videogaming. In *Proceedings of the 25th Australian Computer-Human Interaction Conference: Augmentation, Application, Innovation, Collaboration (OzCHI'13)*. ACM, Adelaide, Australia, 217–220.

F. Figueiredo, F. Benevenuto, and J. M. Almeida. 2011. The tube over time: Characterizing popularity growth of youtube videos. In *Web Search and Data Mining (WSDM'11)*. ACM, Hongkong, China, 745–754.

D. Fitzgerald and D. Wakabayashi. 2014. Apple quietly builds new networks. *Wall Street Journal*. http://www.wsj.com/news/articles/SB10001424052702304851104579361201655365302. (2014).

P. Gill, M. Arlitt, Z. Li, and A. Manhanti. 2007. YouTube traffic characterization: A view from the edge. In *Proceedings of the 7th ACM SIGCOMM Conference on Internet Measurement (IMC'07)*. ACM, San Diego, USA, 15–28.

Y. Guo and A. Iosup. 2012. The game trace archive. In *Proceedings of the 11th Annual Workshop on Network and Systems Support for Games (NetGames'12)*. IEEE press, Venice, Italy, 4:1–4:6.

W. A. Hamilton, O. Garretson, and A. Kerne. 2014. Streaming on twitch: Fostering participatory communities of play within live mixed media. In *Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems (CHI'14)*. ACM, Toronto, Canada, 1315–1324.

T. Henderson and S. Bhatti. 2001. Modelling user behaviour in networked games. In *Proceedings of the Ninth ACM International Conference on Multimedia (MM'01)*. ACM, Ottawa, Canada, 212–220.

Y. Huang, T. Z. J. Fu, D. M. Chiu, J. C. S. Lui, and C. Huang. 2008. Challenges, design and analysis of a large-scale P2P-VoD system. In *Proceedings of the ACM SIGCOMM 2008 Conference on Data Communication (SIGCOMM'08)*. ACM, Seattle, USA, 375–388.

A. Iosup, R. van de Bovenkamp, S. Shen, A. L. Jia, and F. Kuipers. 2014. An analysis of implicit social networks in multiplayer online games. *Internet Computing* 18, 3 (2014), 36–44.

A. L. Jia, S. Shen, R. van de Bovenkamp, F. Kuipers, A. Iosup, and D. H. J. Epema. 2015. Socializing by gaming: Revealing social relationships in multiplayer online games. *Transactions on Knowledge Discovery from Data* 10, 2 (2015), 11:1–11:29.

C. Kang. 2014. He wants to make it playing video games on Twitch. But will people pay to watch? http://www.washingtonpost.com/. (2014).

M. Kaytoue, A. Silva, and C. Raissi. 2012. Watch me playing, I am a professional: A first study on video game live streaming. In *Proceedings of the 21st International Conference on World Wide Web (WWW'12 Companion)*. ACM, Lyon, France, 1181–1188.

H. Kim and K. Kim. 2014. Learning to recommend game contents for real-time strategy gamers. In *2014 IEEE Conference on Computational Intelligence and Games (CIG'14)*. IEEE press, Dortmund, Germany, 1–8.

J. Li, D. Li, Y. Ye, and X. Lu. 2015. Efficient multi-tenant virtual machine allocation in cloud data centers. *Tsinghua Science and Technology* 20, 1 (2015), 81–89.

X. Lu, H. Wang, J. Wang, J. Xu, and D. Li. 2013. Internet-based virtual computing environment: Beyond the data center as a computer. *Future Generation of Computer Systems* 29, 1 (2013), 309–322.

A. Mahanti, N. Carlsson, A. Mahanti, M. Arlitt, and C. Williamson. 2013. A tale of the tails: Power-laws in internet measurements. *IEEE Network* 27, 1 (2013), 59–64.

J. McGonigal. 2011. *Reality Is Broken: Why Games Make Us Better and How They Can Change the World*. Jonathan Cape. 1–400 pages.

V. Nae, A. Iosup, and R. Prodan. 2011. Dynamic resource provisioning in massively multiplayer online games. *Transactions on Parallel and Distributed Systems* 22, 3 (2011), 380–395.

M. E. J. Newman. 2006. Power laws, pareto distributions and Zipf's law. *Contemporary Physics* 46 (2006), 323–351.

H. Pinto, J. M. Almeida, and M. A. Goncalves. 2013. Using early view patterns to predict the popularity of youtube videos. In *Proceedings of the Sixth ACM International Conference on Web Search and Data Mining (WSDM'13)*. ACM, Rome, Italy, 365–374.

K. Pires and G. Simon. 2015. YouTube live and twitch: A tour of user-generated live streaming systems. In *Proceedings of the 6th ACM Multimedia Systems Conference (MMSys'15)*. ACM, Portland, Oregon, 225–230.

J. A. Pouwelse, P. Garbacki, D. H. J. Epema, and H. J. Sips. 2005. The BitTorrent P2P file-sharing system: Measurements and analysis. In *4th International Workshop on Peer-to-Peer Systems (IPTPS'05)*. Springer, Ithaca, USA, 205–216.

Y. Qu, G. Y. Keung, C. Lin, J. Liu, B. Li, S. Xie, and X. Zhang. 2008. Inside the new coolstreaming: Principles, measurements and performance implications. In *The 27th Conference on Computer Communications (INFOCOM'08)*. IEEE press, Phoenix, USA, 1031–1039.

S. Shen and A. Iosup. 2011. The XFire online meta-gaming network: Observation and high-level analysis. In *Massively Multiuser Virtual Environments (MMVE'11)*. IEEE press, Qinghuangdao, China, 1–6.

C. Spearman. 1904. The proof and measurement of association between two things. *American Journal of Psychology* 15 (1904), 72–101.

M. Suznjevic and M. Matijasevic. 2013. Player behavior and traffic characterization for MMORPGs: A survey. *Multimedia System* 19, 3 (2013), 199–220.

Twitch. 2011. (2011). http://www.twitch.tv/.

R. van de Bovenkamp, S. Shen, A. Iosup, and F. A. Kuipers. 2013. Understanding and recommending play relationships in online social gaming. In *Fifth International Conference on Communication Systems and Networks (COMSNETS'13)*. IEEE press, Bangalore, India, 1–10.

WoTreplays. 2013. (2013). http://www.wotreplays.com.

H. Yu, D. Zheng, B. Y. Zhao, and W. Zheng. 2006. Understanding user behavior in large-scale video-on-demand systems. In *Proceedings of the 1st ACM SIGOPS/EuroSys European Conference on Computer Systems 2006 (EuroSys'06)*. ACM, Leuven, Belgium, 333–344.

Received September 2015; revised April 2016; accepted April 2016