

Distilling Superior Peers in Large-Scale P2P Streaming Systems

Zimu Liu, Chuan Wu, Baochun Li
Department of Electrical and Computer Engineering
University of Toronto
{zimu, chuanwu, bli}@eecg.toronto.edu

Shuqiao Zhao
Multimedia Development Group
UUSee, Inc.
zhaosq@uusee.com

Abstract—In large-scale peer-to-peer (P2P) live streaming systems with a limited supply of server bandwidth, increasing the amount of upload bandwidth supplied by peers becomes critically important to the “well being” of streaming sessions in live channels. Intuitively, two types of peers are preferred to be kept up in a live session: peers that *contribute* a higher percentage of their upload capacities, and peers that are *stable* for a long period of time. The fundamental challenge is to identify, and satisfy the needs of, these types of “superior” peers in a live session, and to achieve this goal with minimum disruption to the traditional pull-based protocols that real-world live streaming protocols use.

In this paper, we conduct a comprehensive and in-depth statistical analysis based on more than 130 GB worth of run-time traces from hundreds of streaming channels in a large-scale real-world live streaming system, UUSee (among the top three commercial systems in popularity in mainland China). Our objective is to discover critical factors that may influence the longevity and bandwidth contribution ratio of peers, using survival analysis techniques such as the Cox proportional hazards model and the Mantel-Haenszel test. Once these influential factors are found, they can be used to form a *superiority index* to distill superior peers from the general peer population. The index can be used in any way to favor superior peers, and we simulate the use of a simple ranking mechanism in a natural selection algorithm to show the effectiveness of the index, based on a replay of real-world traces from UUSee.

I. INTRODUCTION

Real-world live P2P multimedia streaming systems have been successfully deployed in the Internet at a large scale, with hundreds of channels and hundreds of thousands of users at any given time. A salient advantage of P2P streaming systems is the use of peer upload bandwidth contribution to complement bandwidth supplies from a limited number of dedicated streaming servers, mitigating their load and allowing better scalability.

An important test to evaluate the effectiveness of live P2P streaming protocols is how well a streaming system can scale up to a number of peers beyond the threshold of what the servers can sustain. Even though it is theoretically possible to scale to an unlimited number of peers, as real-world peers suffer from upload bandwidth limitations such as NAT traversal challenges and firewalls, real-world systems with a limited pool of server bandwidth supplies can only sustain a limited number of peers, spreading across hundreds of channels. When such a sustainable threshold is reached and then breached, some of the peers are bound to suffer from reduced streaming quality, characterized by the continuity of playback.

In such critical conditions, such as flash crowds, it is important to preserve and sustain the streaming quality of two types of “superior” peers: those who contribute a higher ratio of their upload bandwidth, and those who are stable for a long

period of time. The best peers would, of course, have both excellent longevity and high upload contribution ratio. While favoring such superior peers may sacrifice a small portion of non-superior peers in terms of streaming quality, it must be realized that if these superior peers are not given preference, the risks of their departure would further disturb the stability and bandwidth supply-demand balance in the system, which may escalate to even higher levels of volatility. Such deterioration of stability and deficit of bandwidth may spiral out of control.

In this paper, we conduct a comprehensive and in-depth statistical analysis based on more than 130 GB worth of run-time traces from hundreds of streaming channels in a large-scale real-world P2P live streaming system, UUSee [1]. The system that UUSee Inc. operates is among the top three commercial systems in mainland China, along with PPLive and PPStream. Our objective is very clear: we wish to identify critical performance metrics as risk factors that may influence the longevity and bandwidth contribution ratio of peers. To achieve this, we have parsed and imported all run-time traces into a database, where we apply survival analysis techniques such as the Cox proportional hazards model and the Mantel-Haenszel test to discover such influential factors.

Why do we need to discover performance metrics that may influence the longevity and upload contribution ratio of peers? Our purpose is to *distill superior peers* from the general peer population, such that when push comes to shove as the system scales up, we are able to give priorities to sustain the streaming quality of these superior peers that we have identified. Our objective is to be able to do this without adding additional protocol overhead by redesigning the traditional pull-based protocols being used in real-world systems, such as the pull-based streaming design of UUSee.

The original highlight of this paper is our ability to discover such influential performance factors from large volumes of real-world traces, and to construct a *superiority index* based on these factors and the results of our regression analysis. We argue that such a *superiority index* can be used in any way to favor superior peers. Based on a replay of real-world traces from UUSee, we simulate the use of a simple ranking mechanism in a natural selection algorithm, to show the effectiveness of such a simple way of using the index. To our knowledge, though the idea of distilling and favoring high-value and stable peers is intuitively useful in real-world systems, there has been no research in the literature that predicts the superiority of a peer at an early stage of its session, especially with observations derived from large-scale traces of commercial P2P streaming systems.

The remainder of this paper is organized as follows. In

Sec. II, we present our research methodologies with respect to collecting, parsing, and importing UUSEE run-time traces. In Sec. III, we show our survival analysis of the traces to identify influential performance factors on the peer longevity. In Sec. IV, we shift our analytical focus to the upload contribution ratio of peers. In Sec. V, we construct a *superiority index* based on both longevity and upload contribution ratio of peers, and simulate a simple ranking mechanism in a natural selection algorithm to show the effectiveness of using the index. We discuss related work and conclude the paper in Sec. VI and Sec. VII, respectively.

II. OVERVIEW OF TRACES FROM UUSEE

A. UUSEE P2P Streaming Solutions

Supported by venture capital funding from established firms, UUSEE Inc. [1] is one of the leading P2P multimedia solution providers in mainland China, featuring both legal contractual rights to most of the channels of CCTV, the official Chinese television broadcaster, and online broadcasting rights to 2008 Summer Olympics. With a large collection of streaming servers around the world, it simultaneously broadcasts over 800 live streaming channels to millions of peers, mostly encoded to high quality streams around 500 Kbps. The users of UUSEE are distributed across all the major ISPs in China and over 40 countries in the world.

Similar to all current-generation pull-based P2P streaming protocols, UUSEE's streaming protocol design is based on the principle of allowing peers to serve each other by exchanging blocks of data, which are received and cached in their local playback buffers. The buffer at each peer represents a sliding window of the media channel, containing blocks to be played in the immediate future. The peer buffer size in UUSEE is 500 media blocks, and each block represents 1/3 second of media playback. Once a new peer joins a channel in UUSEE, an initial set of partners (up to 50) is supplied by one of its tracking servers. The peer then establishes TCP or UDP connections with these partners, and buffer maps are exchanged periodically. During this process, it measures the throughput of the connection, and then selects a number of most suitable peers (around 30), from which it actually requests media blocks.

To dynamically monitor the entire live streaming system, we have implemented detailed measurement and reporting capabilities within the UUSEE client application. Each peer collects a set of its vital statistics, and encapsulates them into "heartbeat" reports to be sent to the tracking servers every 5 minutes via UDP. The statistics include its IP address, the channel it is watching, its buffer availability map, the number of consecutive blocks in its current playback buffer (henceforth referred to as the *buffering level*), instantaneous aggregate download and upload throughput from and to all partners, as well as its download and upload bandwidth capacities. The download and upload capacities of a peer are estimated using measurement-based algorithms at the initial buffering stage of the peer [2].

B. A First Glance at Real-World Traces

Though we have been continuously monitoring the performance of UUSEE, the study in this paper features a most recent set of 130 GB worth of run-time traces, collected between Thursday, May 29, 2008 (GMT+8) and Monday, June 2, 2008

(GMT+8), which contain continuous-time snapshots of the streaming system throughout the period. We believe those recent traces best captured the up-to-date characteristics of peers in the millions-of-users scale, to which the application has expanded over the years. A summary of the traces is listed in Table I, including a few basic statistics for all the streaming channels in the traces and for the most popular channel, CCTV News, as an example. Here, a *peer session* refers to the lifetime duration between the joining and the departure of a peer in a streaming channel.

TABLE I
SUMMARY OF UUSEE TRACES

Statistics	All Channels	CCTV News
Number of Reports	125,901,574	4,016,575
Number of Peer Sessions	16,295,086	411,887
Number of Unique IPs	4,057,325	247,361

Without a doubt, it is not suitable to handle such a large volume of traces with the traditional plaintext-based data analysis such as XML parsing, due to its time complexity. Instead, we have chosen to use MySQL, a high-performance database featuring B-tree indexing, in-memory hash table and a user-friendly SQL language, to effectively manipulate volumes of traces and explore hidden correlations within them. Even with a well-tuned MySQL 5.0, deployed on a state-of-the-art server platform with two Dual-Core Intel Xeon 3.0 GHz processors and 2 GB memory, it has taken more than 10 hours to import all of the traces into the database. Besides, SQL queries intended for our statistical analysis may also easily take minutes or even hours. To our knowledge, such in-depth analysis of traces at this scale has rarely been attempted and documented in the P2P streaming literature.

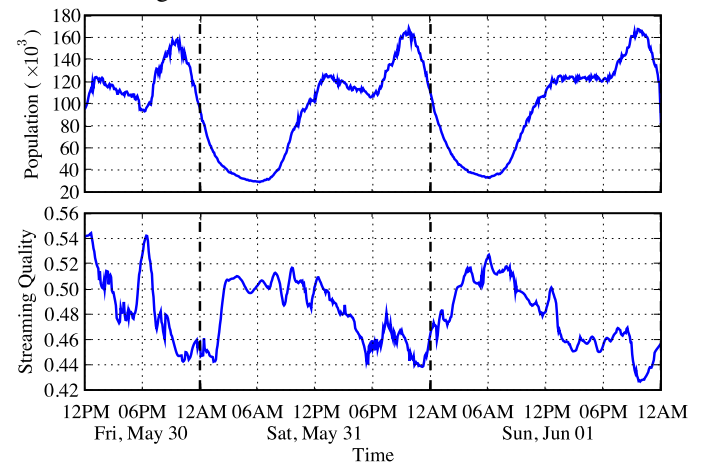


Fig. 1. The evolution of population and streaming quality over all channels.

Why do we still need to refine the P2P streaming protocol while it has been deployed in such a large scale? Our first trace study, the results of which are shown in Fig. 1, exposes an ever existing problem in real-world P2P streaming systems that we cannot ignore: unlike the theoretical expectation of unlimited scalability, the streaming quality in real-world streaming channels downgrades evidently at peak hours of the day with flash crowds of users in the channels. Here, we evaluate the streaming quality in a channel at each time as the *percentage of high-quality peers in the channel*, where a high-quality peer has a buffering level of more than 80% of the total size of its playback buffer. The criterion of the buffering level (*i.e.*,

the number of consecutive blocks in the playback buffer of a peer, starting from the current playback position) has been extensively used in the actual UUsee streaming protocol to evaluate the current streaming quality of a peer. Accordingly, we also use the peer buffering level as our basic streaming quality metric, based on the rationale that the more blocks a peer has cached in its buffer, the higher chance it has to enjoy a smooth playback.

The less than satisfactory streaming performance with an increasing number of peers — mainly ascribed to breaching the sustainable threshold of limited server bandwidth — confirms our supposition on the yet-to-improve scalability and stability of large-scale P2P streaming systems, which constitute the ultimate goal of study in this paper. To promote peer online times for stable bandwidth supplies, we explore critical factors that influence peer longevity; to increase the amount of peer bandwidth supplies for better scalability, we investigate what we can do to improve the percentage of peer bandwidth contribution.

III. PEER LONGEVITY: HAZARD REGRESSION MODELING

In this section, we investigate the influence of various factors on peer longevity in P2P streaming, using survival analysis techniques. We first briefly introduce the basic idea of survival analysis, explore the critical influential factors by survival curve plotting and correlation study, and then derive a Cox proportional hazards model to describe the relationship between peer longevity and its influential factors.

A. Survival Analysis and Censoring

Survival analysis represents a set of statistical methods for the analysis of death or failure events and involves the modeling of time to event data, *i.e.*, the survival time. In our analysis of peer longevity in each streaming channel, a peer's departure represents a *failure* or *death* event, and the time between its joining and departure, the peer longevity (in minutes) is the survival time to be modeled. In survival analysis of many practical applications, *censored* survival data is common, in that we do not know the accurate birth and/or death times of an individual, but have only observed the time before which the individual has been born and the time after which the death occurs. Based on our every-5-minute trace collection methodology, our peer longevity data derived from the traces fall into a specific type of censored data, the *grouped* survival data, in that the observations for birth/failure times are made on fixed intervals, and the obtained large set of survival data features many tied survival times. Fig. 2 gives an illustration for such grouped censorship, in which o_{i-1}, o_i, o_{i+1} are the observation times.




Fig. 2. Grouped survival data: an illustration.

In survival analysis, a *survival function* is frequently used to describe the probability that an individual survives to a specific

time t . Let random variable T represent the longevity of a peer session, our survival function is defined as:

$$S(t) = \Pr(T > t) = 1 - \Pr(T \leq t) = 1 - F(t),$$

where $F(t)$ is the cumulative distribution function (CDF) of the longevity. A standard estimator of the survival function is proposed by Kaplan and Meier, referred to as the product-limit estimator or the K-M estimator [3]. Assume there are p distinct failure times t_1, t_2, \dots, t_p in ascending order such that $t_1 < t_2 < \dots < t_p$. The estimator is then defined as follows:

$$\hat{S}(t) = \begin{cases} 1 & \text{if } t < t_1, \\ \prod_{t_i \leq t} \left(\frac{n_i - d_i}{n_i} \right) & \text{otherwise,} \end{cases}$$

where n_i and d_i denote the number of peers alive by time t_i and the number of peers died in the interval of $(t_{i-1}, t_i]$, respectively.

B. Exploring the Risk Factors

We now seek to investigate critical factors that influence the peer longevity, in order to discover useful insights to improve peer stability.

1) **Streaming Quality:** Intuitively, unfavorable streaming quality, to some extent, may result in the premature departure of peers. Therefore, we start by investigating: *Do peer longevity patterns differ significantly under different levels of streaming quality?* To answer this question, we categorize peer sessions according to the average buffering level achieved throughout each session, and plot the *survival curves*, namely the survival function estimated using the K-M algorithm, of the different groups. Fig. 3 shows the survival curves for three session groups with the average buffering level in the ranges of 0 – 25%, 25 – 75% and 75 – 100% of the total buffer size (500 blocks), respectively. We can observe significant differences across the survival curves, that the session duration is generally larger when the streaming quality is better. We further statistically validate our observations using the Mantel-Haenszel test [4], also referred to as the log-rank test. The log-rank test is commonly applied to test the null hypothesis that a set of survival functions are statistically equivalent, in which the null hypothesis is rejected if the result p -value is lower than the significance level of 0.05. We have derived a log-rank test result of $p = 1 - \Pr_{\chi^2, 2}(14755) \approx 0$, which confirms the significance of the differences among the survival functions for the three groups.




Fig. 3. Survival curves for sessions with different levels of streaming quality.

Not satisfied by only revealing the relevance between the streaming quality and session duration, we take one step further to explore the best statistical metrics of the streaming quality,

that represents the most correlation with the session duration. The factors to be compared are: the average buffering level during a peer session, the standard deviation of the buffering level throughout a peer session, the minimum buffering level during a peer session, and the initial buffering level, as the first buffering level measured when a peer starts its playback. In each sub-figure in Fig. 4, we plot the average and median session durations of each session group at different levels of the respective streaming quality factor, as well as the smoothed LOWESS curves [5].

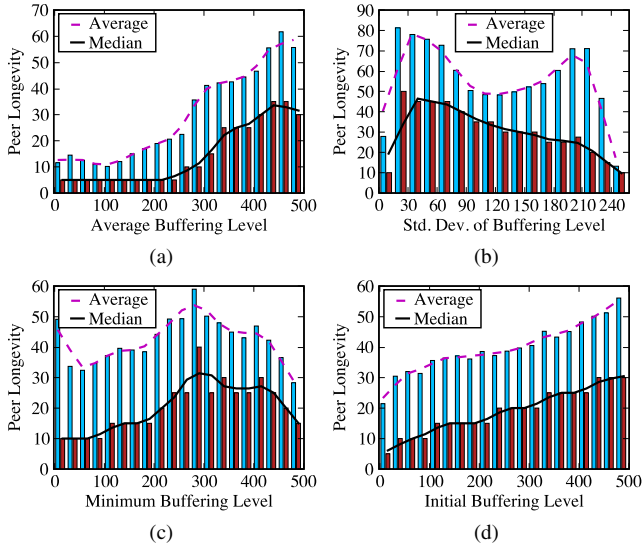


Fig. 4. Correlation of session duration with metrics of streaming quality.

Fig. 4(a) generally reveals a positive correlation between the session duration and the average buffering level, but it is only apparent above a certain threshold value around 225 – 250 for both curves. We also notice that the curves slightly drop in the last range of the buffering level, 475 – 500, which may further reveal the difficulty to maintain a near-full buffer in long streaming sessions, while a slightly less full buffer can already guarantee a smooth viewing experience to keep the peers staying longer.

The negative correlation between the standard deviation of buffering level and session duration, as shown in Fig. 4(b), meets our expectation that the less stable the streaming quality is, the shorter the peers are staying. However, the correlation is not strong and consistent enough for us to select the standard deviation as the best risk factor to decide session duration.

We further investigate the tolerance of peers towards the lowest possible streaming quality and the initial streaming quality, by plotting the correlations with respect to the minimum buffering level and initial buffering level of a session, in Fig. 4(c) and Fig. 4(d), respectively. No dominant correlation is observed in Fig. 4(c), revealing that the minimum buffering level metric is not suitable for our purpose as well. On the other hand, a strong and consistent positive correlation is observed between the session duration and initial buffering level in Fig. 4(d), with respect to both the median and average curves. We further check such linear correlation by computing Pearson’s correlation coefficients; and with the results showing 0.96 and 0.98 for the mean and median respectively, an extremely strong positive linear correlation is suggested. This

observation represents the first interesting discovery in our study that, out of the many possible streaming quality factors, the streaming quality experienced by a peer at the beginning stage of its playback *critically* influences the interest of a peer in a streaming channel, regardless of any original interest towards the possible content of the channel.

2) **Joining Time:** Other than the streaming quality, we expect the time when a peer surfs the Internet also influences its viewing behavior. To investigate how significant the effect of the time is, we categorize the peer sessions according to their start times and compare the survival functions of the resulting session groups.

We first explore any possible effect of the day of the week, using sessions starting at a same time on a weekday, Thursday May 29th, and on the weekend, Saturday May 31st. Fig. 5 exhibits no significant difference between the two session groups, no matter whether the start time is in the morning (10 a.m.) or in the evening (10 p.m.). In addition, the results of log-rank tests, $p = 1 - \Pr_{\chi^2,1}(2.3) \approx 0.06 > 0.05$ and $p = 1 - \Pr_{\chi^2,1}(2.1) \approx 0.07 > 0.05$ for the 10 a.m. and 10 p.m. cases, confirm our observations that there is no significant statistical difference between the two survival functions, respectively.




Fig. 5. Survival curves for detecting the day of the week effect.

Nevertheless, when we zoom into different times on a same day and classify sessions according to the hours they start, the effect of time becomes more evident. Fig. 6 exhibits visible differences among survival curves for session groups of four different starting times on May 30th: 15.3% of the peers starting around 10 p.m. stay for more than an hour, while only 10% of those starting around 10 a.m. have such a longevity. The log-rank test results of $p = 1 - \Pr_{\chi^2,3}(20.7) \approx 0$, rejects the null hypothesis that survival functions are equivalent and validates our observations.

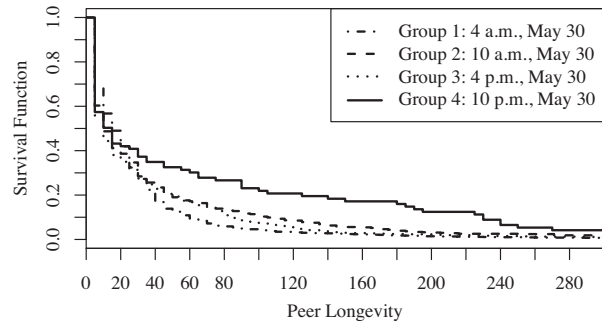


Fig. 6. Survival curves for sessions starting at different times of a day.

In addition, we have investigated the differences among survival curves of other times of the day and among different days. The similar log-rank test results further confirm our second interesting discovery: the peer joining time during one

day's course significantly decides how long it can stay, *e.g.*, peers can afford to watch the streaming channels longer at their evening leisure times; instead, the expected day of the week effect, that peers may stay longer during the weekend, turns out not to be important.

3) **Channel Popularity:** We next investigate whether peer longevity may differ significantly across channels of different popularity, by plotting the correlation between channel population and average session duration across all of the 800 streaming channels in Fig. 7. Each sample in the figures represents one streaming channel. The population of each channel is computed as the average concurrent number of peers in the channel over the trace period. To eliminate the time effect, the average session duration for each channel is evaluated using peer sessions in the channel, which start within one specific hour, *e.g.*, 9 a.m. and 9 p.m. on May 30th, as shown in the figures respectively.




Fig. 7. Correlation between channel population and average longevity in log-log scale, with the linear fitting line and 95% confidence intervals.

Fig. 7 exhibits an evident positive linear correlation between the two quantities at the log-log scale, *i.e.*, the more popular a channel is, the longer the peers may stay. We have computed the Pearson product-moment correlation coefficients, and the p -values for testing the null hypothesis that there is no significant correlation between the two quantities. Both the Pearson's coefficients which gives 0.762 and 0.695 for 9 a.m. and 9 p.m. respectively, and p -values that are extremely close to zero, further validate the significance of the correlation. Therefore, the channel popularity is identified as the third important risk factor to decide peer longevity.

C. Regression Modeling

With three important influential factors identified, we now seek to model the peer longevity as a function of these factors, using the regression techniques in survival analysis.

1) **The Cox Regression Model:** The Cox proportional hazards model [6] is a classical regression model for the analysis of censored survival time with respect to their relationship with covariates (which is the terminology in Cox modeling for influential or risk factors). It models the relationship between the covariates and censored survival time based on the *hazard function*. A hazard function $h(t)$, also referred to as the *hazard rate*, represents the instantaneous failure rate for a session that has survived to time t . Let T denote the duration of a survival session. The hazard function is defined as:

$$h(t) = \lim_{\Delta t \rightarrow 0} \frac{\Pr(t \leq T \leq t + \Delta t | T \geq t)}{\Delta t}.$$

In Cox regression modeling, it models the hazard rate at time t for a session with covariate vector $\mathbf{z} = (z_1, \dots, z_p)$ as a function of a baseline hazard function and the risk factors. The basic Cox model is:

$$h(t; \mathbf{z}) = h_0(t) \exp(\beta^T \mathbf{z}) = h_0(t) \exp\left(\sum_{k=1}^p \beta_k z_k\right) \quad (1)$$

where $h(t; \mathbf{z})$ is the hazard rate at time t for a session with covariate vector \mathbf{z} ; $h_0(t)$ is an arbitrary non-negative baseline hazard function, which is computed during the regression process; and $\beta = (\beta_1, \dots, \beta_p)$ is a column p -vector of coefficients corresponding to the covariates in \mathbf{z} . In our regression modeling, the covariates are selected corresponding to the influential factors on peer longevity we have observed in the previous section: one continuous variable, BUF, is used to represent initial buffering level at the peer. Another continuous variable, POP, corresponds to the population of the channel the peer is in. To represent the time of day effect of peer joins, we divide the time in a day to 24 intervals, each corresponding to one hour, and use 24 indicator variables, $\text{TOD}_i, i = 0, \dots, 23$, to denote the joining time of a peer. For example, if the peer joins its channel between 9 a.m. and 10 a.m., we have $\text{TOD}_9 = 1$ and $\text{TOD}_i = 0, \forall i \in \{0, \dots, 23\} \setminus \{9\}$. The covariates, along with their description and type, are listed in Table II. Altogether, a covariate vector with $p = 26$ components is used in our Cox regression.

TABLE II
COVARIATES IN COX REGRESSION MODEL

Covariate	Description	Type
BUF	Initial buffering level	Continues
POP	Channel population	Continues
$\text{TOD}_i, i = 0, \dots, 23$	Joining time of the day	Binary

The Cox regression based on the model in (1) derives the values of regression coefficients $\beta_i, i = 1, \dots, p$. As a probabilistic model, we can then estimate the probability that a session lasts to any specific time t (*i.e.*, the survival curve of the session), given the values of the covariates for the session and using the derived coefficients. In the following section, we first discuss how we estimate β , and then we show how the derived Cox model can be used in the estimation of session duration in Sec. III-C3.

2) **Maximum Likelihood Estimation:** To derive the regression coefficients in the Cox model in Equation (1) based on our large sets of grouped survival times, *i.e.*, the peer session durations derived from the every-five-minute traces, we employ a maximum likelihood estimation algorithm similar to that proposed by Prentice *et al.* [7] for grouped data version of the Cox model. In this model, peer session durations are grouped into intervals $A_i = [a_{i-1}, a_i), i = 1, \dots, r$ with $a_0 = 0, a_r = \infty$. Assuming all sessions start at a_0 , the durations of sessions ended in A_i are all recored as t_i . From (1), the probability of observing a session duration t_i on a peer session with covariate vector \mathbf{z} is

$$[1 - \alpha_i^{\exp(\beta^T \mathbf{z})}] \prod_{j=1}^{i-1} \alpha_j^{\exp(\beta^T \mathbf{z})} \quad (2)$$

where

$$\alpha_j = \exp\left(-\int_{a_{j-1}}^{a_j} h_0(u) du\right)$$

is the conditional survival probability in A_j for an individual with $\mathbf{z} = 0$. The likelihood function of the regression model in (1) is the product of terms (2) over all the peer sessions in the

traces.

Let $\gamma_j = \log(-\log \alpha_j)$. The logarithm of the likelihood contribution from each session, as in (2), can be written as

$$l = \log(1 - \exp(-\exp(\gamma_k + \beta^T \mathbf{z}))) - \sum_{j=1}^{k-1} \exp(\gamma_j + \beta^T \mathbf{z}). \quad (3)$$

In this way, the original Cox regression problem is converted to a maximum likelihood estimation (MLE) problem, in which we estimate the parameters $(\hat{\gamma}, \hat{\beta})$ that maximize the likelihood represented as the sum of terms (3) over all the peer sessions, based on the data of their session durations and risk factor values. To carry out the maximum likelihood estimation, the Newton-Raphson approach can be easily applied.

In Table III, we present the results of maximum likelihood estimation, using data from 231570 sessions in our traces. The sessions used are from 20 channels randomly selected from over 800 channels in the traces. The purpose for such sampling is not only to expedite the speed of maximum likelihood estimation, but also to exhibit the usefulness of our model, trained using only a limited set of samples, as is to be illustrated in the following section. Due to space constraints, we have chosen to show the estimated values of selected components of $\hat{\beta}$ and $\hat{\gamma}^1$, along with their standard errors. The negative $\hat{\beta}$ values corresponding to BUF and POP validate our earlier observations on their positive effects on the session duration: the higher the initial buffering level is and the more popular the respective channel is, the lower the failure probability of a peer session is, and thus the session duration could be longer. Although the p -value for each β is not given in the table, due to the space limit, all p -values are far below 0.05, suggesting all covariates are significant.

TABLE III
MAXIMUM LIKELIHOOD FIT OF COX REGRESSION MODEL IN (1)

Covariate	$\hat{\beta}$	Std. Err.	Intervals	$\hat{\gamma}$	Std. Err.
BUF	-0.0310	1.19e-07	0-1	-0.2953	1.13e-04
POP	-0.2787	1.53e-05	2-3	-1.3747	1.62e-04
TOD ₀	0.8034	4.45e-04	4-5	-1.5993	1.95e-04
TOD ₄	0.0586	1.05e-03	6-7	-1.7477	2.40e-04
TOD ₈	0.4421	3.57e-04	8-9	-1.6648	2.58e-04
TOD ₁₂	0.6915	3.02e-04	10-11	-1.7799	3.13e-04
TOD ₁₆	0.7540	3.40e-04	12-13	-1.7433	3.49e-04
TOD ₂₀	0.7000	3.11e-04	14-15	-1.8224	4.22e-04

The use of the Cox regression model is based on a presumed proportional hazards assumption [6], that the ratio of the hazard rates of two sessions is only dependent on their covariate values, but independent of time. As a check on the proportional hazards assumption, we apply the likelihood ratio test [8]. The result of the test gives a very small χ^2 value of 376 on 26 d.f., which is far below the 0.05 significance level. Therefore, the proportional hazards assumption stands, and the appropriateness of using the Cox regression model with the covariates in Table II is confirmed.

3) **Longevity Prediction and Model Validation:** With the Cox regression model established, we can now derive the survival curve of a session with covariate vector \mathbf{z} . For the grouped Cox model, the estimator of the survival function with

¹The maximum index of $\hat{\gamma}$ is decided by dividing the maximum session duration from the traces by the trace collection interval of 5 minutes.

covariate vector \mathbf{z} at time t_k can be written as:

$$\hat{S}(t_k; \mathbf{z}) = \prod_{j=1}^{k-1} \exp(-\exp(\hat{\gamma}_j + \hat{\beta}^T \mathbf{z})).$$

We use the expected session time in the survival curve of a session with \mathbf{z} , as the most probable session duration of the session. In this way, given a covariate vector \mathbf{z} , we are able to predict the most probable session duration using our Cox regression model.

Recall that our regression model is trained using only a limited set of session data from 20 random channels. We now evaluate its accuracy in estimating the duration of sessions from other channels, with representative results in Fig. 8.

Fig. 8 plots the measured median durations, along with the predicted durations in highlighted (solid) curves, of peer sessions at different levels of initial buffering level (BUF), with four different cases of POP and TOD: Fig. 8(a) and (b) plot sessions started around 10 p.m. (TOD₂₂ = 1), in a popular channel (POP = 3162) and in a less popular channel (POP = 251), respectively; and Fig. 8(c) and (d) plot sessions in the same channel (with POP = 632), starting at 10 a.m. (TOD₁₀ = 1) and 10 p.m. (TOD₂₂ = 1), respectively. Predictions in all four figures fall into the 75% confidence intervals (in dash lines), which validates the usefulness of our regression model — derived using a small portion of session data — in the accurate prediction of session durations in the entire traces.

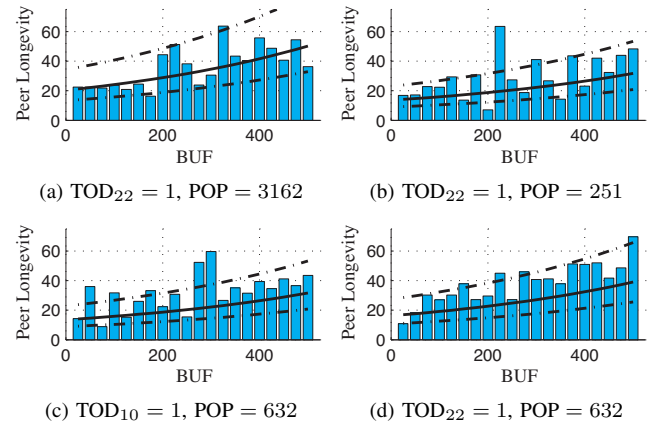


Fig. 8. Regression model validation by prediction of the peer longevity.

The peer longevity model derived in this section provides a useful tool for the estimation and promotion of peer session duration in P2P streaming, in order to enhance the stability and eventually the streaming quality in the system. Practically, we may only be interested in promoting the longevity of high-value peers, *i.e.*, the ones who can contribute a high level of upload bandwidth. A question arises: how can we decide, especially at the early stage after peer joins, which peer can contribute significantly throughout its session time? In the following section, we seek to find the answer to this question.

IV. BANDWIDTH CONTRIBUTION RATIO: INFLUENTIAL FACTORS

Intuitively, a peer may upload more when it inherently has a large upload capacity, *e.g.*, the case of an Ethernet peer. Nevertheless, it is not necessarily so that a peer is always able

to contribute all its upload bandwidth for streaming. In seeking the causes, we do not consider the selfishness of peers, as none of the well-known real-world P2P streaming applications implement any mechanism allowing peers to decide its bandwidth contribution level; we also exclude the possibility of significant protocol inefficiency in UUse P2P streaming, based on our knowledge of its elaborated peer selection and NAT/firewall traverse algorithms implemented to maximize peer bandwidth utilization. Indeed, our focus is on the “objective” factors, such as the network condition and the streaming quality, that may have led to low levels of peer bandwidth contribution. Considering most ISPs in China confine the upload capacity to 512 Kbps, we define 512 Kbps as the *threshold capacity* in this paper, and introduce the *bandwidth contribution ratio*, the ratio of contributed bandwidth at a peer over the threshold capacity, to measure the normalized upload bandwidth utilization in the P2P system. Note that in special cases where the ratio is greater than 1, they are considered as 1 for the convenience of our study. In this section, we seek to explore the important factors that decide the average bandwidth contribution ratio over a session’s course, in order to derive useful insights to identify potential high-contribution peers using the combination of the ratio and upload capacity of each peer.

A. Peer Longevity

Long-lived peers, who are online for a long period of time, are generally regarded as superior peers, not only due to their stability, but also based on a hidden assumption: those stable peers may contribute more of their upload capacity to the P2P streaming. But is this assumption true? To investigate this question, we plot in Fig. 9 the average and median bandwidth contribution ratios of sessions at different levels of the session duration, as well as the smoothed LOWESS curves. The bandwidth contribution ratio of each session is computed by dividing the average upload bandwidth of the peer during the session, whose measurement methodology was mentioned in Sec. II-A, by the threshold capacity of 512 Kbps. No significant correlation has been observed in Fig. 9, which suggests long session duration and high bandwidth contribution level do not necessarily happen altogether at the same peers, and also turns away peer longevity from our influential factor candidate set.




Fig. 9. Correlation of bandwidth contribution ratio with session duration.

We now turn to factors that may fundamentally limit the upload ability of a peer, *e.g.*, low buffering level that represents a limited number of blocks to serve other peers, or inter-peer bandwidth bottlenecks that prohibit the full utilization of last-mile capacities.

B. Streaming Quality

We first investigate the effect of buffering level on the bandwidth contribution ratio of a peer. Fig. 10 plots the

average and median bandwidth contribution ratios of session groups at different levels of streaming quality, in terms of the average buffering level throughout the session duration and the initial buffering level of the sessions, respectively. A positive correlation can be observed in both plots, confirming our guess on the impact of peer buffering level on the utilization of its upload capacity. Similar to Fig. 4(a), the correlation in Fig. 10(a) is more evident when the average buffering level is beyond a threshold value around 225 – 250. On the other hand, Fig. 10(b) reveals a stronger linear positive correlation between bandwidth contribution ratio and the initial buffering level, with respect to both the median and average curves. This surprising observation, that the *initial* buffering level of a peer session exhibits better correlation with the *average* bandwidth contribution ratio throughout the session, other than the *average* buffering level, represents an interesting discovery in our study, which works in favor of us: the average bandwidth contribution level of a peer during its lifetime may be quite accurately predicated using its initial streaming quality upon joining, such that a number of measures can be taken immediately to favor those high-contribution peers.




Fig. 10. Correlation of bandwidth contribution ratio with streaming quality.

C. ISP Membership

We next investigate any bottleneck effect on the inter-peer links, that may limit the utilization of the last-mile capacities at the peers. Our previous study in [2] has revealed the existence of bandwidth bottlenecks along inter-ISP P2P links in UUse P2P streaming network. Considering that peers in larger ISPs have a higher percentage of intra-ISP links while peers in small ISPs may mostly connect to partners in other ISPs, we wonder whether the achievable bandwidth utilization ratios represent any difference among peers in ISPs of different sizes. To investigate this issue, we categorize peer sessions according to their ISP membership, and compare the bandwidth contribution ratios of different session groups. As we have identified the significant impact of initial buffering level on the ratio, we plot the correlation of average bandwidth contribution ratio and initial buffering level for each session group respectively, and compare the resulting plots in Fig. 11. In each of the figures, we not only show the smoothed LOWESS curves, but also plot the fitted linear regression line (in dash) of the average bandwidth contribution ratios at different levels of initial buffering level. The ISP each session group corresponds to is marked at the upper-left corner. Note that the ISPs are listed in the descending order of their peer population. Statistics from the linear regression analysis of each figure are given in Table IV.

From Table IV, we observe small values of the slopes for the regression lines in Fig. 11, due to the large difference in the

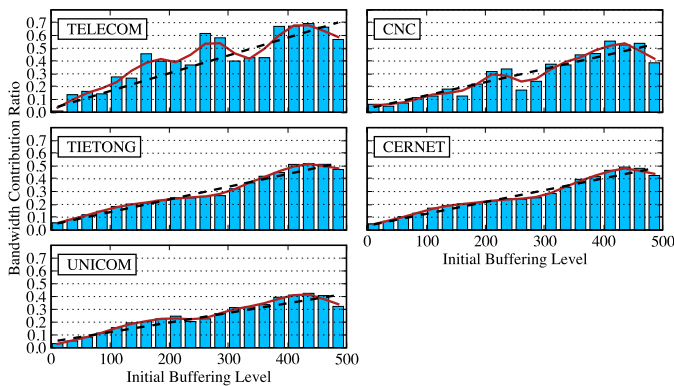


Fig. 11. Correlation of bandwidth contribution ratio with initial buffering level in different ISPs.

TABLE IV
LINEAR REGRESSION STATISTICS FOR SESSIONS IN DIFFERENT ISPS

ISP	Size	Slope	Intercept	Corr. Coef.	p -value
TELECOM	8×10^5	1.39e-3	4.24e-2	0.67	≈ 0
CNC	5×10^5	1.03e-3	3.93e-2	0.41	≈ 0
TIETONG	2×10^4	9.91e-4	5.14e-2	0.49	1e-279
CERNET	7×10^3	9.25e-4	4.36e-2	0.50	7e-119
UNICOM	6×10^2	7.51e-4	5.51e-2	0.52	3e-95

magnitudes of bandwidth contribution ratios and buffering levels. Nevertheless, the positive correlation coefficients, together with the near-0 p -values, validate the significance of the slopes, and also confirm the significance of the correlation between initial buffering level and bandwidth contribution ratio in the cases of each ISP. The most interesting observation we can make from Fig. 11 and Table IV is: the slope of the regression line is steeper for ISPs of a larger size, and is flatter in the cases of smaller ISPs. In another word, at a same initial buffering level, peers in larger ISPs may have a higher bandwidth contribution ratio than those in smaller ISPs. Considering peers in larger ISPs may have more neighbors in the same ISPs than those in small ISPs, such an observation can be explained by the less impact of inter-ISP bandwidth bottleneck on the upload capacity utilization at peers in larger ISPs than those in small ISPs. This represents another interesting discovery in our study of influential factors to bandwidth contribution ratio, that the ISP membership of a peer also decides its ability to utilize its upload capacity, which we may also make use of in the selection of high-contribution peers.

V. SUPERIORITY INDEX: A SIMPLE PEER RANKING MECHANISM

The regression model for peer longevity and influential factors to the bandwidth contribution ratio bring useful insights towards the improvement of stability and scalability of large-scale P2P streaming systems. As an important application, we propose a *Superiority Index*, for distilling superior peers during streaming, which can potentially stay in the system for a long time and contribute a high level of upload bandwidth. The index is defined in the following fashion:

$$\text{Superiority Index} = \text{Predicted Peer Longevity} \times \text{Estimated Average Upload Bandwidth}$$

As the product of estimated peer session duration and average upload bandwidth during the session, the Superiority Index

estimates the potential overall bandwidth contribution at a peer during its session time. Such a Superiority Index can be used to design a simple ranking mechanism in a natural peer selection algorithm, that augments the current P2P streaming protocol, as follows.

In a P2P streaming protocol such as UUSee, new peers connect to a set of existing peers randomly assigned by the tracking server, and each existing peer would treat all new connected peers equivalently and divide its upload bandwidth among them. In our proposed natural selection algorithm, after each new peer has connected and obtained its initial streaming bandwidth, the upstream peer will decide on the potential contribution of the peer by computing its Superiority Index; it then ranks all the new peers connected to itself by their superior indices, and only those peers with large superior indices are kept as the neighbors, while those with small index values will be disconnected. The rationale behind this peer selection process is that, given the limited upload bandwidth in the system, we may only wish to keep the superior peers with potentially better stability and upload bandwidth contribution, and cut off inferior peers as soon as possible, for better stability and scalability of the entire system.

In computing the Superiority Index of each peer, the peer longevity is predicated using the Cox regression model we derived in Sec. III-C1, based on the initial buffering level the peer has experienced, the channel it is in and its joining time of the day. The average upload bandwidth is estimated by multiplying the bandwidth contribution ratio we discussed in Sec. IV, with the threshold capacity of 512 Kbps if the peer is an ADSL peer (as in China the upload capacity of a major portion of ADSL peers are limited to 512 Kbps by ISPs), or with the total upload capacity of the peer in the case of Ethernet peers. The bandwidth contribution ratio of a peer can be derived using its ISP membership and initial buffering level, based on the linear relationships we have derived Sec. IV between the ratio and the initial buffering level in each ISP.

To evaluate the effectiveness of the Superiority Index in distilling superior peer and promoting the “well-being” of the streaming system, we have implemented the natural selection algorithm in a P2P streaming system, that emulates UUSee protocols and replays the real-world scenarios captured by the traces. In our experiments, we emulate the streaming of channel CCTV1 over one day’s course with the same number and ISP distribution of participating peers, as captured in the traces of May 30, 2008. We also emulate peer dynamics by having peers join and depart from the channel following the peer arrival times and session durations derived from the traces. The upload and download capacities of each peer are generated according to their respective distributions summarized from the traces as well.

In our experiments, we run the P2P streaming system without and with the natural selection in place, and compare the average streaming quality in the channel achieved over time. Given the evolution of the peer population in the channel shown in Fig. 12(a), Fig. 12(b)(c) plot the achieved average streaming quality in the two cases, respectively. We can observe that while the streaming quality in Fig. 12(b) drops significantly whenever there is a large peer crowd, the streaming quality downgrade in a system with the natural selection protocol,

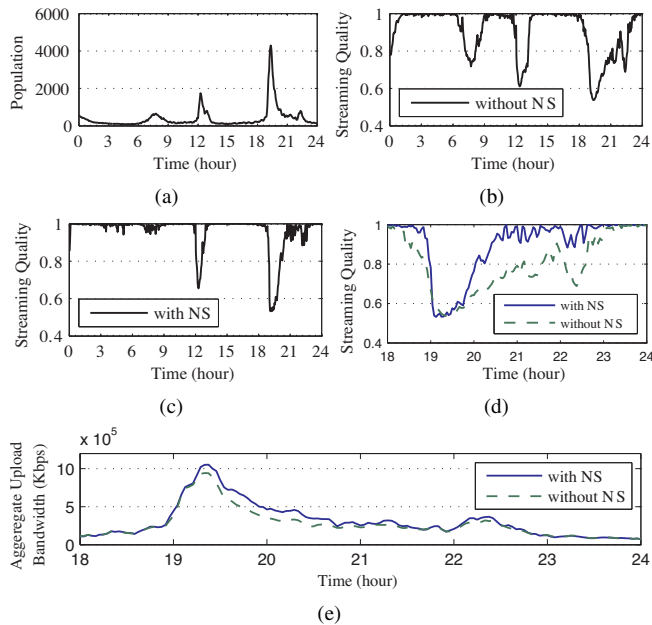


Fig. 12. Comparison of streaming performance with and without the natural selection (NS) algorithm.

as shown in Fig. 12(c), is significantly less. In addition, the streaming quality can much more quickly recover its original level whenever a flash crowd occurs, which can be more clearly observed from a zoom-in to the peak hours of 6 p.m. to 12 a.m. in the two sub figures, as plotted in Fig. 12(d).

In addition, Fig. 12(e) plots the aggregate upload bandwidth of peers in the streaming channel during the peak hours, in the cases with and without the natural selection. We can see the aggregate bandwidth contribution in the system with the natural selection is generally larger than that without it, due to the promotion of high bandwidth peers in the system.

All the above observations exhibit the effectiveness of the proposed superiority index in a natural selection algorithm, which effectively promotes the session duration of high contribution peers, thus enhancing the overall streaming quality in the P2P streaming system.

VI. RELATED WORK

In recent years, significant research efforts have been devoted to the measurement and improvement of real-world P2P systems. To investigate the scalability and stability of these systems, existing studies mostly focus on the measurement and characterization of the P2P topologies [9], throughput levels [2], and peer churns [10].

With respect to measurements related to peer longevity, Wang *et al.* [11] justified the importance of stable peers in an analytical model for the P2P system. Stutzbach *et al.* [12] characterized peer arrivals and departures in three popular P2P file-sharing systems (BitTorrent, Kad and Gnutella). Hei *et al.* [10] measured the distribution of peer life time in a commercial P2P streaming system, PPLive, while Li *et al.* [13] also revealed that the peer session durations follow a heavy-tailed distribution. Chen *et al.* [14] have characterized the influence of network QoS metrics on peer sessions in a P2P VoIP application, Skype.

Our work distinguishes from all existing measurement works in the following two important aspects. *First*, rather than simply

characterizing their distributions, our work focuses on the identification of critical influential factors to the peer longevity and bandwidth contribution ratio. *Second*, we use statistical regression techniques to model the relationship between the influential factors and peer longevity or bandwidth contribution ratio, the insights from which can be used to design better P2P streaming protocols that promote both stability and bandwidth contribution of the peers.

VII. CONCLUDING REMARKS

Our focus in this paper is to explore the influential factors to peer longevity and bandwidth contribution level in practical large-scale P2P streaming applications, utilizing over 130 GB worth of real-world traces from a real-world P2P streaming system, UUSEE. As our key contributions, we have not only identified the key influential factors that decide the duration of peer sessions, but also modeled their relationship into a Cox regression model, using a survival analysis approach. Similarly, with respect to the bandwidth contribution ratio at each peer, we have discovered the impact of the peer initial buffering level and ISP membership, and have linearly modeled their correlations. As an important application of our discoveries, we have designed a superiority index for distilling superior peers from the general peer population, and have applied the index in a natural selection algorithm to promote the session duration of high contribution peers. Our evaluations based on a replay of real-world streaming traces validate the effectiveness of this superiority index in improving the overall stability and scalability of the P2P streaming system.

REFERENCES

- [1] UUSEE Inc. [Online]. Available: <http://www.uusee.com>
- [2] C. Wu, B. Li, and S. Zhao, "Characterizing peer-to-peer streaming flows," *IEEE J. Sel. Areas Commun.*, vol. 25, no. 9, pp. 1612–1626, 2007.
- [3] E. L. Kaplan and P. Meier, "Nonparametric estimation from incomplete observations," *Journal of the American Statistical Association*, vol. 53, no. 282, pp. 457–481, 1958.
- [4] M. Tableman and J. Kim, *Survival Analysis Using S: Analysis of Time-To-Event Data*. Chapman & Hall/CRC, 2004.
- [5] W. S. Cleveland, "LOWESS: A program for smoothing scatterplots by robust locally weighted regression," *The American Statistician*, vol. 35, no. 1, p. 54, 1981.
- [6] D. R. Cox, "Regression models and life-tables," *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 34, no. 2, pp. 187–220, 1972.
- [7] R. L. Prentice and L. A. Gloeckler, "Regression analysis of grouped survival data with application to breast cancer data," *Biometrics*, vol. 34, no. 1, pp. 57–67, 1978.
- [8] T. Therneau and P. Grambsch, *Modeling Survival Data: Extending the Cox Model*. Springer, 2000.
- [9] C. Wu, B. Li, and S. Zhao, "Magellan: Charting large-scale peer-to-peer live streaming topologies," in *Proc. the 27th International Conference on Distributed Computing Systems (ICDCS '07)*, 2007, p. 62.
- [10] X. Hei, C. Liang, J. Liang, Y. Liu, and K. Ross, "A measurement study of a large-scale P2P IPTV system," *IEEE Trans. Multimedia*, vol. 9, no. 8, pp. 1672–1687, 2007.
- [11] F. Wang, J. Liu, and Y. Xiong, "Stable peers: Existence, importance, and application in peer-to-peer live video streaming," *Proc. IEEE INFOCOM '08*, pp. 1364–1372, 2008.
- [12] D. Stutzbach and R. Rejaie, "Understanding churn in peer-to-peer networks," in *Proc. 6th ACM SIGCOMM Conference on Internet Measurement (IMC '06)*, 2006, pp. 189–202.
- [13] B. Li, S. Xie, G. Keung, J. Liu, I. Stoica, H. Zhang, and X. Zhang, "An empirical study of the Coolstreaming+ system," *IEEE J. Sel. Areas Commun.*, vol. 25, no. 9, pp. 1627–1639, 2007.
- [14] K.-T. Chen, C.-Y. Huang, P. Huang, and C.-L. Lei, "Quantifying Skype user satisfaction," in *Proc. ACM SIGCOMM '06*, 2006, pp. 399–410.