

Unifying Models of Churn and Resilience for Unstructured P2P Graphs

Zhongmei Yao, *Student Member, IEEE*, Daren B.H. Cline, Xiaoming Wang, *Student Member, IEEE*, and Dmitri Loguinov, *Senior Member, IEEE*

Abstract—Previous analytical results on the resilience of unstructured P2P systems have not explicitly modeled heterogeneity of user churn (i.e., difference in online behavior) or the impact of in-degree on system resilience. To overcome these limitations, we introduce a unifying model of heterogeneous user churn and derive the distribution of the various metrics observed in prior experimental studies. We also show that the arrival process of in-edges to each user converges to Poisson when system size tends to infinity, model transient behavior of in-degree, and apply these results to obtain the joint in/out-degree isolation probability.

Index Terms—Peer-to-peer, Pareto, stochastic lifetime resilience

1 INTRODUCTION

DURING the last decade, peer-to-peer (P2P) networks have rapidly evolved and have become a highly scalable and efficient platform for large-scale distributed applications. One of the central problems in understanding how these systems behave is the analysis of their properties during *user churn*, which is a general term describing arrival/failure of individual nodes in the system and repair algorithms applied by existing users to counteract the effects of abrupt departures. Unlike other distributed systems where failures may be rare, P2P networks constantly remain in the state of churn and embrace frequent failures as part of their normal operation.

One commonly studied problem in the literature is the ability of P2P networks to stay connected in the face of random failures [1], [3], [8], [10], [12], [13], [14], [15], [17], [19], [21], [23], [25]. Compromised connectivity is a fundamental byproduct of churn that directly affects routing efficiency, search latency, overhead, user satisfaction, and many other performance metrics. Unfortunately, most traditional analytical models that consider churn [13], [15], [19], [21] do not capture the inherent heterogeneity of users, the impact of in-degree on the resilience of the system, or the behavior of P2P networks under non-exponential lifetimes. We overcome these limitations in a set of models described in more detail next.

- Z. Yao is with the Department of Computer Science, The University of Dayton, Dayton, OH 45469 USA. E-mail: zyao@udayton.edu.
- D.B.H. Cline is with the Statistics Department at Texas A&M University, College Station, TX 77843 USA. E-mail: dcline@stat.tamu.edu.
- X. Wang is with Amazon.com, Seattle, WA 98101 USA. E-mail: xmwang@gmail.com.
- D. Loguinov is with the Department of Computer Science and Engineering, Texas A&M University, College Station, TX 77843 USA. E-mail: dmitri@cs.tamu.edu.

Manuscript received 11 Feb. 2013; revised 13 June 2013; accepted 15 June 2013. Date of publication 25 June 2013; date of current version 13 Aug. 2014. Recommended for acceptance by D. Epema.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below. Digital Object Identifier no. 10.1109/TPDS.2013.163

1.1 Churn Model

Heterogeneity of lifetimes is an important property of P2P networks [26], where some users consistently spend significant periods of time in the system and others very little. This observation prompts the question of whether *P2P systems can be modeled using a single homogeneous lifetime distribution (as done in prior work) without sacrificing model accuracy?* In addition to lifetimes, churn is characterized by the distribution of offline durations, which together with lifetimes define the *availability* of each user [4], [24], i.e., the average fraction of time a user is logged in. It is therefore important to understand how off-times contribute to the dynamics of the system and which peer characteristics affect local graph-theoretic properties of each user.

To answer these questions, we offer a novel churn model for the *heterogeneous* behavior of end-users that captures the various differences in their online/offline habits. We view each user as an alternating renewal process that is ON when the user is logged in and OFF otherwise, where online/offline durations of each user i are drawn from peer-specific distributions $F_i(x)$ and $G_i(x)$, respectively. Armed with this model, we obtain the lifetime distribution $F(x)$ of users who visit the system, lifetime distribution $J(x)$ of peers who are concurrently online, and residual lifetime distribution $H(x)$ of a randomly selected user in the graph.

Our results show that $F(x)$ is a weighted function of individual lifetime CDFs $F_i(x)$; however, the other two distributions are rather complex and currently intractable unless neighbor selection is performed *uniformly among the live participants*¹, in which case we show that $J(x)$ and $H(x)$ can be directly obtained from $F(x)$. This is an important conclusion that demonstrates that instead of measuring n individual lifetime distributions, where n is the total number of users participating in the system, one can measure lifetimes of joining peers to obtain $F(x)$, which is then *sufficient to entirely model the effect of churn on P2P graphs*.

1. This can be implemented by picking users from uniformly random subsets of cached nodes or using special random walks [30].

We also revisit the observation of [26] that the users already present in Gnutella and BitTorrent networks exhibit larger average lifetimes than those joining the system. We show that this effect is a consequence of $J(x)$ being the *spread* [28] of $F(x)$, which allows us to prove that random users currently in the system have stochastically larger lifetimes than random arriving users for any set of distributions $\{F_i(x)\}_{i=1}^n$ and $\{G_i(x)\}_{i=1}^n$. We additionally show that while $F(x)$ may be heavy-tailed in practice [6], [11], [18], it is possible that individual lifetime distributions $F_i(x)$ are all exponential, or contain a mixture of exponential and heavy-tailed distributions. Occurrence of this effect depends on random availability of each user and shows that individual peer habits cannot be inferred from their aggregate behavior $F(x)$.

1.2 Local Resilience Model

In the second half of the paper, we tackle the issue of node resilience to isolation in the presence of churn (i.e., loss of all neighbors within a user's lifetime). Prior work [15], [16] has shown that many P2P graphs stayed connected if and only if they did not develop isolated nodes during churn. Despite its importance, the approach taken in [15] modeled node isolation only considering *out-degree* neighbors and did not examine the increased resilience arising from *in-degree* edges.

We overcome this shortcoming by building a complete closed-form model characterizing the evolution of in-degree in unstructured systems. As in [15], we assume that each joining node selects k out-degree neighbors and then constantly replaces them upon detecting their failure. We first establish that the combined arrival process of initial and replacement edges to each user converges to Poisson as system size $n \rightarrow \infty$. We then obtain a model for the expected in-degree of a user at different age t as a function of $F(x)$, including cases with non-exponential peer lifetimes. This result shows that users who survive the initial period accumulate non-trivial in-degree and become more resilient to isolation over time. In contrast, [15] reached the conclusion that P2P users became progressively more susceptible to isolation as their age increased.

We finish the paper by merging the in/out-degree isolation models into a single approximation that clearly shows the contribution of in-degree to the resilience of the graph. Denoting by ϕ the isolation probability of a user and by ϕ_{out} the same metric only considering the out-degree [15], we show that for exponential $F(x)$ the ratio of the two isolation probabilities converges to the following as search delays become asymptotically small:

$$\frac{\phi}{\phi_{out}} \rightarrow \frac{1 - e^{-2k}}{2k}, \quad (1)$$

where k is the initial number of neighbors requested by each arriving user. This result illustrates that the amount of improvement from the in-degree is approximately a factor of $2k$ reduction in the isolation probability. We finish the paper with examples that demonstrate this effect and discussing its implications.

This paper is organized as follows. Section 2 overviews related work. We introduce our churn model in Section 3

and study its properties in Section 4. Section 5 models edge arrival to individual users and derives the expected in-degree of each user. Section 6 models joint in/out-degree isolation. We conclude the paper in Section 7.

2 RELATED WORK

One direction in the literature for modeling churn treats P2P systems as an $M/M/1$ queue (i.e., Poisson arrivals/departures and exponential lifetimes L). For example, [21] assumes a system in which neighbor replacement occurs in direct response to failures and is *instantaneous* (i.e., without any search delay), where the possibilities for replacement are limited to the nodes currently live in a centralized cache. The paper shows that under user churn the graph remains connected and exhibits a logarithmic diameter, both with high probability. Later models of churn [13], [19] assume a DHT-like system in which repair algorithms are run *independently* of user failures and at exponentially distributed intervals (i.e., as Poisson processes). This approach models the consistency-check algorithm in Chord that periodically verifies the successor list and corrects invalid pointers.

The benefit of $M/M/1$ models lies in their memoryless properties that dramatically simplify analysis (e.g., residual lifetimes R are exponential and all system dynamics are Markovian); however, measurement studies [6], [26], [27] show that real networks exhibit a Pareto distribution of L with extremely heavy tails ($\alpha \approx 1.1$). To address this limitation, [15] specifically targets non-exponential user lifetimes and additionally allows *non-zero* neighbor-replacement delays, albeit still under the homogeneous assumption on user lifetimes. In this model, each joining user randomly selects k out-neighbors from the graph and then monitors their online presence using keep-alive messages. Once the failure of an existing neighbor is detected, a uniformly random replacement is sought from among the currently live users in the system. Under these conditions, the paper shows that each user becomes isolated along the out-degree with probability no larger than $\phi_{out} = k\rho/(1 + \rho)^k$, where ρ is the ratio of the average lifetime to the average replacement delay, for all lifetime distributions with an exponential or heavier tail. This result is later generalized in [16] to show that the probability of non-partitioning in many P2P networks converges as $n \rightarrow \infty$ to that of avoiding isolation for each online user.

Our work encompasses all previous cases by allowing heterogeneous user behavior (as seen in real networks [26]), considering the in-degree process, and modeling resilience to disconnection using the joint in/out degree of each node, none of which has been done before.

3 CHURN MODEL

To understand the dynamics of churn and performance of P2P systems, we start by creating a model of user behavior and specifying assumptions on peer arrival, departure, and selection of neighbors. The focus of this section is to formalize recurring user participation in P2P systems in a simple model that takes into account heterogeneous browsing habits and explains the relationship between the various lifetime distributions observable in P2P networks.

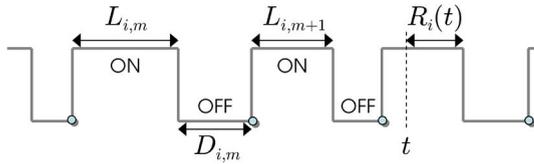


Fig. 1. Process $\{Z_i(t)\}$ depicting ON/OFF behavior of user i .

3.1 Basics

Consider a P2P system with n participating users, where each user i is either alive (i.e., present in the system) at time $t \geq 0$ or dead (i.e., logged off). This behavior can be modeled by a right-continuous ON/OFF process

$$Z_i(t) := \begin{cases} 1 & \text{user } i \text{ is alive at time } t, \\ 0 & \text{otherwise} \end{cases}, \quad 1 \leq i \leq n. \quad (2)$$

This framework is illustrated in Fig. 1, where parameter m stands for the cycle number and random variables $L_{i,m} > 0$, $D_{i,m} > 0$ are durations of user i 's ON (life) and OFF (death) periods, respectively. The figure also shows the *residual process* $R_i(t)$, which is the duration of user i 's remaining online presence from time t , conditioned on its being alive at t .

Assumption 1. Set $\{Z_i(t)\}_{i=1}^n$ consists of mutually independent, alternating renewal processes.

To elaborate, the renewal-process condition means that ON durations $\{L_{i,m}\}_{m=1}^{\infty}$ of user i are iid random variables with some cumulative distribution function (CDF) $F_i(x)$ and OFF durations $\{D_{i,m}\}_{m=1}^{\infty}$ are also iid with another CDF $G_i(x)$. Mutual independence between the users means that they do not synchronize their arrivals or departures and generally exhibit uncorrelated lifetime characteristics.

The main analytical limitation of staying with just Assumption 1 is that it allows user lifetimes to depend on n , which leads to difficulties in proving convergence of various metrics as $n \rightarrow \infty$. We thus offer an alternative method for building P2P graphs that overcomes this problem, but still allows for diversity of user behavior.

Assumption 2.

1. There exist $T \geq 1$ user types representing different churn profiles, which are modeled by a set \mathcal{F} of distinct pairs of non-lattice CDFs defining non-negative random variables

$$\mathcal{F} := \left\{ \left(F^{(1)}(x), G^{(1)}(x) \right), \dots, \left(F^{(T)}(x), G^{(T)}(x) \right) \right\}, \quad (3)$$

where $F^{(j)}(x) > 0$ and $G^{(j)}(x) > 0$ for all $x > 0$, all $j = 1, \dots, T$. Further, the corresponding means

$$l^{(j)} := \int_0^{\infty} (1 - F^{(j)}(x)) dx \quad (4)$$

$$d^{(j)} := \int_0^{\infty} (1 - G^{(j)}(x)) dx \quad (5)$$

satisfy $0 < l^{(j)}, d^{(j)} < \infty$ for all $j = 1, \dots, T$;

2. At time $t = 0$, the pair of ON/OFF-duration CDFs $(F_i(x), G_i(x))$ of each user $i = 1, \dots, n$ is independently drawn from set \mathcal{F} and remains fixed for all $t > 0$, where type- j pairs are selected with probability $p_j \geq 0$ and $\sum_{j=1}^T p_j = 1$;
3. Defining \mathcal{S} to be set of selections made by each user and conditioning on \mathcal{S} , Assumption 1 holds.

Assumption 2(a) uses T as the "diversity" factor of user lifetime/offtime behavior, where $T = 1$ reduces the system to a network of homogeneous users, while $T \geq 2$ allows heterogenous dynamics (e.g., desktop users with large lifetimes and wireless clients with frequent disconnections). This part of the assumption also mandates that all expected online/offline durations be both positive and finite. Part (b) allows for bias in the selection process and lets certain user types be more popular than others. Part (c) ensures that once users have chosen their types, the corresponding ON/OFF processes are independent.

3.2 Properties

For convenience of notation, define the mean lifetime of each user $l_i := E[L_{i,m}]$ and the mean offline duration $d_i := E[D_{i,m}]$, where the average is taken over all cycles $m = 1, 2, \dots$. We next explain the ON/OFF distributions commonly considered in this paper and obtain basic properties of the system.

The first lifetime CDF is exponential

$$F_i(x) = 1 - e^{-x/l_i} \quad (6)$$

and the second CDF is Pareto(α, β)

$$F_i(x) = 1 - \left(1 + \frac{x}{\beta_i} \right)^{-\alpha_i}, \quad \alpha_i > 1, \beta_i > 0, \quad (7)$$

with $l_i = \beta_i/(\alpha_i - 1)$. As we discover below, only the mean of offline distribution $G_i(x)$ affects the studied parameters of the system, which allows us to keep this CDF general. From Smith's theorem [22], the asymptotic *availability* of each user i , i.e., probability that it is in the system at an arbitrary instance t , is given by

$$a_i := \lim_{t \rightarrow \infty} P(Z_i(t) = 1) = \frac{l_i}{l_i + d_i}. \quad (8)$$

We are now ready to obtain the distribution of system size. Denote by $N(n, t) := \sum_{i=1}^n Z_i(t)$ the number of peers in the graph at time t and notice that it is a random process that fluctuates with time. Since P2P networks typically have a non-trivial operating history, analysis is often limited to cases when network age $t \rightarrow \infty$, which we call the *equilibrium state of the system*.

Replacing process $Z_i(t)$ by a random variable Z_i with the same equilibrium distribution, i.e., $P(Z_i = 1) = a_i$, the stationary system size is $N(n) := \sum_{i=1}^n Z_i$. Since Z_i 's are independent and identically distributed under Assumption 2, the central limit theorem shows that $N(n)$ becomes Gaussian as $n \rightarrow \infty$, whose mean and variance are determined solely by the availability set $\{a_1, \dots, a_n\}$.

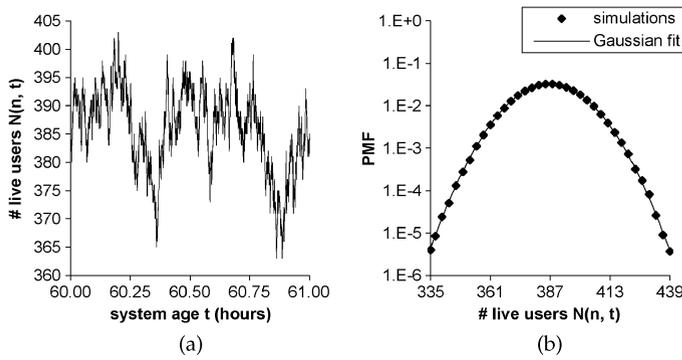


Fig. 2. Sample path and distribution of $N(n, t)$ in system \mathcal{H} with $n = 1000$ users and after 10^6 iterations. The Gaussian fit is from Theorem 1. (a) evolution. (b) PMF.

Theorem 1. Under Assumption 2 and $n \rightarrow \infty$

$$\frac{N(n) - \mu_n}{\sigma_n} \xrightarrow{D} \mathcal{N}(0, 1), \quad (9)$$

where $\mu_n := E[N(n)] = \sum_{i=1}^n a_i$, $\sigma_n^2 := \text{Var}[N(n)] = \sum_{i=1}^n a_i(1 - a_i)$, and $\mathcal{N}(0, 1)$ is the standard normal random variable.

We next show simulations explaining this result and confirming its accuracy in systems with finite age and size. We generate a network of n users whose arrival/departure follows the introduced churn model. The system evolves for at least 50 virtual hours before being examined. We start by generating $\mathcal{T} = 1000$ pairs of means $l^{(j)}$ and $d^{(j)}$, drawn randomly from two Pareto distributions with $\alpha = 3$. For mean ON durations, we use $\beta = 1$ and obtain $E[l^{(j)}] = 1/2$ hour; for mean OFF durations, we use $\beta = 2$ and get $E[d^{(j)}] = 1$ hour.

We study three cases throughout the paper: 1) *heavy-tailed* system \mathcal{H} with $F^{(j)}(x) \sim \text{Pareto}(3, 2l^{(j)})$ and $G^{(j)}(x) \sim \text{Pareto}(3, 2d^{(j)})$; 2) *very heavy-tailed* system \mathcal{VH} with $F^{(j)}(x) \sim \text{Pareto}(1.5, l^{(j)}/2)$ and $G^{(j)}(x) \sim \text{Pareto}(1.5, d^{(j)}/2)$; and 3) *exponential* system \mathcal{E} with $F^{(j)}(x) \sim \exp(1/l^{(j)})$ and $G^{(j)}(x) \sim \text{Pareto}(3, 2d^{(j)})$, where notation $\text{Pareto}(\alpha_i, \beta_i)$ refers to (7). The actual pairs $(F_i(x), G_i(x))$ are selected uniformly randomly from \mathcal{F} , i.e., $p_j = 1/\mathcal{T}$.

Fig. 2(a) shows a single sample-path of system size $N(n, t)$ as a function time t . Part (b) of the figure shows the probability mass function (PMF) of $N(n, t)$ at $t \gg 0$ and a Gaussian fit from Theorem 1, confirming accuracy of the model and validity of approximating $N(n, t)$ with $\mathcal{N}(\mu_n, \sigma_n^2)$.

3.3 Aggregate Lifetimes

Prior measurement studies [26], [27] sampled lifetimes of all users seen in the system over long periods of time to characterize the dynamics of P2P networks. We are now interested in what metric they estimated and how it can be expressed under our churn model.

For each instance of user i present in the system during interval $[0, t]$, suppose we place its ON duration $L_{i,m}$ into set $C_i(t)$ and define $C(t) = \cup_{i=1}^n C_i(t)$, which is a collection of lifetimes across all users seen in $[0, t]$. Then let $F(n, t, x)$ be the CDF of values collected in set $C(t)$ (i.e., the probability that the obtained lifetimes are less than or equal to x). Finally, given n participating users, define $F(n, x) := \lim_{t \rightarrow \infty} F(n, t, x)$ to be the *aggregate lifetime distribution* of the system and $l(n)$ to be its mean in equilibrium.

Next, suppose the average join rate of user i is

$$\lambda_i := \frac{1}{l_i + d_i}, \quad (10)$$

and the fraction of joins that i is responsible for is

$$b_i := \frac{\lambda_i}{\sum_{j=1}^n \lambda_j}. \quad (11)$$

Then, our next result shows that the aggregate lifetime CDF $F(n, x)$ is a weighted average of individual distributions $F_i(x)$, where the weights are biased toward those peers who frequently join and leave the system. This was intuitively expected as their sessions constitute higher fractions of overall peer arrival into the system.

Theorem 2. With Assumption 1, the aggregate lifetime CDF $F(n, x)$ and its mean $l(n)$ are given by

$$F(n, x) = \sum_{i=1}^n b_i F_i(x), \quad l(n) = \sum_{i=1}^n b_i l_i. \quad (12)$$

Proof. See Section 8.1 in Supplemental Material which is available in the Computer Society Digital Library at <http://doi.ieeecomputersociety.org/10.1109/TPDS.2013.163>. \square

To perform a self-check, define $\Lambda(n) = \sum_{j=1}^n \lambda_j$ to be the combined join rate across all users and observe that the expected duration for which peers stay in the system

$$l(n) = \sum_{i=1}^n b_i l_i = \frac{\sum_i \lambda_i l_i}{\sum_j \lambda_j} = \frac{\sum_i a_i}{\Lambda(n)} = \frac{E[N(n)]}{\Lambda(n)} \quad (13)$$

is the mean system size divided by the aggregate user arrival rate, which is consistent with Little's Law.

The next result is stated without proof as it is similar to that of Theorem 2.

Corollary 1. With Assumption 1, the CDF of aggregate offline durations is $G(n, x) := \sum_{i=1}^n b_i G_i(x)$ and its mean is $d(n) := \sum_{i=1}^n b_i d_i$.

Note that Theorem 2 and Corollary 1 hold under the more general Assumption 1, *but only for finiten*. To guarantee that the sums in (12) converge as $n \rightarrow \infty$, one requires Assumption 2. We examine this issue next.

3.4 Convergence of Lifetime Distribution

In order to model the arrival process of in-degree edges later in the paper, we need to know that $F(n, x)$ exists as $n \rightarrow \infty$. To simplify discussion, define the availability and join rate of user type j as

$$a^{(j)} = \frac{l^{(j)}}{l^{(j)} + d^{(j)}} \quad \lambda^{(j)} := \frac{1}{l^{(j)} + d^{(j)}}. \quad (14)$$

Then, the expected availability and arrival rate across all users can be written as

$$a := \sum_{j=1}^{\mathcal{T}} p_j a^{(j)} \quad \lambda := \sum_{j=1}^{\mathcal{T}} p_j \lambda^{(j)}. \quad (15)$$

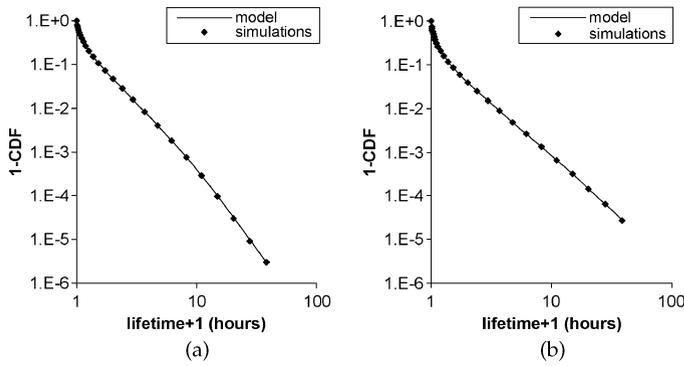


Fig. 3. Comparison of simulation results of $F(n, x)$ to model (12) in a graph with $n = 1000$ nodes. (a) system \mathcal{E} . (b) system \mathcal{H} .

Our next result shows that $F(n, x)$ and its mean $l(n)$, both shown in (12), exhibit convergence as $n \rightarrow \infty$.

Theorem 3. *Under Assumption 2, the following sequences converge almost surely (a.s.) as $n \rightarrow \infty$:*

$$F(n, x) \xrightarrow{a.s.} F(x) := \frac{1}{\lambda} \sum_{j=1}^T p_j \lambda^{(j)} F^{(j)}(x) \quad (16)$$

$$l(n) \xrightarrow{a.s.} l := \frac{a}{\lambda}. \quad (17)$$

Proof. See Section 8.2 in Supplemental Material available online. \square

3.5 Examples

We verify (12) in simulations and discuss several implications of this result. Verification of (16), (17) produces similar results, which we omit for brevity. Two examples are presented in Fig. 3 for systems driven by 1) exponential and 2) heavy-tailed lifetimes. Note that the figures plot the tail $1 - F(n, x)$ on a log-log scale and use lifetime shifted by 1 hour on the x -axis to make the Pareto distribution in (7) appear as a straight line. First, observe that the model is consistent with simulation results in both cases. Second, notice in Fig. 3(a) that system \mathcal{E} produces a heavy-tailed aggregate distribution $F(n, x)$ even though all individual $F_i(x)$ are exponential. This can be explained as follows.

It is known [9] that for a hyper-exponential distribution in the form of (12) and any desired distribution $W(x)$ with a monotonic probability density function (PDF), there exists a set of weights $\{b_1, \dots, b_n\}$ such that (12) converges to $W(x)$ as $n \rightarrow \infty$. Given numerous possibilities for the arrival-rate set $\{\lambda_1, \dots, \lambda_n\}$ in practice, it is possible that $F(n, x)$ may be Pareto, Weibull, or otherwise heavy-tailed, while all individual lifetimes $F_i(x)$ are exponential. It is therefore important to note that Pareto $F(n, x)$ measured experimentally [6], [24] does not necessarily reveal the true nature of individual user behavior.

While our current conclusion shows that one cannot characterize the lifetimes or availability of individual peers by observing their aggregate behavior, the next question we seek to answer is *whether the aggregate behavior $F(n, x)$ can be used to characterize the parameters of a single user selected from the system randomly?*

4 CHARACTERISTICS OF SELECTED USERS

Suppose a node v picks a random live user i as a potential neighbor. Our primary goal in this section is to understand the properties of i in terms of two metrics: its remaining online duration and current session length.

4.1 Residual Lifetimes

Recall from Fig. 1 that $R_i(t)$ is the remaining life of a given user i at time t , i.e., the remainder of the current ON cycle. Variable $R_i(t)$ is important since it determines how long i will remain online after it has been selected as a neighbor. The equilibrium residual lifetime distribution

$$H_i(x) := \lim_{t \rightarrow \infty} P(R_i(t) \leq x | Z_i(t) = 1)$$

can be written in terms of $F_i(x)$ [28]

$$H_i(x) = \frac{1}{l_i} \int_0^x (1 - F_i(u)) du, \quad x \geq 0. \quad (18)$$

Defining $R(n, t)$ to be the residual lifetime of the user randomly selected from among $N(n, t) \geq 1$ users that are alive, suppose $H(n, x)$ is the equilibrium distribution of $R(n, t)$ conditioned on $N(n, t) \geq 1$

$$H(n, x) := \lim_{t \rightarrow \infty} P(R(n, t) \leq x | N(n, t) \geq 1). \quad (19)$$

We next prove that $H(n, x)$ converges as $n \rightarrow \infty$ and obtain its CDF in closed-form.

Theorem 4. *Under uniform selection and Assumption 2, the residual lifetime CDF $H(n, x)$ converges a.s. to the following as $n \rightarrow \infty$:*

$$H(n, x) \xrightarrow{a.s.} H(x) := \frac{1}{l} \int_0^x (1 - F(u)) du, \quad (20)$$

where $F(x)$ and l are given in (16), (17).

Proof. See Section 8.3 in Supplemental Material available online. \square

4.2 Examples

Leveraging Theorem 4 for finite n , one can apply either of the following useful approximations:

$$H(n, x) \approx \frac{1}{l(n)} \int_0^x (1 - F(n, u)) du \quad (21)$$

$$\approx \frac{\sum_{i=1}^n a_i H_i(x)}{\sum_{i=1}^n a_i}, \quad (22)$$

where the first one relies on the knowledge of the aggregate lifetime CDF $F(n, x)$, while the second requires all individual distributions $\{F_i(x)\}_{i=1}^n$.

We next examine (21) in simulations with networks of small size n (in practice, $n \geq 50$ is often sufficient to keep the model very accurate). As shown in Fig. 4 for the exponential and Pareto cases, simulations of $H(n, x)$ match the model very well and demonstrate that \mathcal{E} may again

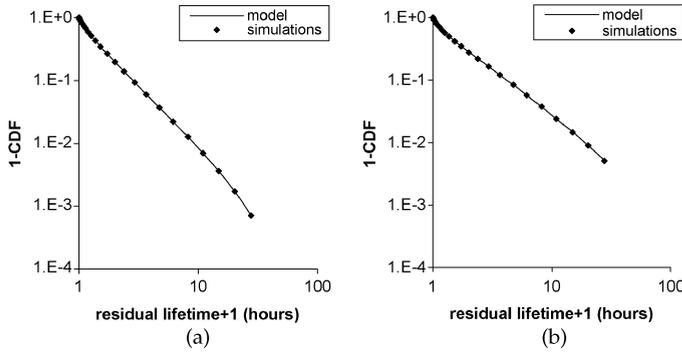


Fig. 4. Comparison of simulation results of $H(n, x)$ to model (21) in a graph with $n = 1000$ nodes. System age 500 h and 10^5 iterations. (a) system \mathcal{E} . (b) system \mathcal{H} .

produce non-exponential dynamics by mixing enough exponential distributions.

Note that (20) is of fundamental importance since it shows that we need to measure only the *aggregate* lifetime distribution $F(x)$ rather than each $F_i(x)$ in order to obtain the residual lifetime distribution of a uniformly selected neighbor. Assuming from measurement studies [6], [11], [18] a Pareto $F(x)$, (20) reduces to

$$H(x) = 1 - (1 + x/\beta)^{-(\alpha-1)}. \quad (23)$$

Comparing (23) to (7), we see that residuals are stochastically larger than user lifetimes, i.e., with shape parameter $\alpha - 1$ instead of α , which implies that a uniformly selected user is more reliable than new arrivals in terms of failure. For other neighbor selection strategies, the distribution of residual lifetimes may be completely different from (20) and should be analyzed using more advanced techniques.

4.3 Lifetime of Users in the System

Denote by $J(n, x)$ the equilibrium lifetime distribution of users *currently* in the system conditioned on $N(n, t) \geq 1$. As observed in [26], distribution $J(n, x)$ is clearly different from $F(n, x)$; however, no closed-form analysis has been offered to explain this. The heuristic rationale behind this difference is that lifetimes of live peers are biased towards larger values, which is commonly known as the *inspection paradox* [28]. Since $J(n, x)$ is the lifetime distribution of a *uniformly randomly selected user from the set of live peers*, we reach the next result by applying Theorem 4.

Corollary 2. *Given Assumption 2, the lifetime distribution $J(n, x)$ of live users converges a.s. as $n \rightarrow \infty$*

$$J(n, x) \xrightarrow{a.s.} J(x) := \frac{1}{l} \left(xF(x) - \int_0^x F(u)du \right). \quad (24)$$

We omit simulations showing the accuracy of (24) due to limited space. Note that exponential lifetimes $F(x)$ imply that $J(x)$ is the Erlang(2) distribution with mean $2E[L]$. For Pareto $F(x)$, spread $J(x)$ has no closed-form expression, but is clearly more heavy-tailed than $F(x)$. The next result summarizes these observations, as well as those of [26], in more formal terms using the fact that spread is the convolution of two distributions $H(x)$.

Corollary 3. *With Assumption 2, spread distribution $J(x)$ is stochastically larger than $F(x)$ and the mean lifetime of a user currently alive in the system is double the mean residual lifetime of a uniformly selected user.*

So far, our results have demonstrated that heterogeneous users employing uniform selection allow both metrics $H(x)$ and $J(x)$ to be reduced to the aggregate behavior $F(x)$ of joining users as long as $n \gg 1$. The rest of the paper shows that $F(x)$ in such systems can be additionally used to obtain the distribution of in-degree as a function of users' age and thus completely characterize local resilience of unstructured P2P networks.

5 IN-DEGREE ANALYSIS

We begin this section by modeling the edge-generation process since this determines both the rate at which each user accumulates incoming neighbors and the stationary in-degree distribution.

5.1 Edge Creation

Assume that each joining node w finds k random initial neighbors and then continuously replaces them as they fail [15]. Let $L \sim F(x)$ be the random lifetime of a user that participates in the system, where $n \rightarrow \infty$ is assumed throughout the section. As the delay S to find replacements is usually negligible compared to ON/OFF durations, it does not have a significant impact on our results. We thus do not explicitly model its effect here.

Focusing on a single out-link of a joining user w , define t_j to be the instance when this link gets its j -th out-neighbor, where $t_0 = 0$ and $t_j = t_{j-1} + R_j$ for $j \geq 1$, where $R_j \sim H(x)$ is the residual of the j -th out-neighbor. Then, suppose that $\{U(t)\}_{t \geq 0}$ is a renewal process whose inter-renewal delays are distributed according to $H(x)$

$$U(t) := \sum_{j=0}^{\infty} \mathbf{1}_{t_j \in [0, t]}. \quad (25)$$

Note that $U(t)$ counts the number of replacements that occur in $[0, t]$, where a) the first renewal always occurs at 0 (i.e., $U(0) = 1$) and b) $U(t) = 0$ for $t < 0$. Then, the expected number of outbound connections generated along a single out-link of w in the interval $[0, t]$ is the *renewal function* $u(t) := E[U(t)]$, which can be expressed as [22]

$$u(t) = \begin{cases} 1 + \sum_{r=1}^{\infty} H^{*r}(t) & t \geq 0 \\ 0 & t < 0, \end{cases} \quad (26)$$

where $H^{*r}(t)$ is the r -fold convolution of $H(t)$. Now, define the expected number of replacements made along all k out-links of a user during its lifetime $L \sim F(x)$ as

$$\theta := kE[u(L) - 1]. \quad (27)$$

Under these definitions, each user arrival inserts on average k *initial* edges into the system and θ *replacement* edges, for a total of $k + \theta$ edges per L time units. In the following, we first examine the *user* arrival process in a P2P system driven by our churn model and then study the behavior of *edge* arrival to each user.

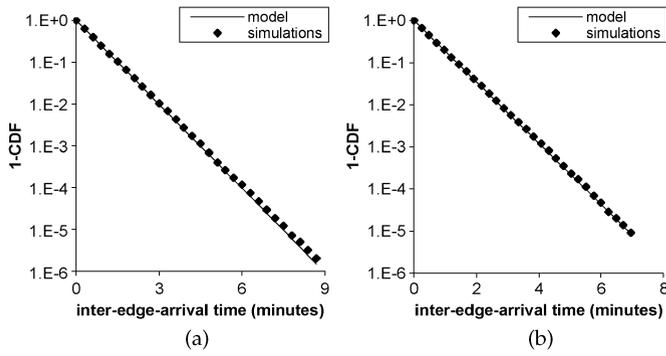


Fig. 5. CCDF of edge inter-arrival delays and exponential model with rate γ in (30) for $n = 1000$, $k = 10$, and $\theta = 10$ using 10^9 iterations. (a) system \mathcal{H} with $n = 1000$. (b) system \mathcal{VH} with $n = 5000$.

5.2 Node Arrival

For each ON/OFF renewal process $\{Z_i(t)\}$, denote by $M_i(t)$ the number of arrivals of user i in the interval $[0, t]$. Note that $\{M_i(t)\}$ is a stationary point process [22] with inter-arrival durations $L_i + D_i$. As before, suppose $\Lambda(n)$ is the combined arrival rate of all peers

$$\Lambda(n) = \sum_{i=1}^n \lambda_i, \quad (28)$$

where λ_i is defined in (10). Then the *superposed* user arrival process into the whole system is $M(n, t) = \sum_{i=1}^n M_i(t)$ with rate $\Lambda(n)$. To understand our next result, define a scaled user ON/OFF process

$$Z'_i(t) = Z_i(t/\Lambda(n)) \quad (29)$$

and assume $\{M'_i(t)\}$ is the point arrival process driven by $Z'_i(t)$, i.e., number of arrivals of user i under $Z'_i(t)$ in the interval $[0, t]$.

Theorem 5. *Given Assumption 2, scaled superposition process $M'(n, t) = \sum_{i=1}^n M'_i(t)$ converges in distribution as $n \rightarrow \infty$ to a Poisson process with unit rate.*

Proof. See Section 8.4 in Supplemental Material available online. \square

Performing the inverse transformation [22], observe that $M(n, t) = M'(n, \Lambda(n)t)$ has properties of a Poisson process for $n \rightarrow \infty$ with rate $\Lambda(n)$, but *only* in sufficiently small windows $[t, t + \Delta t]$. This can be explained by the fact that each renewal process $Z_i(t)$ must appear *sparse* compared to the size of the observation window $[t, t + \Delta t]$ for $M(n, t)$ to be Poisson. At larger timescales [20], the inter-arrival delays of $M(n, t)$ converge to exponential, which is consistent with the measurements in [26], but *the process does not become memoryless even when $n \rightarrow \infty$.*

Since the properties of the scaled *departure* process are identical to those of $M'(n, t)$ in Theorem 5, we omit their formal statement. While user arrivals/departures may be Poisson only locally (i.e., on small timescales), analysis in the remainder of the section shows that the in-degree edge-arrival process does in fact converge to Poisson as $n \rightarrow \infty$. This holds because edge delivery between a pair of live users

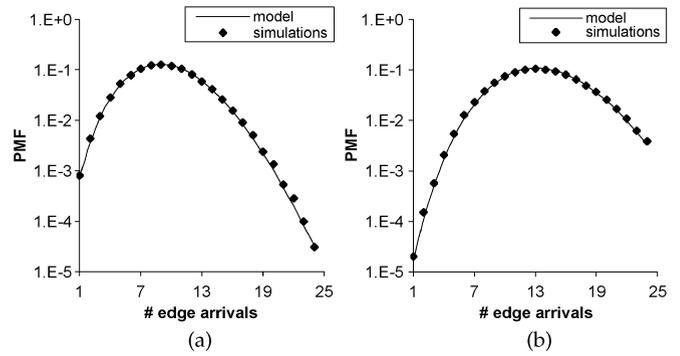


Fig. 6. Distribution of the number of edge arrivals to a node in the interval $[t, t + \Delta t]$ in a system with $n = 1000$ users, $k = 10$, and $\theta = 10$ ($t = 500$ hours). (a) system \mathcal{H} with $\Delta t = 6$ min. (b) system \mathcal{VH} with $\Delta t = 9$ min.

occurs with increasingly large delays (i.e., gets sparser) as $n \rightarrow \infty$, which is the condition that $M(n, t)$ was lacking.

5.3 Edge Arrival

Define $L \sim F(x)$ to be the lifetime of a user and note that its expectation is $E[L] = l$ in (17).

Theorem 6. *Given uniform selection and Assumption 2, the combined arrival process of new and replacement in-degree edges to a random user converges in distribution as $n \rightarrow \infty$ to a Poisson process with constant rate*

$$\gamma := \frac{k + \theta}{E[L]}. \quad (30)$$

Proof. See Section 8.5 in Supplemental Material available online. \square

Note that the Poisson result in Theorem 6 is not an *assumption* of the paper as in prior work [13], [19], [21], but rather a *consequence* of the churn model introduced earlier. To examine how well this works for small n , we show in Fig. 5 the tail distribution of in-edge inter-arrival delays obtained in simulations. Notice that both tails are exponential with the rate given by (30). Additionally, Fig. 6 shows that the number of edge arrivals to a node v in an interval of size Δt closely follows the Poisson distribution with rate $\gamma \Delta t$.

5.4 Expected In-Degree

We are now ready to derive the expected in-degree of each live user as a function of its age. In a stationary system, define $X_n(t)$ to be the in-degree of a random online user v at age $t \geq 0$. Applying Theorems 4 and 6, we show transient properties of $X_n(t)$ next.

Theorem 7. *Under uniform selection and Assumption 2, the expected in-degree of a live user v at fixed age $t \geq 0$ converges as $n \rightarrow \infty$ to a monotonically increasing function of age*

$$E[X_n(t)] \rightarrow k \int_0^{\infty} (u(x) - u(x - t)) dH(x), \quad (31)$$

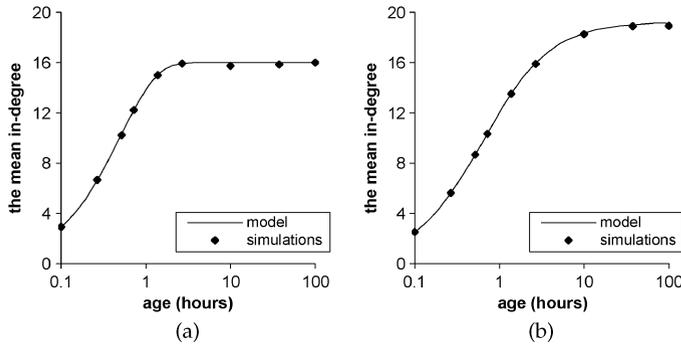


Fig. 7. Comparison of the model for $E[X_n(t)]$ to simulation results for $n = 2000$, $E[L] = 0.5$ hours, and $k = 8$ after 10^6 iterations. (a) exponential (33). (b) Pareto (31) with $\alpha = 3$.

where $H(x)$ is the residual CDF in (20) and $u(t)$ is the renewal function defined in (26).

Proof. See Section 8.6 in Supplemental Material available online. \square

To better interpret the mean in-degree of user v at fixed age t , we transform (31) into

$$\lim_{n \rightarrow \infty} E[X_n(t)] = kE[u(A) - u(A - t)], \quad (32)$$

where variable $A \sim H(x)$ denotes the age of a random user i that is ON at time t , $A - t$ is the duration of time for which i had been alive prior to v 's appearance online, and $u(x)$ represents the expected number of selections made for each out-link in an interval of length x (recall that $U(x) = 0$ for $x < 0$).

It is then not hard to see that the right-hand side of (32) is the expected number of connections that a peer generates into the system while being alive concurrently with v . As shown in Theorem 7, this metric equals v 's mean in-degree, which confirms the symmetry of outgoing/incoming edges. Our last note is that as $t \rightarrow \infty$, (32) tends to $kE[u(A)]$, which provides a simple upper-bound at which the in-degree of each user saturates.

5.5 Examples

We next show that (31) can be expressed in simple closed-form for exponential lifetimes.

Theorem 8. Under uniform selection, Assumption 2, and exponential lifetimes L , in-degree $X_n(t)$ converges to a Poisson random variable with rate

$$\nu(t) = 2k(1 - e^{-t/E[L]}). \quad (33)$$

Furthermore, the number of replacement edges per user lifetime is $\theta = k$, which also equals the expected in-degree at failure $E[\nu(L)]$.

Proof. See Section 8.7 in Supplemental Material available online. \square

Simulation results in Fig. 7 demonstrate that the models are very accurate and indeed saturate at predicted values

$2k$ and $kE[u(A)]$ as age $t \rightarrow \infty$. It is also interesting to observe in the figure that the Pareto curve increases slower, but saturates at larger values, which suggests more in-degree support for users with larger lifetimes. For the examples in the figure, if a node survives for more than 1 hour in the system, it develops an average of 12-15 in-degree neighbors (depending on the distribution of L) and is unlikely to be isolated from the graph from that point on.

Interestingly, analysis above shows that P2P implementations should cap user in-degree at values *no smaller* than $\nu(\infty)$ to avoid rejecting incoming connections. The corresponding upper bound in Gnutella (i.e., 30 in-degree neighbors) satisfies this condition in the two examples shown above.

6 JOINT IN/OUT-DEGREE MODEL

Analytical results in the previous section show that the early stage in a node's life in the network is risky from the isolation point of view as it must rely solely on its out-degree neighbors. However, once a node survives this early stage, it increases its resilience to isolation through constantly arriving in-degree edges.

In this section, we combine the in-degree and out-degree models to derive the *joint* isolation probability. We drop subscript n and implicitly use uniform selection, Assumption 2, and $n \rightarrow \infty$ in all results that follow.

6.1 Preliminaries

Denote by $X^*(t)$ the out-degree of a node v at given age t and define it to be *isolated* when its in-degree $X(t)$ and out-degree $X^*(t)$ are simultaneously zero. Define *time to isolation* T to be the first-hitting time of both processes to state 0

$$T = \inf\{t > 0 : X^*(t) = X(t) = 0 | X^*(0) = k, X(0) = 0\}.$$

Then the probability of node isolation is $\phi = P(T < L)$, where $L \sim F(x)$ is the random lifetime of node v . To prevent the snowball effect, nodes do not replace their *incoming* edges upon failure, which means that the in/out-degree processes are independent of each other.

We use two approaches to derive ϕ assuming exponential lifetimes and search delays. The first approach provides an exact model using matrix algebra, while the second one shows an asymptotically accurate approximation that is available in simple closed-form. Since the exponential case serves as a tight upper bound [15] on the isolation probability for heavy-tailed L , we do not explicitly deal with those cases here.

6.2 Exponential Lifetimes (Exact Model)

Let pair $(X^*(t), X(t))$ be the joint process of out-degree and in-degree at age t and (i, j) denote any admissible state of the joint process for $0 \leq i \leq k$ and $0 \leq j < n$. Recall that edge arrival to a live node occurs according to a Poisson process with rate (30), which for exponential lifetimes becomes $\gamma = 2k/E[L]$. Since current in/out-degree neighbors of v all fail at the same rate $\eta = 1/E[L]$ due to the memoryless property of exponentials, the next result follows immediately.

TABLE 1
Exact Model (36) and Simulations

$E[S]$ min	$k = 6$		$k = 8$	
	Simulations	Model (36)	Simulations	Model (36)
6	3.63×10^{-6}	3.61×10^{-6}	2.80×10^{-8}	2.87×10^{-8}
18	3.15×10^{-5}	3.17×10^{-5}	5.91×10^{-7}	5.98×10^{-7}
30	6.04×10^{-5}	6.08×10^{-5}	1.48×10^{-6}	1.46×10^{-6}
42	8.38×10^{-5}	8.37×10^{-5}	2.30×10^{-6}	2.27×10^{-6}
60	1.06×10^{-4}	1.09×10^{-4}	3.27×10^{-6}	3.28×10^{-6}

Theorem 9. Given lifetimes $L \sim \exp(\eta)$ and search times $S \sim \exp(\zeta)$, the joint process $\{(X^*(t), X(t))\}$ is a homogeneous continuous-time Markov chain with a transition rate matrix $Q = (q_{uu'})$

$$q_{uu'} = \begin{cases} i\eta & (i, j) \rightarrow (i-1, j) \\ (k-i)\zeta & (i, j) \rightarrow (i+1, j), \text{ for } i < k \\ j\eta & (i, j) \rightarrow (i, j-1) \\ 2k\eta & (i, j) \rightarrow (i, j+1) \\ -\Lambda_{ij} & (i, j) \rightarrow (i, j) \\ 0 & \text{otherwise,} \end{cases} \quad (34)$$

where u and u' represent any suitable states of the joint chain satisfying transition requirements on the right side of (34) and $\Lambda_{ij} = i\eta + (k-i)\zeta + j\eta + 2k\eta$.

It is convenient to treat $\{(X^*(t), X(t))\}$ as an absorbing Markov chain in order to derive the PDF of the first-hitting time T on state $(0,0)$. Assuming $(0,0)$ is an absorbing state, we can write Q in canonical form as

$$Q = \begin{pmatrix} 0 & 0 \\ \mathbf{r} & Q_0 \end{pmatrix}, \quad (35)$$

where Q_0 is the rate matrix obtained by removing the rows and columns corresponding to state $(0,0)$ from Q and \mathbf{r} is a column of transition rates into state $(0,0)$.

Applying [29, Theorem 2], we obtain the next result.

Theorem 10. For exponential lifetimes $L \sim \exp(\eta)$ and exponential search delays $S \sim \exp(\zeta)$, the probability of node isolation is given by

$$\phi = \mathbf{y}(0)VBV^{-1}\mathbf{r}, \quad (36)$$

where V and B are matrices derived in [29]. Vector $\mathbf{y}(0)$ contains initial-state probabilities also derived in [29].

We verify (36) in simulations shown in Table 1 using $n = 2000$ and $E[L] = 0.5$ hours. The table shows that our results are indeed very accurate; however, while (36) provides values ϕ that are definitely smaller than isolation probability ϕ_{out} of the out-degree model [15], it is still unclear what impact in-degree has on ϕ and how large the improvement ratio ϕ_{out}/ϕ is. We study these issues below.

6.3 Isolation with Increased Age

To better understand the impact of in-degree on ϕ , let us define the first hitting time T_{out} on state 0 of the out-degree process $\{X^*(t)\}$, i.e., $T_{out} = \inf\{t > 0 : X^*(t) = 0 | X^*(0) = k\}$. Analysis in [16] shows that $\{X^*(t)\}$ is a birth-death Markov chain and derives its CDF function $P(T_{out} < t)$ in matrix form. We use this result and the CDF of T derived in Theorem 10 to compare the distribution of isolation times in the joint in/out-degree model with that studied in [16].

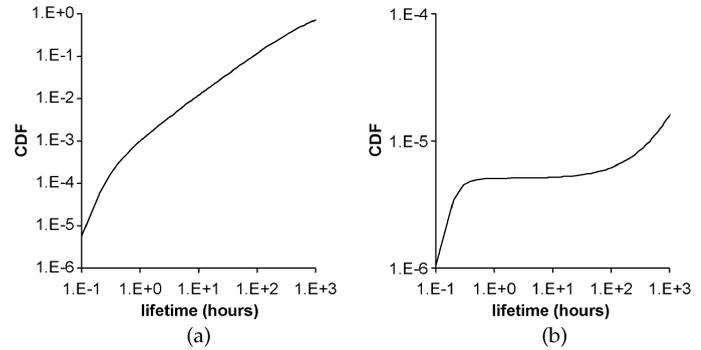


Fig. 8. CDF of T_{out} and T for exponential lifetimes with $E[L] = 0.5$ hours, exponential search delays with $E[S] = 0.1$ hours, and $k = 6$. (a) T_{out} . (b) T .

We plot the exact distributions of both T_{out} and T as functions of user age in Fig. 8. Notice in the figure that $P(T_{out} < t)$ increases almost linearly in time t indicating that users with large lifetimes have proportionally higher probabilities of isolation. In contrast, the CDF curve $P(T < t)$ stays almost flat for $t \in [0.5, 200]$ hours, showing that users with lifetimes in this range exhibit significantly lower isolation probabilities. In fact, once the initial 1/2-hour period is over, isolation probability is orders of magnitude smaller than in the initial phase.

As user age increases above 200 hours, the CDF of T slowly increases in time since in-degree $X(t)$ becomes saturated and can no longer keep up with the increased possibility of neighbor failure.

6.4 Exponential Lifetimes (Asymptotic Model)

Although (36) provides exact results for ϕ , it relies on numerical matrix algebra. Our next task is to obtain a simple closed-form expression for ϕ when the mean search delay $E[S]$ becomes negligible.

Theorem 11. For $L \sim \exp(\eta)$ and $S \sim \exp(\zeta)$, the ration of isolation probabilities converges as

$$\lim_{E[S] \rightarrow 0} \frac{\phi}{\phi_{out}} = \frac{1 - e^{-2k}}{2k}, \quad (37)$$

where $\phi_{out} = \rho k / (1 + \rho)^k$ and $\rho = \zeta / \eta = E[L] / E[S]$.

Proof. See Section 8.8 in Supplemental Material available online. \square

It can be seen from (37) that by considering both in/out-degree, the probability of node isolation is reduced by a factor of approximately $2k$. The reason for this relatively small improvement is that only a fraction of users benefit from the in-degree in their isolation resilience since the majority of users depart before they are able to accumulate a substantial number of in-degree neighbors. Nevertheless, analysis of this section has important consequences as it shows that *the most reliable users of the system (i.e., those with large lifetimes) extract significant benefits from the in-degree process and are thus allowed to continue providing services to others with a much higher probability than possible with just the out-degree.*

TABLE 2
Convergence of (37) to (36)

$E[S]$	Exact model (36)	Approx. model (37)	Relative error
36 sec	8.721×10^{-10}	1.421×10^{-9}	62.91%
3.6 sec	1.498×10^{-14}	1.581×10^{-14}	5.57%
360 ms	1.589×10^{-19}	1.598×10^{-19}	0.55%
36 ms	1.600×10^{-24}	1.600×10^{-24}	0

To complete this section, Table 2 shows the relative approximation error of (37) for $E[L] = 0.5$ hours and $k = 6$. For large S , our numerical results from (36) suggest that (37) provides an upper bound on the isolation probability, where ϕ_{out}/ϕ is several times larger than the $2k$ suggested by (37). However, as $E[S] \rightarrow 0$, the approximate model indeed converges to the exact one as seen in the table.

7 CONCLUSION

This paper introduced a heterogeneous model of user churn and developed numerous closed-form results describing the behavior of the system. This included three types of lifetime distributions, evolution of network size, transient in-degree, and isolation probability under the joint in/out degree model.

Future work involves modeling of non-uniform neighbor selection, study of isolation events in real P2P systems, and analysis of asymptotically small networks.

ACKNOWLEDGMENT

Parts of this paper appeared in ACM Hotmetrics, 2010.

REFERENCES

- [1] R. Albert, H. Jeong, and A.L. Barabási, "Error and Attack Tolerance of Complex Networks," *Nature*, vol. 406, no. 6794, pp. 378-382, July 2000.
- [2] D. Aldous and M. Brown, "Inequalities for Rare Events in Time-Reversible Markov Chains I," in *Stochastic Inequalities*, vol. 22, M. Shaked and Y.L. Tong, Eds. Hayward, CA, USA: Inst. Math. Stat., 1992, pp. 1-16.
- [3] J. Aspnes, Z. Diamadi, and G. Shah, "Fault-Tolerant Routing in Peer-to-Peer Systems," in *Proc. ACM PODC*, July 2002, pp. 223-232.
- [4] R. Bhagwan, S. Savage, and G.M. Voelker, "Understanding Availability," in *Proc. IPTPS*, Feb. 2003, pp. 256-267.
- [5] A.A. Borovkov, *Probability Theory*. New York, NY, USA: Gordon and Breach, 1998.
- [6] F.E. Bustamante and Y. Qiao, "Friendships that Last: Peer Lifespan and its Role in P2P Protocols," in *Web Content Caching and Distribution*. Norwell, MA, USA: Kluwer, Sept. 2003.
- [7] E. Çinlar, "Superposition of Point Processes," in *Stochastic Point Processes: Statistical Analysis, Theory, Applications*. New York, NY, USA: Wiley, 1972, pp. 549-606.
- [8] B.-G. Chun, B. Zhao, and J. Kubiatowicz, "Impact of Neighbor Selection on Performance and Resilience of Structured P2P Networks," in *Proc. IPTPS*, Feb. 2005, pp. 264-274.
- [9] A. Feldmann and W. Whitt, "Fitting Mixtures of Exponentials to Long-Tailed Distributions to Analyze Network Performance Models," *Perform. Eval.*, vol. 31, no. 3/4, pp. 245-279, Jan. 1998.
- [10] K. Gummadi, R. Gummadi, S. Gribble, S. Ratnasamy, S. Shenker, and I. Stoica, "The Impact of DHT Routing Geometry on Resilience and Proximity," in *Proc. ACM SIGCOMM*, Aug. 2003, pp. 381-394.
- [11] K.P. Gummadi, R.J. Dunn, S. Saroiu, S.D. Gribble, H.M. Levy, and J. Zahorjan, "Measurement, Modeling, Analysis of a Peer-to-Peer File-Sharing Workload," in *Proc. ACM SOSP*, Oct. 2003, pp. 314-329.

- [12] M.F. Kaashoek and D. Karger, "Koorde: A Simple Degree-Optimal Distributed Hash Table," in *Proc. IPTPS*, Feb. 2003, pp. 98-107.
- [13] S. Krishnamurthy, S. El-Ansary, E. Aurell, and S. Haridi, "A Statistical Theory of Chord Under Churn," in *Proc. IPTPS*, Feb. 2005, pp. 93-103.
- [14] S.S. Lam and H. Liu, "Failure Recovery for Structured P2P Networks: Protocol Design and Performance Evaluation," in *Proc. ACM SIGMETRICS*, June 2004, pp. 199-210.
- [15] D. Leonard, Z. Yao, V. Rai, and D. Loguinov, "On Lifetime-Based Node Failure and Stochastic Resilience of Decentralized Peer-to-Peer Networks," *IEEE/ACM Trans. Netw.*, vol. 15, no. 3, pp. 644-656, June 2007.
- [16] D. Leonard, Z. Yao, X. Wang, and D. Loguinov, "On Static and Dynamic Partitioning Behavior of Large-Scale P2P Networks," *IEEE/ACM Trans. Netw.*, vol. 16, no. 6, pp. 1475-1488, Dec. 2008.
- [17] J. Li, J. Stribling, T.M. Gil, R. Morris, and M.F. Kaashoek, "Comparing the Performance of Distributed Hash Tables under Churn," in *Proc. IPTPS*, Feb. 2004, pp. 87-99.
- [18] J. Li, J. Stribling, R. Morris, and M.F. Kaashoek, "Bandwidth-Efficient Management of DHT Routing Tables," in *Proc. USENIX NSDI*, May 2005, pp. 99-114.
- [19] D. Liben-Nowell, H. Balakrishnan, and D. Karger, "Analysis of the Evolution of the Peer-to-Peer Systems," in *Proc. ACM PODC*, July 2002, pp. 233-242.
- [20] B. Lindner, "Superposition of Many Independent Spike Trains is Generally not a Poisson Process," *Phys. Rev. E*, vol. 73, no. 2, pp. 022901-1-022901-4, 2006.
- [21] G. Pandurangan, P. Raghavan, and E. Upfal, "Building Low-Diameter Peer-to-Peer Networks," *IEEE J. Sel. Areas Commun.*, vol. 21, no. 6, pp. 995-1002, Aug. 2003.
- [22] S. Resnick, *Adventures in Stochastic Processes*. Cambridge, MA, USA: Birkhäuser, 2002.
- [23] S. Rhea, D. Geels, T. Roscoe, and J. Kubiatowicz, "Handling Churn in a DHT," in *Proc. USENIX ATC*, June 2004, pp. 127-140.
- [24] S. Saroiu, P.K. Gummadi, and S.D. Gribble, "A Measurement Study of Peer-to-Peer File Sharing Systems," in *Proc. SPIE/ACM Multimedia Comput. Netw.*, Jan. 2002, vol. 4673, pp. 156-170.
- [25] I. Stoica, R. Morris, D. Karger, M.F. Kaashoek, and H. Balakrishnan, "Chord: A Scalable Peer-to-Peer Lookup Service for Internet Applications," in *Proc. ACM SIGCOMM*, Aug. 2001, pp. 149-160.
- [26] D. Stutzbach and R. Rejaie, "Understanding Churn in Peer-to-Peer Networks," in *Proc. ACM IMC*, Oct. 2006, pp. 189-202.
- [27] X. Wang, Z. Yao, and D. Loguinov, "Residual-Based Estimation of Peer and Link Lifetimes in P2P Networks," *IEEE/ACM Trans. Netw.*, vol. 17, no. 3, pp. 726-739, June 2009.
- [28] R.W. Wolff, *Stochastic Modeling and the Theory of Queues*. Englewood Cliffs, NJ, USA: Prentice-Hall, 1989.
- [29] Z. Yao, X. Wang, D. Leonard, and D. Loguinov, "On Node Isolation under Churn in Unstructured P2P Networks With Heavy-Tailed Lifetimes," in *Proc. IEEE INFOCOM*, May 2007, pp. 2126-2134.
- [30] M. Zhong, K. Shen, and J. Seiferas, "Non-Uniform Random Membership Management in Peer-to-Peer Networks," in *Proc. IEEE INFOCOM*, Mar. 2005, pp. 1151-1161.



Zhongmei Yao received the BS degree in engineering from Donghua University, Shanghai, China, in 1997, the MS degree in computer science from Louisiana Tech University, Ruston, in 2004, and the PhD degree in computer science from Texas A&M University, College Station, TX, USA, in 2009. Currently, she is an Assistant Professor of computer science at the University of Dayton, Dayton, OH. Her research interests include P2P networks, stochastic modeling, and performance analysis. She is a Student Member of the IEEE.



Daren B.H. Cline received the BS degree in mathematics from Harvey Mudd College, Claremont, CA, USA, in 1978, and the MS and PhD degrees in statistics from Colorado State University, Fort Collins CO, USA, in 1980 and 1983, respectively. He joined the Statistics Department at Texas A&M University, College Station, TX, USA, in 1984, and is now a tenured Professor. His current research interests include applied stochastic processes, Markov chain theory, heavy-tailed distributions, and nonlinear

time series models.



Xiaoming Wang received the BS degree in computer science and the MS degree in electronic engineering from Beijing University of Posts and Telecommunications, Beijing, China, in 1999 and 2002, respectively, and the PhD degree in computer science from Texas A&M University, College Station, TX, USA, in 2009. Currently, he works for Amazon.com, Seattle, WA, USA. His research interests include peer-to-peer systems, probabilistic analysis of computer networks, and topology modeling. He is a Student Member of the IEEE.



Dmitri Loguinov received the BS degree (with honors) in computer science from Moscow State University, Moscow, Russia, in 1995, and the PhD degree in computer science from the City University of New York, New York, in 2002. Currently, he is a Professor in the Department of Computer Science and Engineering at Texas A&M University, College Station, TX, USA. His research interests include P2P networks, information retrieval, congestion control, and Internet measurement and modeling. He is a Senior Member of the IEEE.

▷ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.