

Modeling Heterogeneous User Churn and Local Resilience of Unstructured P2P Networks

Zhongmei Yao, Derek Leonard, Xiaoming Wang, and Dmitri Loguinov
 Department of Computer Science, Texas A&M University
 College Station, TX 77802 USA
 {mayyao, dleonard, xmwang, dmitri}@cs.tamu.edu

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 generic model of heterogeneous user churn, derive the distribution of the various metrics observed in prior experimental studies (e.g., lifetime distribution of joining users, joint distribution of session time of alive peers, and residual lifetime of a randomly selected user), derive several closed-form results on the transient behavior of in-degree, and eventually obtain the joint in/out degree isolation probability as a simple extension of the out-degree model in [13].

I. INTRODUCTION

Peer-to-peer (P2P) networks have recently emerged as an efficient and highly resilient platform for large-scale distributed applications. One of the fundamental problems in comprehending 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 surviving users to counteract the effects of abrupt departures. Unlike other distributed systems where failures may be considered rare or abnormal, most P2P networks constantly remain in the state of churn and embrace frequent failures as part of their normal operation.

While many metrics of a system (e.g., search latency, path existence probability, efficiency of routing, message overhead, file popularity) affect its usefulness to the user, one commonly studied problem in the literature is the ability of P2P networks to stay connected in the face of random failures [1], [2], [6], [8], [10], [11], [12], [13], [15], [17], [21], [19], [23]. It may be argued that compromised connectivity is one of the most fundamental byproducts of churn that directly affects routing efficiency and other metrics observed by the user. However, before resilience and performance of P2P networks can be fully understood, a good model of churn is required since even today most analytical models that consider churn [11], [13], [17], [19] do not completely 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.

A. Churn Model

We start the paper with a goal of modeling churn in P2P systems and striking a balance between model complexity and its fidelity. Our focus for the time being lies with networks in which neighbor selection occurs only during the initial join into the system and failure of an existing neighbor. This

framework is commonly used in *unstructured* P2P implementations (such as the ultra-peer layer of Gnutella) where replacement takes place as a reactive response to failed keep-alive messages. Certain DHTs in which neighbor pointers do not have to change after arrival of new users (e.g., Randomized Chord [18]) fall into this category as well.¹ This type of churn was originally formalized in [13], where Leonard *et al.* equipped users with random lifetimes L_i that determined the duration of their presence in the system and modeled neighbor replacement using random delays S_i that included the timeouts to detect each failure and protocol delays to actually obtain a new neighbor. This model, however, treated P2P users equally in their online characteristics (i.e., all user lifetimes were drawn from the same distribution) and did not consider their offline behavior as having any impact on churn.

Heterogeneity of lifetimes is a fundamental property of P2P systems where some users consistently spend substantial periods of time in the system and others very little. This observation prompts the question of *whether P2P systems can indeed be modeled using a single homogeneous lifetime distribution 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], [22], 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 (e.g., distribution of in and out-degree at each time t , probability that a given neighbor is alive, isolation probability within a lifetime) of each user.

To answer these questions, we offer a generic churn model that captures the *heterogeneous* behavior of end-users, including their difference in online habits and diversity of offline “think time.” 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 are drawn from (potentially different) distributions $F_i(x)$ and $G_i(x)$ unique to each user i . This approach captures the fact that the same user in subsequent sessions is likely to exhibit characteristics resembling those of its prior sessions, while simultaneously allowing for cases where similarity between different users is arbitrarily low.

Armed with this model, we obtain the distribution $F(x)$ of lifetimes of *newly arriving* users, the joint lifetime distribution $J(x)$ of the users *currently in the system*, and the *residual*

¹Analysis of other types of DHTs where arriving users actively replace existing neighbors will be covered in a separate paper.

lifetime distribution $H(x)$ of a randomly selected alive user in the network. Our results show that all three metrics are weighted functions of individual lifetime distributions $F_i(x)$, where $H(x)$ is additionally dependent on the number of alive users currently in the graph, the probability that user i is picked by a joining peer, and the conditional residual lifetimes of the neighbors chosen by the selection method. The model for $H(x)$ is extremely complex and generally intractable unless neighbor selection is performed *uniformly among the alive users*², in which case we show that $H(x)$ 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 users to obtain $F(x)$, which is then *sufficient to entirely model the effect of churn* on P2P graphs.

We also revisit the observation of [24] 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* [25] of distribution $F(x)$, which allows us to prove that random users currently in the system have stochastically larger lifetimes than random arriving users *regardless of the shape of distributions* $F_i(x)$ and $G_i(x)$. We additionally show that while $F(x)$ may be heavy-tailed as observed in practice [5], [9], [16], it is possible that individual lifetime distributions $F_i(x)$ may *all* be exponential, or contain a mix of exponential and heavy-tailed distributions. Occurrence of this effect depends on random availability of each user and shows that conclusions on the individual habits of peers may not be drawn from their aggregate behavior $F(x)$.

B. Local Resilience Model

In the second half of the paper, we tackle the issue of node resilience to isolation in the presence of churn under the assumption that neighbors provide *mutual* resilience benefits to each other (i.e., both outgoing and incoming edges increase resilience of peers). Prior work [13], [14] showed that many P2P graphs stayed connected if and only if they did not develop isolated nodes during churn and derived that individual node isolation probabilities were functions of $\rho = E[L_i]/E[S_i]$, where $E[L_i]$ was the mean lifetime of homogeneous users and $E[S_i]$ was the mean search (i.e., node-replacement) delay. However, despite its importance, this approach only modeled the *out-degree* of each user and did not consider the increased resilience arising from additional *in-degree* edges arriving “in the background” to each user during its stay in the system.

We overcome this shortcoming and build a complete closed-form model characterizing the evolution of in-degree in unstructured systems under the assumption of uniform neighbor selection. We first show that under certain mild assumptions the edge arrival process to each user tends to a Poisson distribution when the system size becomes sufficiently large, which is consistent with recent observation of this phenomenon in

certain real networks [24]. We then derive the transient distribution of in-degree as a simple function of $F(x)$, *including cases with non-exponential peer lifetimes*, and show that users who stay online longer quickly accumulate non-trivial in-degree and become much more resilient to isolation over time. This outcome was intuitively expected as it makes sense that current unstructured P2P networks have been designed such that users with more contribution to the system (i.e., longer lives) become better connected over time and provide more search capabilities to their neighbors. In contrast, the original model of [13] showed that P2P users became progressively more susceptible to isolation as their age increased.

We finish the paper by combining the in and 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 (i.e., loss of all neighbors within its lifetime) and by ϕ_{out} the same metric with only the out-degree being considered [13], we show that for exponential $F(x)$ the following holds as search delays become asymptotically small (i.e., tend to zero):

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

where k is the initial number of neighbors obtained by each arriving user. This result illustrates that the amount of improvement from the in-degree amounts to approximately a factor of $2k$ reduction in the isolation probability. Also note that for non-negligible search delays, ratio ϕ_{out}/ϕ is larger than implied by the worst-case bound in (1), but does not admit a simple closed-form expression. We finish the paper with examples and simulations that demonstrate this effect.

II. RELATED WORK

One of the first models of churn was proposed in [19], which assumed an unstructured P2P system with Poisson arrivals and departures that could be modeled as an $M/M/1$ queue. Neighbor replacement in this system was in direct response to failures and was assumed to be instantaneous, where the possibilities for replacement were limited to the nodes currently alive in a certain centralized cache. The paper showed that under user churn the graph remained connected and exhibited a logarithmic diameter, both with high probability.

Later models of churn [17] and recently [11] assumed a DHT-like system in which repair algorithms were run *independently* of user failures and at exponentially distributed intervals (i.e., as Poisson processes). This approach modeled the consistency check algorithm in Chord, which periodically verified the successor list and corrected invalid pointers. While viable in general, this model required Chord to change neighbors frequently (i.e., $\Omega(\log n)$ new neighbors per half-life) and stabilize at rather high rates [11] since repairs were not correlated with actual successor failures. Additionally, these models assumed homogeneous exponential lifetimes and Poisson arrival/departure processes with no way of generalizing their results to non-exponential system dynamics.

A different approach was undertaken in [13], where neighbor replacements were explicitly initiated in response to failed links. In this setup, each joining user randomly selected k

²This can be implemented by picking users from uniformly random subsets of cached nodes or using special random walks on the graph [27].




Fig. 1. Churn model depicting ON/OFF behavior $Z_i(t)$ of user i .

neighbors from the graph and then monitored their online presence using keep-alive messages.³ Once the failure of an existing neighbor was detected, a uniformly random replacement was sought from among the currently alive users in the system. Detection and replacement delays were also random, but explicitly non-zero. Under these conditions, the paper showed that each user became isolated with probability no larger than $\phi_{out} = k\rho/(1+\rho)^k$, where ρ was the ratio of the average lifetime to the average replacement delay, for all lifetime distributions with an exponential or heavier tail. This result was later generalized in [14] to show that the probability of non-partitioning in many P2P networks converged as $n \rightarrow \infty$ to that of avoiding isolation for each alive user.

III. 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.

Due to limited space, all proofs omitted from this paper can be found in the technical report [26].

A. Churn Model

Consider a P2P system with n participating users, where each user i is either alive (i.e., present in the system) at time t or dead (i.e., logged off). This behavior can be modeled by an alternating renewal process $\{Z_i(t)\}$ for each user i :

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

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

We next make several modeling assumptions about this system. First, we suppose that users behave independently of each other and that processes $\{Z_i(t)\}$ and $\{Z_j(t)\}$ for any $i \neq$

j are independent. This means that users do not synchronize their arrival or departures and generally exhibit uncorrelated lifetime characteristics. This assumption also implies that users simultaneously present in the system with multiple identities are not very common and have no large-scale impact on the dynamics of the network. Second, we assume for each process $\{Z_i(t)\}$ that its ON durations $\{L_{i,c}\}_{c=1}^{\infty}$ have some joint distribution $F_i(x)$ and that its OFF durations $\{D_{i,c}\}_{c=1}^{\infty}$ have another joint distribution $G_i(x)$. While this model is generic enough to allow dependency between cycle lengths, the results developed later in the paper are sensitive *not* to the correlational structure of processes $\{L_{i,c}\}_{c=1}^{\infty}$ and $\{D_{i,c}\}_{c=1}^{\infty}$, but to the *distribution* of random variables comprising each of them. We thus, without loss of generality, omit discussion of correlated sequences and treat all lifetime and offtime processes as *i.i.d.* sets of variables. This allows us to replace $\{L_{i,c}\}_{c=1}^{\infty}$ with a random variable $L_i \sim F_i(x)$ and $\{D_{i,c}\}_{c=1}^{\infty}$ with $D_i \sim G_i(x)$.

Examples of ON/OFF distributions commonly considered in this paper are the exponential:

$$F_i(x) = 1 - e^{-\lambda_i x} \quad (3)$$

with mean $1/\lambda_i$ and the shifted Pareto:

$$F_i(x) = 1 - (1 + x/\beta_i)^{-\alpha_i}, \quad x > 0, \alpha_i > 1 \quad (4)$$

with mean $\beta_i/(\alpha_i - 1)$. For convenience of notation, define the average lifetime $l_i = E[L_i]$ and the average offline duration $d_i = E[D_i]$.

From Smith's theorem, we easily obtain that the asymptotic *availability* [4], [22] of each user i , i.e., the probability that it is in the system at a random instance $t \gg 0$, is given by:

$$a_i = \lim_{t \rightarrow \infty} P(Z_i(t) = 1) = \frac{l_i}{l_i + d_i}. \quad (5)$$

The final metric related to our churn model is the distribution of the number of users in the graph. Denote by $N(n, t) = \sum_{i=1}^n Z_i(t)$ the number of users in the network at time t and notice that it is also a random process that fluctuates with time. In cases when the size of the system is understood, we will write $N(t)$ instead of $N(n, t)$.

Since most P2P properties of interest require a sufficiently evolved system with $t \rightarrow \infty$, our analysis below examines the distribution of the various metrics in the *equilibrium*. Our next result shows that the limiting distribution of $N(t)$ is Gaussian with mean and variance solely determined by the availability set $\{a_1, \dots, a_n\}$.

Lemma 1: The number of users $N = \lim_{t \rightarrow \infty} N(t)$ observed in the equilibrium tends to a Gaussian random variable with mean $\sum_{i=1}^n a_i$ and variance $\sum_{i=1}^n a_i(1 - a_i)$ as $n \rightarrow \infty$.

Before we show simulations with our churn model, we define the rules for producing heterogeneous users and their lifetime/offtime distributions. The only parameters that control each user's availability are l_i and d_i , which we draw for each user randomly from two Pareto distributions with $\alpha = 3$ as described next. For the mean ON duration, we use $\beta = 1$ and obtain $E[l_i] = 1/2$ hour; for the mean OFF duration, we use $\beta = 2$ and get $E[d_i] = 1$ hour. Note that once pair (l_i, d_i) is selected for user i , it remains constant for the entire

³Gnutella, for example, sends a ping message every 3 seconds and detects link failure when TCP declares the connection aborted, which happens after several (e.g., 5 in Windows) subsequently failed retransmission attempts.




Fig. 2. Sample path and distribution of $N(t)$ in system \mathcal{H} with $n = 1000$ users. The Gaussian fit is from Lemma 1.

evolution of the system. The actual lifetimes and offtimes of each cycle are drawn from distributions $F_i(x)$ and $G_i(x)$, which are selected such that their means are equal to the corresponding l_i and d_i . We study three cases throughout the paper: 1) heavy-tailed system \mathcal{H} with $F_i(x) \sim \text{Pareto}(3, 2l_i)$ and $G_i(x) \sim \text{Pareto}(3, 2d_i)$; 2) very heavy-tailed system \mathcal{VH} with $F_i(x) \sim \text{Pareto}(1.5, l_i/2)$ and $G_i(x) \sim \text{Pareto}(1.5, d_i/2)$; and 3) exponential system \mathcal{E} with $F_i(x) \sim \exp(1/l_i)$ and $G_i(x) \sim \text{Pareto}(3, 2d_i)$, where notation $\text{Pareto}(\alpha_i, \beta_i)$ refers to (4) and $\exp(\lambda_i)$ to (3).

Fig. 2(a) shows one example for the evolution of system size $N(t)$ as a function time t , which as $n \rightarrow \infty$ tends to simple Brownian motion (also known as the limiting distribution of random walks). Part (b) of the figure shows the PMF of $N(t)$ at $t \gg 0$ and a Gaussian fit from the previous lemma (note the log-scale of the y -axis). Numerous additional results confirming that $N(t)$ is a Gaussian process are omitted for brevity.

B. Lifetime Distribution of Joining Users

The goal of this subsection is to derive the distribution of user lifetimes when they join the system. While each user has a different distribution of ON/OFF durations, we show below that the aggregate behavior of the system can be reduced to that of individual users in a simple closed-form expression. Suppose that random variable L describes the lifetime of the next user that enters the system and define $F(x) = P(L < x)$ to be its CDF. Note that L can also be viewed as the lifetime distribution of *all* users who have ever visited the system.

Before we proceed to the next result, let λ_i be the reciprocal of the mean duration of a life-death cycle of process $\{Z_i(t)\}$:

$$\lambda_i = \frac{1}{l_i + d_i}, \quad (6)$$

which can also be interpreted as the average arrival (or departure) rate of user i . Moreover, denote by b_i the ratio of λ_i to the total rate of arrival in the system:

$$b_i = \frac{\lambda_i}{\sum_{j=1}^n \lambda_j} = \left(\sum_{j=1}^n \frac{l_j + d_j}{l_j + d_j} \right)^{-1}. \quad (7)$$

The next lemma shows that the lifetime distribution of all users put together is a weighted average of individual lifetime

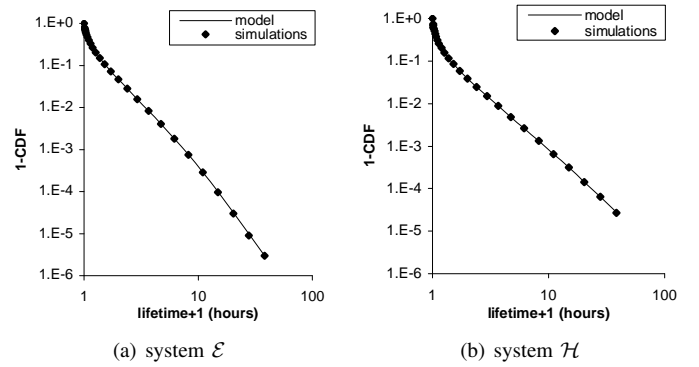


Fig. 3. Comparison of simulation results of $F(x)$ to model (8) in a graph with $n = 1000$ nodes.

distributions, where the weights are biased toward those peers who frequently join and leave the system since their sessions constitute the majority of overall peer arrival into the system.

Theorem 1: The distribution of lifetime of the next joining user is given by the following weighted average:

$$F(x) = \sum_{i=1}^n b_i F_i(x), \quad (8)$$

where $F_i(x)$ is the lifetime CDF of user i .

Similarly, we obtain the following results.

Corollary 1: Random offline duration D of the next failing user is given by $G(x) = P(D < x) = \sum_{i=1}^n b_i G_i(x)$, where $G_i(x)$ is the offtime CDF of user i . Furthermore, the mean session time of a joining user and the mean OFF period of a failing user are respectively $E[L] = \sum_{i=1}^n b_i l_i$ and $E[D] = \sum_{i=1}^n b_i d_i$.

We next verify (8) and discuss several implications of this result. Two typical simulations are presented in Fig. 3 for exponential and heavy-tailed lifetimes, both of which show that the model is very consistent with simulation results. Note that both figures are on log-log scale and plot $1 - F(x)$ vs. $1 + x$ to make the shifted Pareto distribution in (4) appear as a straight line. Notice in Fig. 3(a) that system \mathcal{E} produces a heavy-tailed aggregate distribution $F(x)$ even though *all individual lifetime distributions $F_i(x)$ are exponential*. This can be explained as follows. It is well-known [3], [7] that for a hyper-exponential distribution in the form of (8) and *any* desired distribution $W(x)$ with a monotonic PDF, there exists a set of weights $\{b_1, \dots, b_n\}$ such that (8) 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 one can observe a nicely shaped Pareto, Weibull, or other distribution $F(x)$, which is produced by a mixture of exponentials $F_i(x)$. It may therefore be preliminary to conclude that Pareto $F(x)$ measured experimentally [5], [22] provides any indication as to the true nature of $F_i(x)$ or 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 $F(x)$ can be used to characterize the parameters of a single user selected from the system randomly?* We focus on two such metrics

below: the remaining online duration of a randomly selected user from among the peers currently alive and the joint lifetime distribution of all alive peers at sufficiently large times t .

C. Residual Lifetime Distribution

Assume that a user v selects some number of neighbors from among the existing peers in the system and does not replace them until they die. The mechanism for initial selection and subsequent replacement may be uniformly random, based on the age of existing users, or influenced by the degree and structure of the graph (e.g., random walks). Let $R_i(t)$ denote the length of the remaining life cycle of a given user i at time t , i.e., the remainder of the current ON cycle illustrated in Fig. 1. Note that $R_i(t)$ is important to peer v since this metric determines how long this neighbor will be online *after it is selected by v* . As shown in [13], the residual lifetime distribution solely determines the resilience of the out-degree isolation model.

Define the equilibrium residual distribution of user i :

$$H_i(x) = \lim_{t \rightarrow \infty} P(R_i(t) < x | Z_i(t) = 1) \quad (9)$$

and note that the offtime distribution $G_i(x)$ has no impact on $H_i(x)$ since (9) is conditioned on the fact that user i is alive at t (i.e., when a user is OFF, it is never selected by other peers for any purposes). Then from renewal theory [20], the residual lifetime distribution of user i is:

$$H_i(x) = \frac{1}{E[L_i]} \int_0^x (1 - F_i(u)) du, \quad (10)$$

where $F_i(x)$ is the lifetime distribution of user i as before.

Next, let $R(t)$ be the residual lifetime of the user *randomly selected* by a peer v at time t , $R = \lim_{t \rightarrow \infty} R(t)$ be its equilibrium version, and $H(x) = P(R < x)$ be its CDF. To understand the next result, denote by:

$$Y_i(t) = \begin{cases} 1 & \text{user } i \text{ is selected at } t \\ 0 & \text{otherwise} \end{cases} \quad (11)$$

the indicator process that shows whether user i is picked by v at time t and by:

$$V_i(x) = \lim_{t \rightarrow \infty} P(R_i(t) < x | Y_i(t) = 1, Z_i(t) = 1) \quad (12)$$

the equilibrium CDF of residual lifetimes of the users selected for replacement. Further define:

$$s_{ij} = \lim_{t \rightarrow \infty} P(Y_i(t) = 1 | Z_i(t) = 1, N(t) = j) \quad (13)$$

to be the probability that i is selected among j existing users. Note that our prior assumptions on the churn model dictate that $V_i(x)$ cannot depend on $N(t)$, i.e., users do not change their lifetime behavior based on the size of the system.

We next elaborate on definitions (12)-(13). In systems where the residual lifetime distribution of a user does not change based on whether that peer has been chosen or not, $V_i(x)$ is identical to $H_i(x)$. This means that $V_i(x) = H_i(x)$ holds only in cases when neighbor selection is *independent* of the lifetimes (or ages) of the selected users (e.g., this model was used in [13]). Examples that satisfy this condition include

uniform selection, selection based on content similarity or random hashing space, age-independent popularity, etc. On the other hand, selection based on the age of potential neighbors or random walks (which depend on the in-degree of each user, which in turn depends on age) do not fall into this category.

To understand s_{ij} , we have the following examples. For uniform selection, s_{ij} is simply $1/j$. For stationary random walks, $s_{ij} = d_i / \sum_{m=1}^j d_m$, where d_i is the stationary in-degree of user i conditioned on the fact that it is alive. For content-based selection, assume that each user shares w_i files with others and that each peer is selected to be a neighbor proportionally to its ‘‘content utility’’ w_i . Then, s_{ij} may be equal to $w_i / \sum_{m=1}^j w_m$.

As must be evident, (12)-(13) can model arbitrarily complex rules for choosing neighbors; however, tractability of the resulting distribution $H(x)$ is questionable for all except the simplest cases. Below, we first derive $H(x)$ for the most generic case and show that it can be expressed as a sum of weighted individual residual distributions, where the weights are biased towards users with large availability a_i and high probability s_{ij} of being selected. We then simplify this expression for uniform selection.

Theorem 2: The residual lifetime distribution of a random neighbor in an equilibrium P2P system is given by:

$$H(x) = \sum_{i=1}^n V_i(x) a_i \sum_{j=1}^n s_{ij} P(N(n-1) = j-1), \quad (14)$$

where $N(n-1) = \lim_{t \rightarrow \infty} N(n-1, t)$.

The above theorem shows that $H(x)$ depends on a number of complex factors, including the distribution of $N(n-1)$, CDF function $V_i(x)$, and metric s_{ij} . While for large n , the distribution of $N(n-1)$ is very close to that of N , the resulting Gaussian distribution has no suitable expressions for expanding (14) in closed-form. Furthermore, without the knowledge of $V_i(x)$ (which depends on $F_i(x)$ as well as all other $F_j(x)$, $j \neq i$), computation of $H(x)$ is intractable.

Fortunately, for certain neighbor-selection strategies we can circumvent the difficulty of obtaining $P(N(n-1) = j-1)$ in closed-form, as shown in the next lemma.

Lemma 2: The residual lifetime distribution of neighbors under uniform selection is given by:

$$H(x) = \frac{1}{\sum_{i=1}^n a_i} \sum_{i=1}^n a_i H_i(x), \quad (15)$$

where $H_i(x)$ is the residual CDF of user i in (10).

Next, we run simulations to confirm Lemma 2. As shown in Fig. 4 for the exponential and Pareto cases, simulation results of $H(x)$ match the model very well and also demonstrate that \mathcal{E} may produce non-exponential residual lifetime distributions. While yet again $H(x)$ in (15) requires the knowledge of $F_i(x)$ and a_i for $i = 1, \dots, n$, our next result shows that $H(x)$ under uniform selection may be directly obtained from $F(x)$.

Theorem 3: Under uniform selection, the equilibrium residual distribution $H(x)$ of random neighbors can be reduced to the lifetime distribution of joining users:

$$H(x) = \frac{1}{E[L]} \int_0^x (1 - F(u)) du, \quad (16)$$

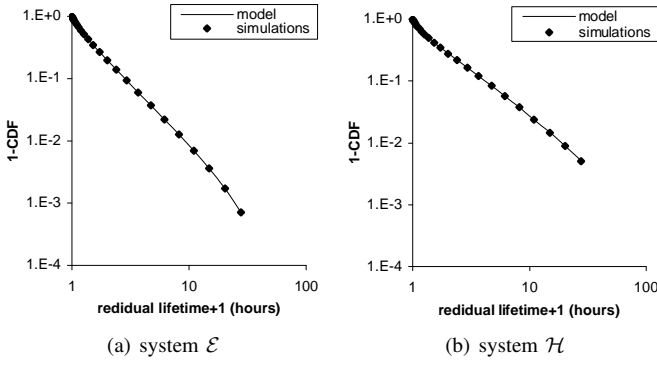


Fig. 4. Comparison of simulation results of $H(x)$ to model (15) in a graph with $n = 1000$ nodes.

where $E[L]$ is given in Corollary 1 and $F(x)$ is in (8).

Note that (16) is an extremely important result since it shows that in practice one only needs to measure the aggregate lifetime distribution $F(x)$ and its mean $E[L]$ rather than each ON distribution $F_i(x)$ and each user availability a_i in order to obtain the residual lifetime distribution of a uniformly selected neighbor. Assuming from measurement studies [5], [9], [16], that $F(x)$ is Pareto with $F(x) = 1 - (1 + x/\beta)^{-\alpha}$, (16) reduces to:

$$H(x) = 1 - (1 + x/\beta)^{-(\alpha-1)} \quad (17)$$

as also shown in [13]. Comparing (17) to $F(x)$, we see that R is stochastically larger than user lifetimes L , i.e., $P(R > x) \geq P(L > x)$, which implies that a uniformly randomly selected user is more reliable than new arrivals in terms of failure. For other neighbor selection strategies, it is important to realize that the distribution of residual lifetimes R may be completely different from (16) and should be analyzed accordingly.

D. Lifetime Distribution of Users in the System

Denote by $J(x)$ the joint lifetime distribution of users *currently* in the system. As observed in [24], distribution $J(x)$ is clearly different from $F(x)$; however, no closed-form analysis has been made available to date. The intuitional rationale behind this difference is that lifetimes of the peers observed in the system are biased towards larger values, which is commonly known as the *inspection paradox* [25]. In the next result, we formally derive $J(x)$ as a function of $F(x)$.

Theorem 4: The joint lifetime distribution $J(x)$ of alive users in the equilibrium is:

$$J(x) = \frac{1}{E[L]} \left(xF(x) - \int_0^x F(u)du \right), \quad (18)$$

where $E[L]$ is given in Corollary 1 and $F(x)$ is in (8).

The accuracy of (18) was confirmed in simulations, but is omitted here for brevity. Since spread is the convolution of two distributions $H(x)$, it follows that exponential lifetimes $F(x)$ imply that $J(x)$ is the Erlang(2) distribution with mean $2E[L]$. For Pareto $F(x)$, spread $J(x)$ is a convolution of two distributions in (17), which 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 [24], in more formal terms.

Corollary 2: For any set of lifetime distributions $\{F_i(x)\}_{i=1}^n$, the spread $J(x)$ is stochastically larger than $F(x)$ and the mean lifetime of a user currently alive in the system is double the mean residual lifetime $E[R]$ of a uniformly selected user.

In conclusion, our results demonstrate that given heterogeneous users and uniform selection of neighbors, both metrics $H(x)$ and $J(x)$ can be reduced to the aggregate behavior $F(x)$ of arriving users. 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.

IV. IN-DEGREE MODEL

Before analyzing node in-degree, we study the process of edge arrival into each user under uniform selection since this metric determines both the rate at which each user accumulates incoming neighbors and the stationary in-degree distribution. In the following, we first examine the *user* arrival and failure process in a P2P system driven by our churn model, from which we obtain the behavior of *edge* arrival to each user.

A. Node Arrival Process

Recall that our churn model prescribes that each arriving node find k random initial neighbors and then continuously replace them as they fail. Let $X(t)$ be the number of in-degree neighbors accumulated by a random user by age t , L be its random lifetime, and $\theta = E[X(L)]$ its mean number of in-degree neighbors when it dies. Under these definitions, each new arrival inserts an average of k edges and each failure brings an average of θ edges into the system.

Define *new edges* to be those added by the arrival of new users, which occurs when a process $\{Z_i(t)\}$ transitions from 0 to 1, and *old edges* to be those added in response to the failure of existing users, which happens S time units after $\{Z_i(t)\}$ transitions from 1 to 0. Since search delay S is usually negligible compared to ON/OFF cycle durations, we do not explicitly model its effect to save space.

For each ON/OFF renewal process $\{Z_i(t)\}$ of user i , denote by T_k the k -th arrival instance of user i . Then let:

$$M_i^a(t) = \max\{k : T_{k-1} \leq t\} \quad (19)$$

be the total number of arrivals of user i in the interval $[0, t]$ for $t > 0$. Hence, $\{M_i^a(t) : t \geq 0\}$ is a counting process with inter-arrival cycles given by $L_i + D_i$.

Theorem 5: Superposition arrival process: $M^a(t) = \sum_{i=1}^n M_i^a(t)$ converges as $n \rightarrow \infty$ to a homogeneous Poisson process with rate:

$$\lambda = \sum_{i=1}^n \lambda_i < \infty, \quad (20)$$

where λ_i is the reciprocal of $l_i + d_i$ given by (6).

This conclusion is confirmed for the very heavy-tailed system \mathcal{VH} (i.e., $\alpha = 1.5$) in Fig. 5. Notice in the figure that for small n , the arrival process is not Poisson; however, as n becomes sufficiently large, the inter-arrival delay




Fig. 5. Distribution of user inter-arrival delays in system $\mathcal{V}\mathcal{H}$ approaches exponential with rate λ given in (20) as n becomes sufficiently large.




Fig. 6. Distribution of the number of user arrivals in interval $[t, t + \Delta t]$ in a system with $n = 1000$ users. The lines show Poisson fits with λ in (20).

is perfectly exponential. Furthermore, the rate of arrival λ follows (20) very well. Next, we verify that the distribution of $M^a(t, t + \Delta t)$, i.e., the number of arrivals from n sources in any fixed interval $[t, t + \Delta t]$, converges to the Poisson distribution with parameter λ as $n \rightarrow \infty$. We show two typical simulation results of the distribution of the number of arrivals from $n = 1000$ sources in the interval $[t, t + \Delta t]$ in Fig. 6. In part (a), we have heavy-tailed system \mathcal{H} with $\Delta t = 36$ seconds, while in part (b) we have very heavy-tailed system $\mathcal{V}\mathcal{H}$ with $\Delta t = 43.2$ seconds. Observe from Fig. 6 that in each case, simulation results of $M^a(t, t + \Delta t)$ show a good match with a Poisson distribution with parameter λ . For other ON/OFF distributions, simulations are also consistent with the properties of a Poisson process and are not shown here for brevity.

We next follow the same steps to study the user failure (i.e., departure) process. Denote by $M_i^d(t)$ the number of departures of user i in the time interval $[0, t]$. Then it is clear that $\{M_i^d(t) : t \geq 0\}$ is a counting process with random inter-departure delays $D_i + L_i$.

Corollary 3: Superposition departure process $M^d(t) = \sum_{i=1}^n M_i^d(t)$ converges as $n \rightarrow \infty$ to a homogeneous Poisson process with rate λ in (20).

Finally, note that the Poisson results above are not an *assumption* of the paper as in prior work [11], [17], [19], but rather a *consequence* of the churn model introduced earlier. Armed with these observations, we next focus on the edge arrival process to each user.




Fig. 7. Distribution of edge inter-arrival delays approaches exponential with rate λ_v in (21) for $k = 10$ and $\theta = 10$.




Fig. 8. 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 $\theta = 10$. The lines show Poisson fits with λ_v in (21).

B. Edge Arrival Process

In this subsection, we analyze the arrival process of edges to individual users assuming uniform selection. For other strategies for obtaining users, the analysis is significantly more complex and will be considered in future work.

Theorem 6: Edge arrival into a random user under uniform selection converges as $n \rightarrow \infty$ to the Poisson process with rate:

$$\lambda_v = \frac{k + \theta}{E[L]}, \quad (21)$$

where k is the number of neighbors a user selects upon joining the system, θ is the mean in-degree of a user when it leaves the system, and $E[L]$ is the mean peer lifetime.

Fig. 7 shows the distribution of edge inter-arrival delays to a single node obtained in simulations in a system with 1000 users. Notice in its two sub-figures, the distribution of inter-arrival delay approaches exponential with rate given by (21) for both systems. Additionally, Fig. 8 shows that the distribution of the number of edge arrivals to a node in a random interval of size Δt approaches the Poisson distribution with the same average rate λ_v in (21).

Equipped with Theorem 6, the evolution of a node's in-degree becomes tractable as we show below.

C. General Model of Expected Degree

As before, define $X(t)$ to be the in-degree of a random user v at age $t \geq 0$. In this subsection, we focus on transient

and stationary distributions of $X(t)$ under uniform selection of neighbors and assuming the main churn model of the paper.

Recall that the lifetime distribution of a new arrival is given by $F(x)$ in (8). Thus, once a *new* edge arrives to the current node v , it lives for a random time L whose CDF is given by $F(x)$. Moreover, each *old* edge arriving to v lives for a random duration R whose CDF is $H(x)$ in (16). As discussed earlier in this section, user v dies at some random time L drawn from $F(x)$ and its average degree at that time is $\theta = E[X(L)]$.

The following theorem solves the mean in-degree $E[X(t)]$ of a node at age t for any lifetime distribution.

Theorem 7: The expected in-degree of a random node v at age t is a monotonically increasing function of age given by:

$$E[X(t)] = \int_0^t \frac{k(1 - F(t - z)) + \theta(1 - H(t - z))}{E[L]} dz, \quad (22)$$

where θ is derived in [26].

Although $E[X(t)]$ in (22) appears complicated, it can be computed if the lifetime distribution $F(x)$ and residual lifetime distribution $H(x)$ are known. We next carry out this task using two distributions of user lifetimes and confirm the accuracy of the model.

D. Verification of Expected Degree

In the following, we use exponential and Pareto user lifetimes to compute $E[X(t)]$ and then verify the obtained result in simulations.

Theorem 8: For exponential lifetimes $L \sim \exp(\mu)$, the mean in-degree of a node at current age t is

$$E[X(t)] = 2k(1 - e^{-\mu t}) \quad (23)$$

and the mean in-degree at failure time is $\theta = k$.

From (23), observe that the mean in-degree of a node increases monotonically from $X(0) = 0$ when it arrives into the system to $X(\infty) = 2k$ when its age tends to infinity.

Next, we compute θ and $E[X(t)]$ for Pareto lifetimes.


Theorem 9: For Pareto $L \sim 1 - (1 + x/\beta)^{-\alpha}$ with $\alpha > 2$, the mean in-degree of a node at current age t is given by:

$$E[X(t)] = \theta \frac{E[R]}{E[L]} Q(t, \alpha - 2) + kQ(t, \alpha - 1), \quad (24)$$

where $Q(t, \alpha) = 1 - (1 + t/\beta)^{-\alpha}$, $E[L] = \beta/(\alpha - 1)$, $E[R] = \beta/(\alpha - 2)$, and the mean in-degree during failure is:

$$\theta = \frac{2k\beta}{2\beta + E[L]}. \quad (25)$$

Simulation results in Fig. 9 demonstrate that the models are very accurate and indeed saturate at predicted values $2k$ and $\theta E[R]/E[L] + k$ as age $t \rightarrow \infty$. It is also interesting to observe in the figure that 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. The saturation effect illustrated in Fig. 9 suggests that P2P implementations should cap user in-degree at values no smaller than the limit of (22) for $t \rightarrow \infty$. The corresponding upper bound in Gnutella (i.e., 30 in-degree neighbors) satisfies this condition for the two examples shown above.



(a) exponential lifetimes with $\mu = 2$

(b) Pareto lifetimes with $\alpha = 3$

Fig. 9. Comparison of models (23), (24) to simulation results for $E[L] = 0.5$ hours and $k = 8$.

In the next section, we focus on the probability that $X(t)$ is zero at any time t and later consider the impact of $X(t) = 0$ on the isolation probability.

E. Zero In-Degree

In this subsection, we use metric $P(X(t) = 0)$, i.e., the probability that a node's in-degree is zero at age t , to determine local in-degree resilience. The following theorem shows that regardless of the lifetime distribution, $P(X(t) = 0)$ is a simple function of expected in-degree $E[X(t)]$.

Theorem 10: The probability that the in-degree of a node is zero at given age t is:

$$P(X(t) = 0) = e^{-E[X(t)]}, \quad (26)$$

where $E[X(t)]$ is given by (22). Furthermore, $X(t)$ tends to a Poisson distribution with rate $E[X(t)]$ as $n \rightarrow \infty$.

Notice from (26) that $P(X(t) = 0)$ decreases in age t since $E[X(t)]$ increases with t as shown in (22). In other words, the longer a user stays alive in the system, the smaller the probability that it is disconnected because its in-degree is zero (of course, after $E[X(t)]$ becomes saturated, $P(X(t) = 0)$ stays almost constant). Substituting (23) and (24) into (26), we verify $P(X(t) = 0)$ for exponential and Pareto lifetimes in Table I, which demonstrates that both models of $P(X(t) = 0)$ are very accurate. Also observe in the table that after a user spends 2 hours in the system and with just 5 initial neighbors, its probability to have no in-degree links at time $t \geq 2$ is very small (i.e., less than 5.3×10^{-5}).

V. 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 actually very 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 incoming edges. In this section, we combine the in-degree and out-degree models to derive the *joint* isolation probability of an arriving user.

A. Preliminaries

As in [13], denote by $W(t)$ the out-degree of a node v at given age t and define it to be *isolated* when its in-degree and

TABLE I

MODEL (26) AND SIMULATIONS WITH $E[L] = 0.5$ HOURS AND $k = 5$

Age t hours	Exponential lifetimes		Pareto lifetimes	
	Simulation	Model	Simulation	Model
0.1	1.63×10^{-1}	1.63×10^{-1}	2.03×10^{-1}	2.03×10^{-1}
1.1	1.35×10^{-4}	1.37×10^{-4}	3.16×10^{-4}	3.17×10^{-4}
2.1	5.26×10^{-5}	5.27×10^{-5}	5.04×10^{-5}	5.02×10^{-5}
4.1	4.55×10^{-5}	4.55×10^{-5}	1.30×10^{-5}	1.31×10^{-5}
7.1	4.55×10^{-5}	4.54×10^{-5}	6.54×10^{-6}	6.55×10^{-6}

out-degree are simultaneously zero. Define *time to isolation* T to be the first-hitting time of both processes to state 0:

$$T = \inf\{t > 0 : W(t) = X(t) = 0 | W(0) = k, X(0) = 0\}. \quad (27)$$

Then the probability of node isolation is simply $\phi = P(T < L)$, where L is the random lifetime of node v . Note that unlike in the out-degree process, a node does not replace its incoming edges, which means that the in-degree and out-degree processes are independent of each other.

In the next subsections, we derive ϕ for systems with exponential user lifetimes and exponential search delays using two methods. 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.

B. Exponential Lifetimes (Exact Model)

Let pair $(W(t), X(t))$ be the joint process of out-degree and in-degree of a node 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 from Section IV that edge arrival at any node occurs according to a Poisson process with rate (21). Therefore, under uniform selection, new in-degree neighbors arrive to v at rate:

$$\lambda_v = \frac{k + \theta}{E[L]} = \frac{2k}{E[L]} \quad (28)$$

since $\theta = k$ for exponential lifetimes. Also note that the existing in-degree neighbors fail at rate $\mu = 1/E[L]$ due to the memoryless property of exponential distributions. This leads to the following result.

Lemma 3: For exponential lifetimes $L \sim \exp(\mu)$ and exponential search times $S \sim \exp(\sigma)$, the joint process $\{(W(t), X(t))\}$ is a homogeneous continuous-time Markov chain with a transition rate matrix Q derived in [26].

It is convenient to treat $\{(W(t), X(t))\}$ 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:

$$Q = \begin{pmatrix} 0 & 0 \\ \mathbf{r} & Q_0 \end{pmatrix}, \quad (29)$$

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 vector of transition rates into state $(0, 0)$. Before we proceed to the next result, define $\pi(0) = (\pi_{(i,j)}(0))$ to be the initial state distribution where each cell $\pi_{(i,j)}(0)$ of the row vector $\pi(0)$ is the probability that the chain starts with state

TABLE II

EXACT MODEL (30) AND SIMULATIONS FOR $E[L] = 0.5$ HOURS

$E[S]$ min	$k = 6$		$k = 8$	
	Simulations	Model (30)	Simulations	Model (30)
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}

(i, j) at time 0. Since chain $\{(W(t), X(t))\}$ always starts with state $(k, 0)$, we have $\pi_{(k,0)}(0) = 1$ and $\pi_{(i,j)}(0) = 0$ in all other initial states.

With this result in hand, we next obtain ϕ .

Theorem 11: For exponential lifetimes $L \sim \exp(\mu)$ and exponential search delays $S \sim \exp(\sigma)$, the probability of node isolation is given by:

$$\phi = \pi(0)VBV^{-1}\mathbf{r}, \quad (30)$$

where V is a matrix of eigenvectors of Q_0 , $B = \text{diag}(b_j)$ is a diagonal matrix with $b_j = 1/(\mu - \xi_j)$, ξ_j is the j -th eigenvalue of Q_0 , and Q_0 and \mathbf{r} are given in (29).

We verify (30) in simulations shown in Table II, which shows that our results are indeed very accurate. While (30) provides values ϕ that are smaller than isolation probability ϕ_{out} of the out-degree model [13] by several orders of magnitude, it is still unclear what impact in-degree has on the probability that a user gets isolated as its age increases and how large the improvement ratio ϕ_{out}/ϕ is. We study these issues below.

C. 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 $\{W(t)\}$, i.e., $T_{out} = \inf\{t > 0 : W(t) = 0 | W(0) = k\}$. Analysis in [14] shows that $\{W(t)\}$ 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 the proof of Theorem 11 to compare the distribution of isolation times in the joint in/out degree model with that studied in [14]. We plot the exact distributions of both T_{out} and T as functions of user age in Fig. 10. 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 curve of $P(T < t)$ becomes almost flat as time t increases beyond 0.5 hours showing that users with lifetimes in the range $[0.5, 200]$ hours exhibit almost the same 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 $X(t)$ becomes saturated and can no longer keep up with the increased possibility of isolation.

D. Exponential Lifetimes (Asymptotic Model)

Although (30) provides exact results for ϕ , it relies on numerical matrix algebra. Our next task is to combine the earlier




Fig. 10. The 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 $k = 6$.

derived in-degree model (26) with the out-degree isolation result in [13] to obtain a simple closed-form solution to ϕ when the mean search delay $E[S] \rightarrow 0$.

Theorem 12: For $L \sim \exp(\mu)$ and $S \sim \exp(\sigma)$, isolation probability converges to the following as $E[S] \rightarrow 0$:

$$\phi = \frac{1 - e^{-2k}}{2k} \phi_{out}, \quad (31)$$

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

It can be seen from (31) that by considering both in-degree and out-degree, the probability of node isolation is reduced by a factor of approximately $2k$ for non-trivial k . The reason for this relatively small improvement is that only a handful of users benefit from the in-degree in their isolation resilience since the majority of users depart very quickly and are unable to accumulate any 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 huge benefits from the in-degree process and are thus allowed to continue providing services to others with much higher probability than possible with just the out-degree.*

To complete this section, Table III shows the relative approximation error of (31) and indeed confirms its asymptotic accuracy. Note that for large search delays, (31) provides an upper bound on the isolation probability, where the ratio ϕ_{out}/ϕ is 3-10 times larger than the $2k$ suggested by (31). For instance, for fixed $E[L] = 0.5$ hours and $k = 6$, ratio ϕ_{out}/ϕ is 39 when $E[S] = 2$ minutes and 120 when $E[S] = 6$ minutes.

VI. CONCLUSION

This paper introduced a simple model of churn and developed numerous closed-form results describing the behavior of users including their joint and residual lifetime distributions, evolution of system size, transient in-degree distribution, and isolation probability under the joint in/out degree model. Future work involves modeling of non-uniform neighbor selection and churn of traditional DHTs.

REFERENCES

[1] R. Albert, H. Jeong, and A. Barabási, "Error and Attack Tolerance of Complex Networks," *Nature*, vol. 406, 2000.
 [2] J. Aspnes, Z. Diamadi, and G. Shah, "Fault Tolerant Routing in Peer to Peer Systems," in *Proc. ACM PODC*, Jul. 2002.

TABLE III
CONVERGENCE OF (31) TO (30) FOR $E[L] = 0.5$ HOURS AND $k = 6$

$E[S]$	Exact model (30)	Approx. model (31)	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

[3] S. Bernstein, "Sur Les Fonctions Absolument Monotones," *Acta Math*, vol. 51, 1928.
 [4] R. Bhagwan, S. Savage, and G. Voelker, "Understanding Availability," in *Proc. IPTPS*, Feb. 2003.
 [5] F. Bustamante and Y. Qiao, "Friendships that Last: Peer Lifespan and its Role in P2P Protocols," in *Proc. Intl. Workshop on Web Caching and Distribution*, Sept. 2003.
 [6] B. Chun, B. Zhao, and J. Kubiatowicz, "Impact of Neighbor Selection on Performance and Resilience of Structured P2P Networks," in *Proc. IPTPS*, Feb. 2005.
 [7] A. Feldmann and W. Whitt, "Fitting Mixtures of Exponentials to Long-Tailed Distributions to Analyze Network Performance Models," *Performance Evaluation*, vol. 31, no. 3-4, pp. 245-279, Jan. 1998.
 [8] 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.
 [9] K. Gummadi, R. Dunn, S. Saroiu, S. Gribble, H. Levy, and J. Zahorjan, "Measurement, Modeling, and Analysis of a Peer-to-Peer File-Sharing Workload," in *Proc. ACM SOSP*, October 2003.
 [10] M. Kaashoek and D. Karger, "Koorde: A Simple Degree-Optimal Distributed Hash Table," in *Proc. IPTPS*, Feb. 2003.
 [11] S. Krishnamurthy, S. El-Ansari, E. Aurell, and S. Haridi, "A Statistical Theory of Chord under Churn," in *Proc. IPTPS*, Feb. 2005.
 [12] S. Lam and H. Liu, "Failure Recovery for Structured P2P Networks: Protocol Design and Performance Evaluation," in *Proc. ACM SIGMETRICS*, Jun. 2004.
 [13] D. Leonard, V. Rai, and D. Loguinov, "On Lifetime-Based Node Failure and Stochastic Resilience of Decentralized Peer-to-Peer Networks," in *Proc. ACM SIGMETRICS*, Jun. 2005.
 [14] D. Leonard, Z. Yao, X. Wang, and D. Loguinov, "On Static and Dynamic Partitioning Behavior of Large-Scale Networks," in *Proc. IEEE ICNP*, Nov. 2005.
 [15] J. Li, J. Stribling, T. Gil, R. Morris, and M. Kaashoek, "Comparing the Performance of Distributed Hash Tables under Churn," in *Proc. IPTPS*, Feb. 2004.
 [16] J. Li, J. Stribling, R. Morris, and M. Kaashoek, "Bandwidth-Efficient Management of DHT Routing Tables," in *Proc. USENIX NSDI*, May 2005.
 [17] D. Liben-Nowell, H. Balakrishnan, and D. Karger, "Analysis of the Evolution of the Peer-to-Peer Systems," in *Proc. ACM PODC*, Jul. 2002.
 [18] G. Manku, M. Naor, and U. Weider, "Know thy Neighbor's Neighbor: the Power of Lookahead in Randomized P2P Networks," in *Proc. ACM STOC*, Jun. 2004.
 [19] G. Pandurangan, P. Raghavan, and E. Upfal, "Building Low-Diameter Peer-to-Peer Networks," *IEEE J. Sel. Areas Comm.*, vol. 21, no. 6, Aug. 2003.
 [20] S. Resnick, *Adventures in Stochastic Processes*. Birkhäuser, 2002.
 [21] S. Rhea, D. Geels, T. Roscoe, and J. Kubiatowicz, "Handling Churn in a DHT," in *Proc. USENIX Annual Technical Conference*, Jun. 2004.
 [22] S. Saroiu, P. Gummadi, and S. Gribble, "A Measurement study of Peer-to-Peer File Sharing Systems," in *Proc. MMCN*, 2002.
 [23] I. Stoica, R. Morris, D. Karger, M. Kaashoek, and H. Balakrishnan, "Chord: A Scalable Peer-to-Peer lookup Service for Internet Applications," in *Proc. ACM SIGCOMM*, Aug. 2001.
 [24] D. Stutzbach and R. Rejaie, "Understanding Churn in Peer-to-Peer Networks," in *Proc. ACM IMC*, Oct. 2006.
 [25] R. Wolff, *Stochastic Modeling and the Theory of Queues*. Prentice Hall, 1989.
 [26] Z. Yao, D. Leonard, X. Wang, and D. Loguinov, "Modeling Heterogeneous User Churn and Local Resilience of Unstructured P2P Networks," Texas A&M University, Tech. Rep. 2006-8-1, Aug. 2006. [Online]. Available: <http://irl.cs.tamu.edu/publications/>.
 [27] M. Zhong, K. Shen, and J. Seiferas, "Non-Uniform Random Membership Management in Peer-to-Peer Networks," in *Proc. IEEE INFOCOM*, Mar. 2005.