

# On Node Isolation under Churn in Unstructured P2P Networks with Heavy-Tailed Lifetimes

Zhongmei Yao, Xiaoming Wang, Derek Leonard, and Dmitri Loguinov\*

Department of Computer Science, Texas A&M University

College Station, TX 77843 USA

{mayyao, xmwang, dleonard, dmitri}@cs.tamu.edu

**Abstract**—Previous analytical studies [12], [18] of unstructured P2P resilience have assumed exponential user lifetimes and only considered age-independent neighbor replacement. In this paper, we overcome these limitations by introducing a general node-isolation model for heavy-tailed user lifetimes and arbitrary neighbor-selection algorithms. Using this model, we analyze two age-biased neighbor-selection strategies and show that they significantly improve the residual lifetimes of chosen users, which dramatically reduces the probability of user isolation and graph partitioning compared to uniform selection of neighbors. In fact, the second strategy based on random walks on age-weighted graphs demonstrates that for lifetimes with infinite variance, the system monotonically *increases* its resilience as its age and size grow. Specifically, we show that the probability of isolation converges to zero as these two metrics tend to infinity. We finish the paper with simulations in finite-size graphs that demonstrate the effect of this result in practice.

## I. INTRODUCTION

Resilience of P2P networks under random user arrival and departure (i.e., churn) has recently become an active research area [9], [10], [11], [12], [13], [14], [15], [22]. One of the primary metrics of resilience is *graph disconnection* during which a P2P network partitions into several non-trivial sub-graphs and starts to offer limited service to its users. However, as shown in [13], most partitioning events in well-connected P2P networks are single-node isolations, which occur when the immediate neighbors of a node  $v$  fail before  $v$  is able to detect their departure and then replace them with other alive users. For such networks, node isolation analysis has become the primary method for quantifying network resilience in the presence of user churn.

Traditional analysis of node isolation [12], [13] focuses on the effect of average neighbor-replacement delay  $E[S]$ , average user lifetime  $E[L]$ , and fixed out-degree  $k$  on the resilience of the system. These results show that probability  $\phi$  with which each arriving user is isolated from the system during its lifetime is proportional to  $k\rho(1+\rho)^{-k}$ , where  $\rho = E[L]/E[S]$ . While this result is asymptotically exact under *exponential* user lifetimes and *uniform* neighbor selection, it remains to be investigated whether stronger results can be obtained for heavy-tailed lifetimes and/or non-uniform neighbor selection. We study these questions below.

### A. Paper Structure and Contributions

The main focus of this paper is to understand node isolation in the context of unstructured networks (such as Gnutella)

where neighbor selection is not constrained by fixed rules. As in [12], we assume that each arriving user is assigned a random lifetime  $L$  drawn from some distribution  $F(x)$  and is given  $k$  initial neighbors randomly selected from the system. The user then constantly monitors and replaces its neighbors to avoid isolation from the rest of the system. Random replacement delay  $S$  is needed to detect the failure of an old neighbor and find a new one from among the remaining alive users. Unlike [12], we allow  $L$  to come from any distribution (e.g., Pareto, Weibull), as long as  $E[L] < \infty$ , and neighbor selection to be arbitrary, as long as the stationary distribution  $H(x)$  of residual lifetimes  $R$  of selected neighbors is known.

We first build a generic isolation model that allows computation of  $\phi$  with arbitrary accuracy for any completely monotone density function of residual lifetimes  $R$ . This result is achieved by replacing  $H(x)$  with a hyper-exponential distribution, which can be performed with any accuracy, and then solving the resulting Markov chain for the probability of absorption into the isolation state before the user decides to leave the system. While this model only admits a numerical solution through matrix manipulation, it allows very accurate computation of  $\phi$  for very heavy-tailed cases when the exponential upper bound  $\phi \leq k\rho(1+\rho)^{-k}$  [12] is rather loose. The model is also necessary for studying isolation behavior of the various neighbor-selection strategies examined in later parts of the paper where simulations are impractical or impossible due to the small values of  $\phi$ .

The second part of the paper verifies the model of  $\phi$  under uniform neighbor replacement and analyzes its performance for very heavy-tailed lifetimes (i.e.,  $Var[L] = \infty$ ). We show that as the age  $\mathcal{T}$  of the system becomes infinite and the Pareto shape parameter  $\alpha$  approaches 1, the isolation probability decays to zero proportionally to  $(\alpha - 1)^k$ , which holds for *any* number of neighbors  $k \geq 1$  and *any* search delay  $S$ , implying that such systems may achieve arbitrary resilience without replacing any neighbors. In practice, however,  $\alpha$  is a fixed number bounded away from 1 (common studies [1] suggest  $\alpha = 1.06$ ) and  $\mathcal{T}$  is finite, which cannot guarantee high levels of robustness without neighbor replacement.

As an improvement over the uniform case, we next study the so-called *max-age* neighbor selection [1], [11], [25], in which a user samples  $m$  uniformly random peers per link it creates and selects the one with the largest current age to be its neighbor. We show that larger values of  $m$  lead to stochastically larger  $R$  and improve the expected remaining lifetimes of found neighbors by a factor proportional to  $m^{1/(\alpha-1)}$  for non-trivial  $m$ . For example,  $\alpha = 3$  increases  $E[R]$  as  $\sqrt{m}$ ,  $\alpha \approx 2$

\*Supported by NSF grants CCR-0306246, ANI-0312461, CNS-0434940, and CNS-0519442.

increases  $E[R]$  linearly in  $m$ , and  $\alpha < 2$  results in  $E[R] = \infty$  regardless of  $m$  as long as  $\mathcal{T} = \infty$ . We do not obtain a closed-form factor of reduction for  $\phi$  compared to the purely uniform case, but note that it is a certain monotonic function of  $m$ . This does not change, however, the qualitative behavior of  $\phi$  under the no-replacement policy and still requires  $\alpha \rightarrow 1$  to achieve  $\phi \rightarrow 0$  for any fixed  $m$ .

While the max-age approach is viable and very effective in general, it relies on the system's ability to sample  $m$  peers uniformly randomly per created link. This can be accomplished using Metropolis-style random walks [31]; however, the overhead may not scale well for large  $m$ . To build a distributed solution that requires only *one* sample per link, the last part of the paper proposes a novel technique based on random walks over directed graphs, in which the weight of in-degree edges at each node is kept proportional to the age of the corresponding user. Under these conditions, we derive a model for the residual distribution  $H(x)$  and show that isolation probability  $\phi$  converges to 0 for any  $1 < \alpha \leq 2$  as system size  $n \rightarrow \infty$  and age  $\mathcal{T} \rightarrow \infty$ , which holds for any number of neighbors  $k \geq 1$  and any search delay  $S$ . Compared to the uniform and max-age cases, this is a much stronger result that shows that with just  $k = 1$  neighbor and no replacement of failing neighbors, large P2P systems with  $\alpha \leq 2$  can guarantee arbitrarily low values of  $\phi$ . We finish the paper by studying in simulations the approach rate of  $\phi$  to 0 and its effect in practice.

The rest of the paper is organized as follows. Section II introduces a generic isolation model for non-exponential lifetimes and verifies it in simulations. In section III, we formalize max-age neighbor replacement and derive the corresponding residual lifetime distributions. Section IV discusses random walks on age-weighted graphs and studies the asymptotic effect of  $1 < \alpha \leq 2$  on isolation probability. Section V discusses related work and section VI concludes the paper.

## II. GENERAL NODE ISOLATION MODEL

In this section, we build a model for the probability  $\phi$  that a node  $v$  becomes isolated due to all of its neighbors simultaneously reaching the failed state during its lifetime. While closed-form derivation of  $\phi$  for systems with non-exponential user lifetimes is difficult, certain cases identified below can be solved with arbitrary accuracy by replacing residual lifetimes and search delays with their hyper-exponential equivalents.

For the churn model, we adopt the conventions of [12], but relax the assumptions of uniform neighbor selection and exponential lifetimes. Denote by  $W(t)$  the out-degree of a node  $v$  at time  $t$ . The rest of this section deals with constructing a continuous-time Markov chain for the degree process  $\{W(t)\}$ , which allows us to obtain an accurate model of  $\phi$ .

### A. Hyper-Exponential Approximation

Recall that the hyper-exponential distribution  $H_m$  is a mixture of  $m$  exponential random variables with probability density function (PDF) in the form of [28]:

$$f_H(x) = \sum_{j=1}^m p_j \mu_j e^{-\mu_j x}, \quad (1)$$

where  $\mu_j, p_j \geq 0$  for all  $j$  and  $\sum_{j=1}^m p_j = 1$ . The above distribution can be interpreted as generating each exponential random variable  $\exp(\mu_j)$  with probability  $p_j$ . It is well-known [5] that any *completely monotone* density function  $f(x)$  can be represented with any desired accuracy using (1), i.e.,  $f_H(x) \rightarrow f(x)$  as  $m \rightarrow \infty$ . In the analysis below, we leverage this property of hyper-exponentials and the fact that Pareto and Weibull residual PDFs are completely monotone. While some of the prior literature [5] has used as many as 14 exponentials to approximate Pareto  $f(x)$ , our analysis suggests that as few as 3 are usually sufficient for achieving very accurate results on  $\phi$  (see below).

Before we proceed with the derivations, it is useful to visualize the meaning of hyper-exponential distributions in our lifetime model. Assume that there are  $r$  different types of neighbors, where residual lifetimes of peers of type  $1 \leq i \leq r$  are exponentially distributed with rate  $\mu_i$ . When  $v$  requires a new neighbor, it selects a node of type  $i$  with probability  $p_i$ . Similarly, assume that there are  $s$  types of searches that can be currently in progress. A search of type  $1 \leq j \leq s$  is instantiated by  $v$  with probability  $q_j$  and has duration exponentially distributed with rate  $\lambda_j$ . As long as neighbor residual lifetimes  $R$  and search delays  $S$  can be reduced to the hyper-exponential distribution, the resulting process  $\{W(t)\}$  can be viewed as a homogenous continuous-time Markov chain as we show next.

Denote each state  $u$  of  $\{W(t)\}$  for any given user  $v$  by:

$$u = (x_1, \dots, x_r, y_1, \dots, y_s), \quad (2)$$

where  $x_i$  is the number of  $v$ 's neighbors of type  $i$ ,  $y_j$  is the number of searches in progress of type  $j$ ,  $0 \leq x_i \leq k$ ,  $0 \leq y_j \leq k$ , and  $\sum_{i=1}^r x_i + \sum_{j=1}^s y_j = k$ . Also notice that out-degree  $W(t)$  can be represented as  $\sum_{i=1}^r x_i$ . The next theorem specifies that  $\{W(t)\}$  is a Markov chain and derives its transition rate matrix.

*Theorem 1:* Given that the density function of residual lifetimes  $f_R(t) = \sum_{j=1}^r p_j \mu_j e^{-\mu_j t}$  and the density function of search times  $f_S(t) = \sum_{j=1}^s q_j \lambda_j e^{-\lambda_j t}$ ,  $\{W(t)\}$  is a homogeneous continuous-time Markov chain with a transition rate matrix  $Q = (q_{uu'})$ , where

$$q_{uu'} = \begin{cases} q_j x_i \mu_i & (x_i, y_j) \rightarrow (x_i - 1, y_j + 1) \\ p_i y_j \lambda_j & (x_i, y_j) \rightarrow (x_i + 1, y_j - 1) \\ -\Lambda_u & u' = u \\ 0 & \text{otherwise} \end{cases}, \quad (3)$$

$u$  and  $u'$  represent any suitable states in the form of (2) that satisfy transition requirements on the right side of (3), and  $\Lambda_u = \sum_{i=1}^r x_i \mu_i + \sum_{j=1}^s y_j \lambda_j$ .

The next step is to obtain the initial state distribution of  $\{W(t)\}$  and derive the PDF of the first-hitting time on state  $W(t) = 0$  based on the transition rate matrix  $Q$  in (3). For small values of  $k$ , the matrix can be easily represented in memory and manipulated in software packages such as Matlab. For example, when  $r = s = 3$  commonly used in this work, the size of  $Q$  is  $252 \times 252$  for  $k = 5$  and  $792 \times 792$  for  $k = 7$ .

The initial state distribution  $\pi(0)$  is in form of:

$$\pi(0) = (\pi_{(x_1, \dots, x_r, y_1, \dots, y_s)}(0)), \quad (4)$$

where each entry in the vector represents the probability that the chain starts in state  $(x_1, \dots, x_r, y_1, \dots, y_s)$  for all possible permutations of variables  $x_i$  and  $y_j$ . Note, however, that the only valid starting states are those in which the number of alive neighbors  $\sum_{i=1}^r x_i$  is exactly  $k$  and the number of searches in progress  $\sum_{j=1}^s y_j$  is zero.

After rather straightforward manipulations,  $\pi(0)$  can be obtained as follows.

*Lemma 1:* Valid starting states have initial probabilities:

$$\pi_{(x_1, \dots, x_r, 0, \dots, 0)}(0) = \prod_{i=1}^r \binom{k - \sum_{j=1}^{i-1} x_j}{x_i} p_i^{x_i}, \quad (5)$$

and all other states have initial probability 0.

Model (5) is verified in simulations that are not shown here for brevity. Armed with this result, we next focus our attention on deriving  $\phi$ .

### B. Isolation Probability

Denote by  $T$  the first-hitting time of  $\{W(t)\}$  onto state 0:

$$T = \inf(t > 0 : W(t) = 0 | W(0) = k). \quad (6)$$

The goal of this subsection is to derive the distribution of  $T$ , based on which we obtain  $\phi = P(T < L)$ .

In general, it is convenient to treat  $\{W(t)\}$  as an absorbing Markov chain in order to derive the PDF of  $T$ . To this end, let  $\Omega$  denote the set of all valid states (i.e., in the form of (2) and satisfying all constraints following the equation) and

$$E = \left\{ (0, \dots, 0, y_1, \dots, y_s) : \sum_{j=1}^s y_j = k \right\} \quad (7)$$

be the set of all absorbing states. Then, for each non-absorbing state  $u \in \Omega \setminus E$ , its transition rate to  $E$  is given by:

$$q_{uE} = \sum_{u' \in E} q_{uu'}, \quad (8)$$

where  $q_{uu'}$  is the cell of matrix  $Q$  corresponding to transitions from state  $u$  to  $u'$ . We can then write  $Q$  in canonical form as:

$$Q = \begin{pmatrix} 0 & 0 \\ \mathbf{r} & Q_0 \end{pmatrix}, \quad (9)$$

where  $\mathbf{r} = (q_{uE})^T$  for  $u \notin E$  is a column vector representing the transition rates to the absorbing set  $E$  and  $Q_0$  is the rate matrix obtained by removing the rows and columns corresponding to states in  $E$  from  $Q$ . The following lemma shows that the PDF of  $T$  is fully determined by  $\pi(0)$  and  $Q$ .

*Lemma 2:* For residual lifetimes and search delays with hyper-exponential distributions, the PDF of  $T$  is given by:

$$f_T(t) = \pi(0)VD(t)V^{-1}\mathbf{r}, \quad (10)$$

where  $\pi(0)$  is the initial state distribution in (5),  $V$  is a matrix of eigenvectors of  $Q_0$ ,  $D(t) = \text{diag}(e^{\xi_j t})$  is a diagonal matrix,  $\xi_j \leq 0$  is the  $j$ -th eigenvalue of  $Q_0$ , and  $Q_0$  and  $\mathbf{r}$  are in (9).

With Lemma 2 in hand, integrating  $f_T(t)$  using the distribution of user lifetimes immediately leads to the following theorem.

*Theorem 2:* For hyper-exponential residual lifetimes and search delays, the probability of isolation is:

$$\phi = \pi(0)VBV^{-1}\mathbf{r}, \quad (11)$$

where  $B = \text{diag}(b_j)$  is a diagonal matrix with:

$$b_j = \int_0^\infty (1 - F(t))e^{\xi_j t} dt, \quad (12)$$

$F(t)$  is the CDF of user lifetimes, and all other parameters are the same as in Lemma 2.

Using rate matrix  $Q_0$ , vector  $\mathbf{r}$ , and (11)-(12), the solution to node isolation probability  $\phi$  can be easily computed using numerical packages such as Matlab. We perform this task next.

### C. Verification of Isolation Model

We examine the accuracy of (11)-(12) using the simplest example of uniform selection. We first explore the exponential case for comparison purposes and then derive the same metric for Pareto lifetimes. For exponential lifetimes, the next lemma immediately follows upon recalling that neighbor residual lifetimes  $R$  are also exponentially distributed with  $m = 1$  in (1) due to the memoryless property of the distribution.

*Lemma 3:* For exponential  $L \sim \exp(\mu)$  and search delays with a hyper-exponential density  $f_S(x)$ , (12) is simply:

$$b_j = 1/(\mu - \xi_j). \quad (13)$$

Our next theorem derives  $\phi$  for Pareto lifetimes with CDF  $1 - (1 + x/\beta)^{-\alpha}$  for  $\alpha > 1$ ,  $x \geq 0$ . Denote by  $R$  the residual lifetime of a uniformly random user in the system. Assuming a sufficiently large system age  $\mathcal{T}$ , it follows from [12] that the CDF of  $R$  under uniform selection is given by:

$$P(R < x) = 1 - \left(1 + \frac{x}{\beta}\right)^{-(\alpha-1)}. \quad (14)$$

It is clear from (14) that the PDF of Pareto residuals is completely monotone and thus can be fitted with its hyper-exponential equivalent. Invoking Theorem 2, we immediately obtain the following.

*Lemma 4:* For Pareto  $L \sim 1 - (1 + x/\beta)^{-\alpha}$  and hyper-exponential search delays, (12) becomes:

$$b_j = \beta e^{-\xi_j \beta} E_\alpha(-\xi_j \beta), \quad (15)$$

where  $E_\alpha(x) = \int_1^\infty e^{-xu} u^{-\alpha} du$  is the generalized exponential integral.

Next, we run simulations to verify these results. Simulations are performed over different distributions of search times on a graph with  $n = 1,000$  nodes,  $k = 7$ , and mean lifetime  $E[L] = 0.5$  hours (additional simulations produce similar results and are omitted for brevity). The first search time distribution is Pareto with  $\alpha = 3$  and  $\beta = E[S](\alpha - 1)$  to keep the mean equal to  $E[S]$ . The second distribution is Weibull with CDF  $1 - e^{-(x/a)^c}$  and mean  $E[S] = a\Gamma(1 + 1/c)$ . The third is exponential with rate  $1/E[S]$ . To compute the model, Pareto residual lifetime  $R$  is fitted with a hyper-exponential mixture model (1) using  $r = 3$  and each non-exponential search distribution is fitted with model (1) using  $s = 3$ .

Exponential and Pareto models of  $\phi$  are compared to simulation results in Table I. Notice in the table that both (13)

TABLE I  
COMPARISON OF MODEL  $\phi$  TO SIMULATIONS UNDER UNIFORM SELECTION WITH  $E[L] = 0.5$  HOURS AND  $k = 7$

$E[S]$ hours	Pareto $L$ with $\alpha = 3$						Exponential $L$	
	Pareto $S$ with $\alpha = 3$		Weibull $S$ with $c = 0.7$		Exponential $S$		Pareto $S$ with $\alpha = 3$	
	Simulations	Model (15)	Simulations	Model (15)	Simulations	Model (15)	Simulations	Model (13)
.001		$1.11 \times 10^{-16}$		$1.12 \times 10^{-16}$		$1.12 \times 10^{-16}$		$4.40 \times 10^{-16}$
.01		$8.49 \times 10^{-11}$		$8.45 \times 10^{-11}$		$9.05 \times 10^{-11}$		$3.70 \times 10^{-10}$
.05	$4.56 \times 10^{-7}$	$4.49 \times 10^{-7}$	$4.93 \times 10^{-7}$	$4.96 \times 10^{-7}$	$6.27 \times 10^{-7}$	$6.28 \times 10^{-7}$	$2.31 \times 10^{-6}$	$2.31 \times 10^{-6}$
.1	$1.13 \times 10^{-5}$	$1.14 \times 10^{-5}$	$1.21 \times 10^{-5}$	$1.25 \times 10^{-5}$	$1.75 \times 10^{-5}$	$1.74 \times 10^{-5}$	$6.01 \times 10^{-5}$	$6.04 \times 10^{-5}$
.4	$1.64 \times 10^{-3}$	$1.64 \times 10^{-3}$	$1.60 \times 10^{-3}$	$1.58 \times 10^{-3}$	$2.57 \times 10^{-3}$	$2.59 \times 10^{-3}$	$6.80 \times 10^{-3}$	$6.78 \times 10^{-3}$
.8	$7.78 \times 10^{-3}$	$7.78 \times 10^{-3}$	$7.14 \times 10^{-3}$	$7.16 \times 10^{-3}$	$1.12 \times 10^{-2}$	$1.12 \times 10^{-2}$	$2.56 \times 10^{-2}$	$2.56 \times 10^{-2}$




Fig. 1. Impact of shape parameter  $\alpha$  on model  $\phi$  under uniform selection, Pareto lifetimes,  $E[L] = 0.5$  hours,  $\beta = (\alpha - 1)E[L]$ , exponential search delays, and  $k = 7$ .

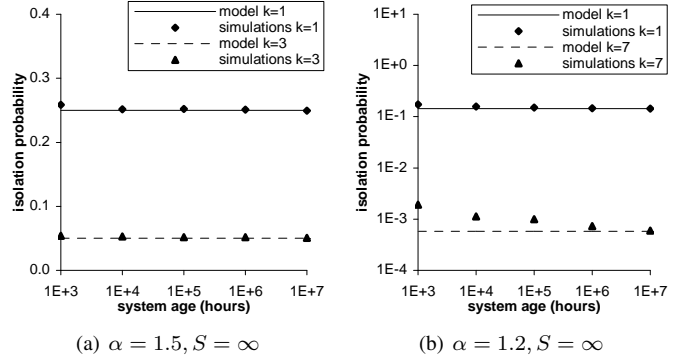


Fig. 2. Convergence of simulation results to model  $\phi$  in (16) as system age  $T \rightarrow \infty$  under uniform selection, no neighbor replacement, and Pareto lifetimes with  $\beta = (\alpha - 1)E[L]$  in a graph with  $n = 1,000$  nodes.

and (15) are indeed very accurate for all examined search and lifetime distributions. The table also confirms that as  $E[S] \rightarrow 0$ , metric  $\phi$  becomes insensitive to the distribution of  $S$ , which was earlier observed in [12] but never verified.

To understand the influence of tail weight of the lifetime distribution  $F(x)$  on isolation, we use (15) to compute  $\phi$  for several values of shape parameter  $\alpha$  and keep  $\beta = (\alpha - 1)E[L]$  to ensure that the mean lifetime  $E[L]$  remains fixed. The result is shown in Fig. 1 for two values of  $E[S]$  and  $k = 7$ . Notice in both sub-figures that the relationship between  $\phi$  and  $\alpha$  is similar and that  $\phi$  appears to be approximately a logarithmic function of  $\alpha$  for  $\alpha \leq 21$ , confirming that the more heavy-tailed the lifetime distribution, the smaller  $\phi$ .

#### D. Necessity of Neighbor Replacement

Fig. 1 suggests that  $\phi$  tends to 0 as  $\alpha$  approaches 1 from above, but it is not clear at what rate this convergence takes place and whether this is indeed true. Furthermore, since  $E[R] = \infty$  for  $\alpha \leq 2$ , a natural question arises about whether a finite system of  $n$  users and finite age  $T$  can in fact exhibit infinite expected residuals or  $\phi = 0$  when  $\alpha = 1$ . We answer these questions next and show that condition  $\alpha \rightarrow 1$  indeed guarantees  $\phi \rightarrow 0$  even in cases when no replacement of failed neighbors is performed; however, it requires that the system be *in equilibrium*<sup>1</sup> by the time it is observed by an arriving user.

<sup>1</sup>The first renewal cycle of each user must be drawn from its residual distribution or system age  $T$  be infinite. See [28, page 65] for a definition.

**Theorem 3:** For an equilibrium system, Pareto lifetimes with  $\alpha > 1$ , and infinitely large search delays (i.e.,  $S = \infty$ ), the isolation probability is:

$$\phi = \frac{k!}{(\gamma + 1) \times \dots \times (\gamma + k)}, \quad (16)$$

where  $\gamma = \alpha/(\alpha - 1)$ . For fixed  $k$  and  $\alpha \rightarrow 1$  (i.e.,  $\gamma \rightarrow \infty$ ), (16) converges to zero as  $\Theta(\gamma^{-k})$ .

This result is very interesting since most prior work [12] does not consider  $\alpha \leq 2$  as such cases result in infinite expected residual lifetimes, which cannot be observed in any finite system. However, if the age of the system tends to infinity, i.e.,  $T \rightarrow \infty$ , or the first lifetime of each user is drawn from the residual distribution (14), the asymptotic bound in (16) is actually achievable. In such cases, as  $\alpha$  tends to 1, the isolation probability will decay to zero proportionally to  $(\alpha - 1)^k$  as given by Theorem 3 and the system will attain any desired level of resilience without replacing neighbors. On the other hand, for  $\alpha$  sufficiently larger than 2 studied in prior work [12], age  $T$  must simply exceed the convergence time to equilibrium of the underlying user-lifetime renewal process, which usually happens very quickly.

Fig. 2 shows simulation results of  $\phi$  with  $S = \infty$  and two values of very heavy-tailed  $\alpha$ . Notice in Fig 2(a) that for  $\alpha = 1.5$ , simulation results converge to model  $\phi$  before system age reaches  $10^4$  hours (i.e., 1.14 years). However, as  $\alpha$  reduces to 1.2, the convergence takes a much longer time as shown in Fig 2(b), where simulations approach the model when system age grows to more than  $T = 10^6$  hours = 114 years.

The above analysis shows that the asymptotic result  $\phi \rightarrow 0$  as  $\alpha \rightarrow 1$  is not readily achievable in finite P2P systems. Furthermore, recent measurement studies of user lifetimes suggest that P2P networks exhibit  $\alpha$  that is bounded away from 1 (e.g.,  $\alpha = 1.06$  in [1]). Hence, most current P2P systems are not likely to satisfy the condition for  $\phi \rightarrow 0$  under uniform selection and thus need to utilize either a large number of neighbors  $k$  or perform dynamic replacement of dead links with  $E[S] \ll \infty$ .

### E. Discussion

While the general form of  $\phi$  in the exact model (11) is very complex, a simple qualitative rule of increasing resilience (i.e., reducing  $\phi$ ) can be formulated based on the properties of residual lifetimes selected by the users of a P2P system. Notice that for a fixed lifetime distribution  $F(x)$ , higher resilience is achieved by selecting neighbors that exhibit larger (in some sense) remaining lifetimes. Thus, given two strategies  $\mathcal{S}_1$  and  $\mathcal{S}_2$  for selecting neighbors, the strategy that obtains a neighbor with a larger residual lifetime during *every* replacement instance  $\tau$  guarantees a lower isolation probability since the chosen neighbors survive longer and increase the chance that the current user will depart before becoming isolated. Since comparison of residual lifetimes of obtained neighbors in  $\mathcal{S}_1$  and  $\mathcal{S}_2$  can be performed only in the *probabilistic* sense, the above discussion can be formalized as following: *strategies that produce stochastically larger<sup>2</sup> distributions of residuals guarantee lower isolation frequency and higher resilience.*

Note, however, that future residual lifetimes of sampled peers are usually not available in practice. Instead, assuming that  $F(x)$  is not memoryless (i.e., non-exponential), current user age  $A$  may be used as a robust predictor of  $R$ . To understand this correlation for Pareto  $F(x)$ , consider the probability that a peer's remaining lifetime is larger than  $y \geq 0$  given that its current age  $A$  is  $x \geq 0$ :

$$P(R > y | A = x) = \left(1 + \frac{y}{\beta + x}\right)^{-\alpha}. \quad (17)$$

Observe that the above conditional probability is a monotonically increasing function of age, i.e., the larger  $x$ , the more likely a node is to survive at least  $y$  time units in the future. This implies that *users with larger age demonstrate stochastically larger residual lifetimes  $R$ .*

This result can be generalized to all heavy-tailed distributions (defined in terms of conditional mean exceedance [8] or tail-decay rate [27], e.g., Pareto, Weibull, and Cauchy), in which the expected remaining lifetime increases and  $R$  becomes stochastically larger with age. In contrast, light-tailed distributions (e.g., uniform and Gaussian), exhibit expected residual lifetimes that are decreasing functions of age. Finally, for the exponential distribution, age does not affect residual lifetimes and hence does not provide any useful information for neighbor selection.

Armed with these observations and prior measurement results that demonstrate heavy-tailed user lifetimes in real

<sup>2</sup>Variable  $X$  is stochastically larger than  $Y$  if  $P(X > x) \geq P(Y > x)$  for all  $x \in \mathbb{R}$  [28].

P2P systems [1], the rest of the paper explores two simple neighbor-selection methods that rely on age of existing peers to increase network resilience.

## III. MAX-AGE SELECTION

Recall that under uniform selection, each alive user is chosen by an arriving node  $v$  with the same probability. To prevent  $v$  from connecting to weak neighbors that are about to depart (i.e., users with short remaining lifetimes), this section leverages the heavy-tailed nature of the lifetime distribution  $F(x)$  and models the *max-age* neighbor-selection strategy proposed in [1], [11], [25]. In this approach, a joining node  $v$  uniformly randomly selects  $m$  alive users from the system and chooses the user with the maximal age. It then repeats this procedure  $k$  times to obtain its  $k$  initial neighbors. The same process is executed every time a dead link is detected.

In what follows in this section, we first analyze the distribution of residuals obtained by the max-age method and then discuss the corresponding isolation probability  $\phi$ .

### A. Residual Lifetime Distribution

Denote by  $\Omega_m$  the set of  $m$  candidate nodes, by  $U_m$  the residual lifetime of the max-age user in  $\Omega_m$ , and by  $H^c(x) = P(U_m > x)$  the CCDF of random variable  $U_m$ . Then, we get:

$$H^c(x) = P\left(R_i > x | A_i = \max_{j \in \Omega_m} \{A_j\}\right), \quad (18)$$

where  $A_i$  is the current age of a user  $i$  in  $\Omega_m$  and  $R_i$  is its residual lifetime. Intuitively, (18) states that  $U_m$  equals  $R_i$  given that user  $i$  has the maximum age in  $\Omega_m$ . Next, from renewal process theory, the equilibrium age distribution of *existing users* in the system is [21]:

$$F_A(x) = P(A < x) = \frac{1}{E[L]} \int_0^x (1 - F(u)) du, \quad (19)$$

where  $E[L] < \infty$ . The following theorem shows that  $H^c(x)$  is fully determined by the number of sampled users, lifetime distribution  $F(x)$ , and age distribution  $F_A(x)$ .

**Theorem 4:** Given that a user's age is larger than that of  $m-1$  uniformly selected alive users in the system, its residual lifetime has the following CCDF:

$$H^c(x) = \frac{m}{E[L]} \int_0^\infty (1 - F(x+y)) F_A^{m-1}(y) dy, \quad (20)$$

where  $F_A(x)$  is given by (19).

Next, we use exponential lifetimes as an example to verify (20). Using  $F(x) = F_A(x) = 1 - e^{-\mu x}$ , (20) reduces to:

$$H^c(x) = m\mu \int_0^\infty e^{-\mu(x+y)} (1 - e^{-\mu y})^{m-1} dy = e^{-\mu x}. \quad (21)$$

Hence, it follows from (21) that for exponential lifetimes:

$$P(U_m > x) = P(L > x) = e^{-\mu x}, \text{ for any } m \geq 1, \quad (22)$$

which is consistent with the memoryless property of the exponential distribution. Substituting Pareto lifetimes into (20), we obtain:

$$H^c(x) = \frac{m}{E[L]} \int_0^\infty \left(1 + \frac{x+y}{\beta}\right)^{-\alpha} \left(1 - \left(1 + \frac{y}{\beta}\right)^{1-\alpha}\right)^{m-1} dy, \quad (23)$$




Fig. 3. Accuracy of models (23) and (28) for Pareto lifetimes with  $E[L] = 0.5$  hours and  $\alpha = 3$  in a graph with  $n = 5,000$  nodes.

where  $E[L] = \beta/(\alpha - 1)$ .

Although no closed-form solution for (23) exists in the general case, we next perform a self-check using  $m = 1$ . Note that for  $m = 1$ , (23) yields:

$$H^c(x) = \frac{\alpha - 1}{\beta} \int_0^\infty \left(1 + \frac{x+y}{\beta}\right)^{-\alpha} dy = \left(1 + \frac{x}{\beta}\right)^{1-\alpha}, \quad (24)$$

which indicates that  $P(U_1 > x) = P(R > x)$  (i.e., max-age selection with  $m = 1$  reduces to single-user uniform selection).

Our next result shows that  $U_m$  is stochastically larger than  $U_{m-1}$  for any heavy-tailed  $F(x)$  and any  $m \geq 2$ .

**Theorem 5:** For any distribution in which larger age implies stochastically larger residuals (i.e., function (17) is monotonically increasing in  $x$ ), the following holds:

$$P(U_m > x) \geq P(U_{m-1} > x), \quad x \geq 0, m \geq 2. \quad (25)$$

Simulation results in Fig. 3(a) show for  $m = 6$  that model (23) is very accurate and random variable  $U_6$  is indeed stochastically larger than  $R$  (simulations with other  $m$  and those confirming (25) are omitted for brevity). Next, we solve for the expectation of  $U_m$  in closed-form for Pareto lifetimes and show the effect of  $m$  on the average residual lifetimes of selected neighbors.

**Lemma 5:** For Pareto  $L \sim 1 - (1 + x/\beta)^{-\alpha}$ ,  $\alpha > 2$ , the expectation of  $U_m$  is given by:

$$E[U_m] = \frac{\beta m! \Gamma(\frac{\alpha-2}{\alpha-1})}{(m(\alpha-1) - 1) \Gamma(m - \frac{1}{\alpha-1})}, \quad m \geq 1, \quad (26)$$

where  $\Gamma(x)$  is the gamma function. For  $\alpha \leq 2$ , the expected residual lifetime converges to infinity as system age  $\mathcal{T}$  becomes large:

$$\lim_{\mathcal{T} \rightarrow \infty} E[U_m] = \infty, \quad m \geq 1. \quad (27)$$

To better understand the effect of  $m$  on the mean of  $U_m$ , we approximate  $E[U_m]$  as follows. Setting  $c = \Gamma(\frac{\alpha-2}{\alpha-1})$  and expanding the gamma function in the denominator, (26) for  $\alpha > 2$  yields:

$$E[U_m] \approx cE[L] \left(m + \frac{1}{\alpha}\right)^{1/(\alpha-1)}. \quad (28)$$

We next discuss several examples that use (28) with different  $\alpha$ . For Pareto lifetimes with  $E[L] = 0.5$  hours and




Fig. 4. Comparison of model  $\phi$  to simulations using the max-age selection strategy for Pareto lifetimes with  $E[L] = 0.5$  hours and  $\alpha = 3$ , exponential search times and  $k = 7$  in a graph with 5,000 nodes.

$\alpha = 3$ , it can be seen from (28) that  $E[U_m]$  follows the curve  $0.886(m + 0.33)^{0.5} \sim \sqrt{m}$  as  $m \rightarrow \infty$ . However, for smaller  $\alpha$ , a more aggressive increase in  $E[U_m]$  can be obtained. For  $\alpha \rightarrow 2$ ,  $E[U_m] \sim m$  is approximately linear, and for  $\alpha < 2$ ,  $E[U_m] = \infty$  for any  $m \geq 1$  (as before, the last results only holds conditioned on  $\mathcal{T} = \infty$ ). It is also apparent from (28) that as shape parameter  $\alpha$  tends to infinity, the impact of  $m$  on  $E[U_m]$  is weakened and  $E[U_m] \rightarrow E[L]$ , which confirms a well-known fact [12] that Pareto lifetimes with very large  $\alpha$  behave as exponential random variables.

Model (26) is confirmed to be exact using simulations not shown here due to limited space. Fig. 3(b) shows the accuracy of the match between  $E[U_m]$  predicted by the exact model (26) and that by the approximate model (28) for  $\alpha = 3$ . Additional examples with smaller  $\alpha$  are omitted for brevity.

## B. Isolation and Resilience

To obtain model  $\phi$ , we approximate the tail of  $U_m$  in (20) with its hyper-exponential equivalent in (1) and then compute  $\phi$  by applying Theorem 2 as in Section II-C. Fig. 4 shows  $\phi$  predicted by the model compared to simulations for Pareto lifetimes with  $E[L] = 0.5$  hours,  $k = 7$ , exponential search delays, and two values of  $m$ . As the figure illustrates, the derived result is very accurate and indeed shows inversely proportional dependency between the number of sampled users  $m$  and  $\phi$ . The influence of  $m$  on isolation probability for Pareto lifetimes is presented more clearly in Fig. 5. As the trendlines show,  $\phi$  is approximately a power-law function  $m^{-a}$  for each fixed  $E[S]$ , where exponent  $a$  is 2.4–5.7 in the figure. Thus, for  $\alpha = 3$ ,  $m = 10$  sampled users reduce  $\phi$  by a factor of 251 and  $m = 30$  by a factor of 3,508; however, for  $\alpha = 2$ ,  $m = 10$  drops  $\phi$  by a factor of 489,000 and  $m = 30$  by a factor of 2.5 billion. Interestingly, while  $E[U_m]$  may exhibit an unimpressive growth as a function of  $m$  (i.e., linear or slower), the corresponding  $\phi$  demonstrates much faster decay rate and almost always provides significant benefits as  $m$  increases.

In systems that do not replace neighbors and  $\alpha \rightarrow 1$ , the limiting isolation probability in (16) is reduced along the corresponding curve in Fig. 5, i.e., proportionally to  $m^{-a}$ . Thus, for any finite  $m$ , (16) does not qualitatively change its decay rate toward zero as a function of  $\gamma = \alpha/(\alpha - 1)$  and

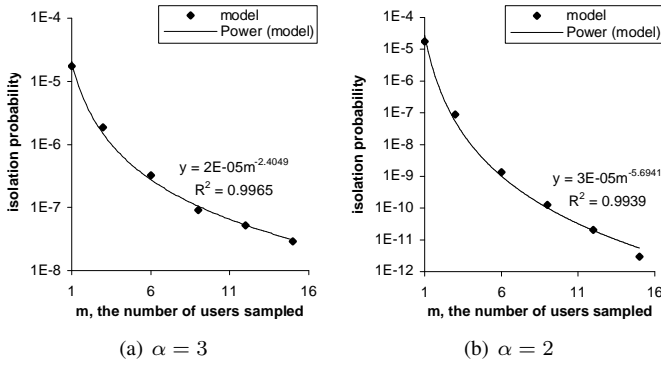


Fig. 5. Influence of  $m$  on model  $\phi$  under max-age selection for Pareto lifetimes with  $E[L] = 0.5$  hours, exponential search times with  $E[S] = 6$  minutes, and  $k = 7$ .

leads to no novel discussion. In the next section, however, we develop another neighbor selection framework that guarantees a much stronger result in which  $\phi$  converges to zero for any  $1 < \alpha \leq 2$ , any number of neighbors  $k \geq 1$ , and any search delay as system age and size tend to infinity. An additional reason for improving the max-age method in the next section is the difficulty of implementing uniform neighbor selection in decentralized P2P networks without global knowledge at each node. Distributed methods of uniform sampling of users exist [6], [31]; however, they require either  $k$ -regular graphs [6] or complex walk patterns [31]. In both cases, max-age selection forces a user to sample  $m$  peers to obtain a single neighbor and may not scale well for large  $m$ . In contrast, the method we describe below needs only *one* sample per neighbor and operates in graphs with irregular degree distributions.

#### IV. AGE-BIASED RANDOM-WALK SELECTION

In this section, we first introduce a new neighbor selection strategy that is based on random walks over weighted directed graphs and then deal with the distribution of neighbor residual lifetimes and the corresponding isolation probability.

##### A. Random Walks on Weighted Directed Graphs

We start by designing a simple random-walk algorithm, which we call *age-weighted*, whose stationary distribution  $\pi$  ensures that the probability that a user  $u$  is selected by another peer is proportional to  $u$ 's current age. Recall that a directed graph  $G = (V, E)$  consists of a vertex set  $V$  and edge set  $E$  (note that we use notation  $G$  instead of  $G(t)$  at time  $t$  under the assumption that  $G$  remains the same while a random walk is performed). Let  $u \rightarrow v$  represent a directed link  $(u, v) \in E$ ,  $N_u^+ = \{v \in V : u \rightarrow v\}$  be the set of out-degree neighbors of  $u$ , and  $N_u^- = \{v \in V : u \leftarrow v\}$  be the set of in-degree neighbors of  $u$ . Further define  $A_u$  to be the age of user  $u$  and set the weight of each incoming edge  $v \rightarrow u$  at node  $u$  to be  $u$ 's age normalized by the number of in-degree neighbors:

$$w(v, u) = \frac{A_u}{|N_u^-|}. \quad (29)$$

It then follows that the in-degree  $d_u^-$  of  $u$  is simply its age:

$$d_u^- = \sum_{v \in N_u^-} w(v, u) = A_u, \quad (30)$$

and its out-degree  $d_u^+$  is the sum of normalized ages of its out-degree neighbors:

$$d_u^+ = \sum_{v \in N_u^+} w(u, v) = \sum_{v \in N_u^+} \frac{A_v}{|N_v^-|}. \quad (31)$$

Then, age-weighted random walks are executed by alternating between walking along incoming and outgoing edges as we describe next. Given that the walk is currently at node  $u$ , the first jump is performed to an *in-degree* neighbor  $h$  of  $u$ ,  $h \in N_u^-$ , with probability

$$p_{uh} = \frac{w(h, u)}{d_u^-}. \quad (32)$$

The second jump is performed to an *out-degree* neighbor  $v$  of  $h$  with probability:

$$p_{hv} = \frac{w(h, v)}{d_h^+}. \quad (33)$$

It is clear that the transition probability from  $u$  to  $v$  is  $p_{uv} = \sum_{h \in N_u^-} p_{uh} p_{hv}$ . After the two jumps,  $v$  becomes the current node and this procedure repeats. Each step consists of two jumps, the node reached after  $l$  steps is selected as the neighbor of the current user. As shown in [32], the stationary distribution of this random walk is given by  $\pi = (\pi_u)$ , where  $\pi_u = d_u^- / \sum_{v \in V} d_v^-$ . Recalling (30), we immediately obtain that age-weighted random walks achieve the desired distribution:

$$\pi_u = \frac{A_u}{\sum_{v \in V} A_v}, \text{ for all } u \in V. \quad (34)$$

The starting point of a random walk is determined as follows. Each new user executes a random walk starting from an alive user obtained through bootstrap, while each existing user uniformly randomly selects one of its currently alive out-degree neighbors as the initial point of the walk. Note that if a node does not have any incoming edges, it will never be selected by our walk. To avoid this situation, we alternate between ending walks with an in-degree and an out-degree jump, which gives new users an opportunity to receive incoming edges. Simulations below use random walks of  $l = 10$  steps as further increasing  $l$  does not result in measurable improvements in  $\pi$  for the cases considered in this paper.

##### B. Residual Lifetime Distribution

Denote by  $Z$  the residual lifetimes of neighbors obtained by age-weighted random walks and by  $H^c(x) = P(Z > x)$  its CCDF. Assuming that the walk is executed for longer than the mixing time of the chain corresponding to the underlying graphs [16], we obtain the distribution of  $Z$  in the next theorem.

*Theorem 6:* Given that (34) holds for age-weighted random walks, neighbor residual lifetime  $Z$  has the following CCDF:

$$H^c(x) = \frac{1}{E[L]E[A]} \int_0^\infty y(1 - F(x + y))dy, \quad (35)$$

for  $E[L] < \infty$  and mean user age  $E[A] < \infty$ .




Fig. 6. Comparison of model  $\phi$  to simulations under age-weighted random walks for Pareto lifetimes,  $E[L] = 0.5$  hours,  $\beta = (\alpha - 1)E[L]$ , exponential search delays, and  $k = 7$  in a graph with  $n = 8,000$  nodes.

It is easy to show that for exponential lifetimes, (35) reduces to  $1 - F(x)$ , again confirming the memoryless property of exponential distributions. For Pareto lifetimes, the CCDF of  $Z$  is also very simple as we show in the next lemma.

*Lemma 6:* For Pareto lifetimes  $L \sim 1 - (1 + x/\beta)^{-\alpha}$  with  $\alpha > 2$ , the CCDF of  $Z$  is given by:

$$H^c(x) = \left(1 + \frac{x}{\beta}\right)^{-(\alpha-2)}. \quad (36)$$

For  $1 < \alpha \leq 2$ ,  $Z$  converges in probability to  $\infty$  as system age  $\mathcal{T}$  and size  $n$  both tend to  $\infty$ .

Note that for  $\alpha > 2$ , the PDF of  $Z$  is completely monotone and thus suitable for our hyper-exponential model. Also notice that  $Z$  is stochastically larger than residual lifetimes  $R$  under uniform selection for all choices of  $\alpha$ . In fact,  $Z$  shifts the shape of the Pareto distribution from  $\alpha$  to  $\alpha - 2$ , which is not achievable under max-age selection even as  $m \rightarrow \infty$ . Furthermore, for  $1 < \alpha \leq 2$ , residuals  $Z$  tend to a defective random variable with all mass concentrated at  $+\infty$  as system size and age become infinite. This shows that in asymptotically large systems,  $Z$  exceeds any lifetime  $L$  with probability 1 and no user suffers isolation (more on this below).

### C. Isolation and Resilience

To obtain model  $\phi$  under age-weighted random-walk selection, we fit the distribution of  $Z$  shown in (36) with its hyper-exponential equivalent and then invoke Theorem 2 to solve for  $\phi$ . Next, we test the accuracy of model  $\phi$  in simulations where  $n = 8,000$  nodes join and leave the system at random instances and each node performs age-weighted random walks to find its neighbors. As shown in Fig. 6, simulation results are very close to the values predicted by theoretical  $\phi$ . Examples showing the relationship between of  $\phi$  and  $\alpha$  are presented in Fig. 7. As shown in Fig. 7(a), simulation results are consistent with model  $\phi$  under a variety of values  $\alpha$  that allow quick simulations and do not require very large  $\mathcal{T}$  or  $n$  (i.e.,  $\alpha \geq 3$ ). It is interesting to observe in the figure that as  $\alpha$  decreases, the gap between  $\phi$  under age-weighted random walks and that under uniform selection drastically increases and reaches a factor of  $10^4$  for  $\alpha = 2.5$ . This shows that age-weighted random walks are extremely effective in systems with very heavy-tailed lifetimes (i.e.,  $\alpha$  below 2.5). Fig. 6(b) shows that

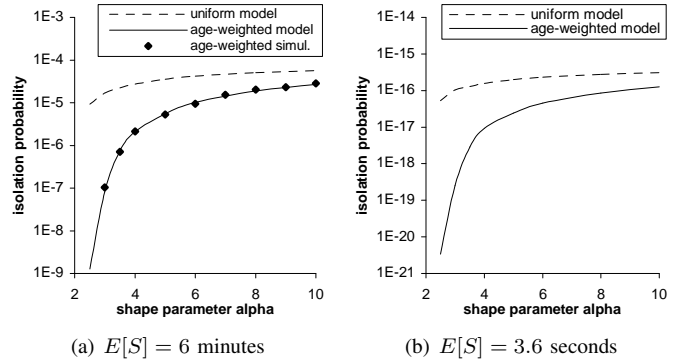


Fig. 7. Impact of  $\alpha$  on  $\phi$  under uniform selection and under age-weighted random walks for Pareto lifetimes,  $E[L] = 0.5$  hours,  $\beta = (\alpha - 1)E[L]$ , exponential search delays, and  $k = 7$ .

the same conclusion holds for  $E[S] = 3.6$  seconds, in which case  $\phi$  is on the order of  $10^{-20}$  and only allows computation using the model since simulations are impractical for such small probabilities.

The most intriguing result shown in Fig. 7 is that  $\phi$  tends to 0 as  $\alpha$  converges to 2 from above. However, as before, this convergence requires that system age tend to infinity. In addition, the following result states that system size  $n$  must also be infinite to obtain  $\phi = 0$ .

*Theorem 7:* For age-weighted random walks, Pareto lifetimes with  $1 < \alpha \leq 2$ , any number of neighbors  $k \geq 1$ , and any type of search delay (including  $S = \infty$ ), the isolation probability converges to zero as system age  $\mathcal{T}$  and size  $n$  approach infinity:  $\lim_{n \rightarrow \infty} \lim_{\mathcal{T} \rightarrow \infty} \phi = 0$ .

Note that Theorem 7 is a much stronger result than Theorem 3 since  $\phi$  under uniform selection does *not* asymptotically approach 0 for any fixed  $\alpha \in (1, 2]$ . However, the asymptotic result of this section is more difficult to achieve since it requires not only an equilibrium system, but also an infinitely large user population.

We finish this section by examining age-weighted random walks under finite  $\mathcal{T}$  and  $n$  using several values of  $1 < \alpha \leq 2$ . For such cases, recall from Lemma 6 that  $Z$  converges in probability to  $\infty$ ; however, initial analysis shows that the convergence rate of  $Z \rightarrow \infty$  and  $\phi \rightarrow 0$  can only be expressed using complex Appell hypergeometric functions [4] of  $\mathcal{T}$  and  $n$  for which no closed-form expansion exists. We leave this task for future work and instead show simulations of  $\phi$  in Fig. 8 as  $\mathcal{T}$  becomes large ( $n$  is kept equal to  $\mathcal{T}/10$ ). For both values of  $\alpha$ , the figure shows that  $\phi$  monotonically decreases as system age  $\mathcal{T}$  increases. In fact, for  $k = 7$ , the system achieves isolation probability below  $10^{-7}$  without replacing neighbors at  $\mathcal{T} = 30,000$  hours and  $n = 3,000$  users. Additional simulations with  $k = 7$  suggest that increasing  $n$  to over 1 million users and keeping the age around 1 year will produce  $\phi$  sufficiently small for most large-scale networks today.

## V. RELATED WORK

Construction and maintenance of overlay networks consists of initial neighbor selection and subsequent replacement of dead links. Many P2P systems, including structured [19], [26], [30], [23] and unstructured [2], [24], [17], perform



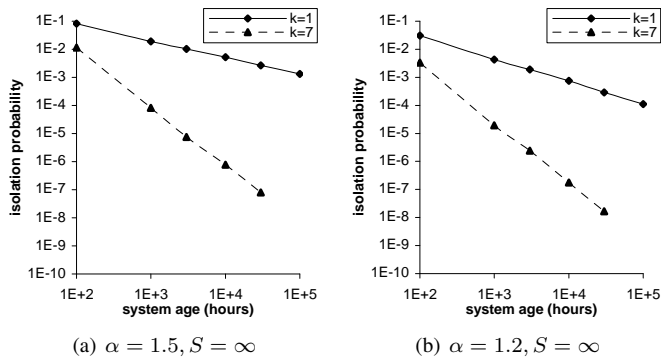


Fig. 8. Simulation results of  $\phi$  under age-weighted selection as system age  $T$  and size  $n$  increase for Pareto lifetimes with  $E[L] = 0.5$  hours.

neighbor selection and replacement to achieve the desired routing efficiency and search performance in the face of node joins and departures. Previous work has used proximity-based neighbor selection to reduce lookup latency [7], [20], [29], [17], capacity-based selection to improve system scalability [2], [11], and age-biased neighbor preference to improve reliability of the system [1], [11], [25]. Additional studies have analyzed the tradeoffs between resilience and proximity [3] as well as studied how well different neighbor selection and recovery strategies could handle churn in DHTs [22]. Finally, only a handful of modeling studies of user isolation and neighbor selection under churn exist [9], [12], [15], [18] and they are mostly limited to exponential user lifetimes and centralized (age-unrelated) user replacement.

## VI. CONCLUSION

This paper derived a general model of resilience for unstructured P2P networks under heavy-tailed user lifetimes and formally analyzed two age-dependent neighbor-selection techniques. Our results show that the proposed random-walk method may achieve *any* desired level of resilience without replacing the neighbors as long as  $1 < \alpha \leq 2$  and system size  $n$  and age  $T$  are sufficiently large. This indicates that P2P systems under proposed neighbor selection and very heavy-tailed lifetimes (i.e.,  $\alpha < 2$ ) become progressively more resilient over time and asymptotically tend to an “ideal” system that never disconnects as users join the network.

Future work includes derivation of residual lifetime distributions in finite systems under age-weighted walks and analysis of the limiting distribution of  $U_m$  as  $m \rightarrow \infty$ .

## REFERENCES

- [1] F. E. Bustamante and Y. Qiao, “Friendships that Last: Peer Lifespan and its Role in P2P Protocols,” in *Proc. Intl. Workshop on Web Content Caching and Distribution*, Sep. 2003.
- [2] Y. Chawathe, S. Ratnasamy, L. Breslau, N. Lanham, and S. Shenker, “Making Gnutella-like P2P Systems Scalable,” in *Proc. ACM SIGCOMM*, Aug. 2003, pp. 407–418.
- [3] 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.
- [4] H. Exton, *Handbook of Hypergeometric Integrals: Theory, Applications, Tables, Computer Programs*. Ellis Horwood, Chichester, England, 1978.
- [5] 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.

- [6] C. Gkantsidis, M. Mihail, and A. Saberi, “Random Walks in Peer-to-Peer Networks,” in *Proc. IEEE INFOCOM*, Mar. 2004, pp. 120–130.
- [7] 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.
- [8] T. Hettmansperger and M. Keenan, “Tailweight, Statistical Inference, and Families of Distributions – A Brief Survey,” *Statist. Distributions in Scientific Work*, G. P. Patil et al. (eds), vol. 1, pp. 161–172, 1980.
- [9] S. Krishnamurthy, S. El-Ansari, E. Aurell, and S. Haridi, “A Statistical Theory of Chord under Churn,” in *Proc. IPTPS*, Feb. 2005, pp. 93–103.
- [10] S. S. Lam and H. Liu, “Failure Recovery for Structured P2P Networks: Protocol Design and Performance Evaluation,” in *Proc. ACM SIGMETRICS*, Jun. 2004, pp. 199–210.
- [11] J. Ledlie, J. Sheidman, M. Amis, and M. Seltzer, “Reliability- and Capacity-Based Selection in Distributed Hash Tables,” Harvard University Computer Science, Tech. Rep., Sep. 2003.
- [12] 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, pp. 26–37.
- [13] 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, pp. 345–357.
- [14] 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.
- [15] D. Liben-Nowell, H. Balakrishnan, and D. Karger, “Analysis of the Evolution of the Peer-to-Peer Systems,” in *Proc. ACM PODC*, Jul. 2002, pp. 233–242.
- [16] L. Lovász, “Random Walks on Graphs: A Survey,” *Combinatorics, Paul Erdős is Eighty*, D. Miklós et al. (eds), vol. 2, pp. 353–398, 1996.
- [17] L. Massoulié, A.-M. Kermarrec, and A. Ganesh, “Network Awareness and Failure Resilience in Self-Organising Overlay Networks,” in *Proc. IEEE SRDS*, Oct. 2003, pp. 47–55.
- [18] 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.
- [19] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Shenker, “A Scalable Content-Addressable Network,” in *Proc. ACM SIGCOMM*, Aug. 2001, pp. 161–172.
- [20] S. Ratnasamy, M. Handley, R. Karp, and S. Shenker, “Topologically-Aware Overlay Construction and Server Selection,” in *Proc. IEEE INFOCOM*, Jun. 2002, pp. 1190–1199.
- [21] S. Resnick, *Adventures in Stochastic Processes*. Birkhäuser, 2002.
- [22] S. Rhea, D. Geels, T. Roscoe, and J. Kubiatowicz, “Handling Churn in a DHT,” in *Proc. USENIX Ann. Tech. Conf.*, Jun. 2004, pp. 127–140.
- [23] A. Rowstron and P. Druschel, “Pastry: Scalable, Decentralized Object Location and Routing for Large-Scale Peer-to-Peer Systems,” in *Proc. IFIP/ACM International Conference on Distributed Systems Platforms (Middleware)*, Nov. 2001, pp. 329–350.
- [24] D. Rubenstein and S. Sahu, “Can Unstructured P2P Protocols Survive Flash Crowds,” *IEEE/ACM Trans. Netw.*, vol. 13, no. 3, pp. 501–512, Apr. 2005.
- [25] K. Sripanidkulchai, A. Ganjam, B. Maggs, and H. Zhang, “The Feasibility of Supporting Large-Scale Live Streaming Applications with Dynamic Application End-Points,” in *Proc. ACM SIGCOMM*, Aug. 2004, pp. 107–120.
- [26] 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.
- [27] M. S. Taqqu, W. Willinger, and R. Sherman, “Proof of a Fundamental Result in Self-Similar Traffic Modeling,” *ACM Comput. Commun. Rev.*, vol. 27, no. 2, pp. 5–23, Apr. 1997.
- [28] R. W. Wolff, *Stochastic Modeling and the Theory of Queues*. Prentice Hall, 1989.
- [29] H. Zhang, A. Goal, and R. Govindan, “Incrementally Improving Lookup Latency in Distributed Hash Table Systems,” in *Proc. ACM SIGMETRICS*, Jun. 2003, pp. 114–125.
- [30] B. Y. Zhao, L. Huang, J. Stribling, S. C. Rhea, A. D. Joseph, and J. Kubiatowicz, “Tapestry: A Resilient Global-Scale Overlay for Service Deployment,” *IEEE J. Sel. Areas Commun.*, vol. 22, no. 1, pp. 41–53, Jan. 2004.
- [31] 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.
- [32] D. Zhou, J. Huang, and B. Schölkopf, “Learning from Labeled and Unlabeled Data on a Directed Graph,” in *Proc. ICML 2005*, Aug. 2005, pp. 1036–1043.