Unstructured P2P Link Lifetimes Redux

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

***Abstract*—We revisit link lifetimes in random P2P graphs under dynamic node failure and create a unifying stochastic model that generalizes the majority of previous efforts in this direction. We not only allow nonexponential user lifetimes and age-dependent neighbor selection, but also cover both active and passive neighbor-management strategies, model the lifetimes of incoming and outgoing links, derive churn-related message volume of the system, and obtain the distribution of transient in/out degree at each user. We then discuss the impact of design parameters on overhead and resilience of the network.**

***Index Terms*—In-degree, stochastic modeling, user churn.**

I. Introduction

P2P NETWORKS organize users into a distributed graph that is jointly maintained and dynamically restructured by its participants under churn [2], [3], [9], [11], [12], [18], [21], [24]. Many P2P properties (e.g., message overhead, resilience to disconnection, and ability to reach other peers with queries) depend on the behavior of node degree, which is determined solely by the lifetime of edges in the graph. Despite the sizeable volume of analytical work on P2P networks [7], [13]–[17], [19], [25], [27]–[31], accurate characterization of link lifetime has been elusive.

We start by defining terminology and our modeling objectives. Suppose L_v is the random lifetime of user v and $R_v(t)$ is its *residual* (i.e., remaining) lifetime at time t, conditioned on v being alive at t. If peer w creates link $w \to v$ during join into the system or repair of broken edges, we call w the *initiator* and v the *recipient* of the connection. For such a link created at time t, there are actually two lifetimes: *out-link* duration $V = R_v(t)$, which is how long the connection stays active from w's perspective, and *in-link* duration $W = R_w(t)$, which is the same from v's perspective.

In the traditional sense, the link remains online only for $\min(V, W)$ time units. However, the degree at each user depends asymmetrically on the individual variables V and W, which makes *them*, rather than $\min(V, W)$, our target in this paper. It should be noted that links are treated as directional for the analysis. However, system performance (e.g., query routing and resilience) is still determined by the combined in/out degree at each user (i.e., edges are undirected for all other purposes).

Link lifetimes depend on how peers select their neighbors during join and replacement of failed edges. If this process is independent of age (e.g., based on geographical proximity, random hash function, presence of certain shared content), then analysis falls under the so-called *uniform selection*, where it has been shown [30] that V is the residual of L_v. However, even under uniform selection, the distribution of in-link lifetime W has remained unexplored.

For *age-biased* neighbor selection, two methods have been proposed in the analytical literature. The first one, called *max-age* [25], [31], selects m uniformly random peers and then picks the one with the largest age. The rationale is that under heavy-tailed user lifetimes, residuals are stochastically larger for users with higher age.[1] The second method, called *age-proportional* [31], selects each user in linear proportion to its current age. This is implemented using a random walk on the graph using a Markov chain whose transition probabilities are functions of current ages of the users adjacent to each link. For these two specific techniques, the distribution of V has been derived in [25] and [31]. However, extension to more general preference functions or analysis of W has not been offered.

A. Contributions

To understand the impact of neighbor choice on the degree of the system and message overhead to maintain the graph, our first contribution is to propose a novel modeling paradigm for out-link churn that allows arbitrary age-biased neighbor selection using a general preference function $p(x)$, where x is the age of potential neighbors at the time of edge creation. We provide a set of conditions under which there exists a simple expression for the asymptotic distribution of V as network size $n \to \infty$ and explain how to select $p(x)$ to obtain the three special cases considered in prior work (i.e., uniform, max-age, and age-proportional).

The new model is flexible enough to cover both *active* and *passive* systems (i.e., with and without neighbor replacement [14]), which represent the two most commonly modeled approaches. Since max-age employs a very complex nonlinear $p(x)$ that does not immediately reveal the impact of m on $E[V]$, we propose an alternative mechanism that performs similarly, but allows closed-form tuning of out-link lifetime.

Our second contribution is to analyze the edge-replacement process and obtain the rate λ at which neighbors are sought in the system. Since each search may require substantial network resources (e.g., flooding and/or random walks), minimization of λ may be beneficial in practice. We show that λ, which depends

Manuscript received June 07, 2013; revised December 10, 2013; accepted February 03, 2014; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor Y. Liu. Date of publication February 27, 2014; date of current version June 12, 2015. An earlier version of the paper appeared in the IEEE International Conference on Computer Communications (INFOCOM), Turin, Italy, April 14–19, 2013.

Z. Yao is with the University of Dayton, Dayton, OH 45469 USA (e-mail: zyao@udayton.edu).

D. B. H. Cline and D. Loguinov are with Texas A&M University, College Station, TX 77843 USA (e-mail: dcline@stat.tamu.edu; dmitri@cs.tamu.edu).

Digital Object Identifier 10.1109/TNET.2014.2306153

[1]Variable X is said to be stochastically larger than Y iff $P(X > x) \geq P(Y > x)$ for all x[20]. We call lifetimes *heavy-tailed* if $P(L > x + y | L > y) \geq P(L > x)$ for all x, y, i.e., given that a user has survived to any positive age y, the remaining lifetime is stochastically larger than L. If the inequality is reversed, we call such distributions *light-tailed*.

on the distribution of V and replacement delay S, can be controlled using $p(x)$ and is automatically minimized by any P2P system whose V is sufficiently heavy-tailed (e.g., Pareto lifetimes with $\alpha \le 2$ and age-proportional selection).

Our third contribution is to develop a novel approach to modeling the distribution of in-link lifetime W. We show that under Pareto lifetimes (often observed in real P2P systems [1], [22], [26]), W is stochastically larger than lifetimes L, but smaller than residuals R. Interestingly, this indicates that in-link users are more reliable than new arrivals, but less so than random live peers in the system. We also observe that increasing the bias toward nodes with large age, i.e., using a more aggressive $p(x)$, leads to a surprising reduction in $E[W]$. This indicates that there exists an inherent tradeoff between in- and out-edge resilience. As V becomes stochastically larger, W gets stochastically smaller and eventually converges in distribution to L. While a somewhat similar result was observed in DHTs [9], [29], the reasons for these phenomena are completely different as we discuss below.

Our fourth contribution is to show that incoming links in the proposed framework are delivered to each peer through a nonhomogeneous Poisson process whose rate is determined by the age-preference function $p(x)$. This allows us to obtain the transient distribution of in-degree $D_{\text{in}}(\tau)$, where τ is the current age of a live peer, extending the result of [27] to nonuniform selection. We discover that bounded preference functions (e.g., uniform, max-age) guarantee finite $E[D_{\text{in}}(\tau)]$ as the user's age $\tau \to \infty$. On the other hand, unbounded preference functions (e.g., age-proportional) grow in-degree to infinity, which inevitably forces popular users to reject incoming requests after they become overloaded (as often seen in Gnutella [8]). This not only increases neighbor search latency and message overhead, but also does not guarantee eventual connection success in asymptotically large networks.

We finish the paper by studying the combined in/out degree $D(\tau)$ in both passive and active systems, making observations on the usage of our models to select parameters of the system to achieve desired performance, which forms our fifth contribution.

II. Out-Link Churn

To model a P2P system, one requires three underlying assumptions—a churn model, neighbor-replacement behavior at each peer, and a preference function during link formation. We outline these next.

A. Active Systems

Consider a network of n participants forming a random P2P graph, where each node w can be modeled by a stationary alternating-renewal process representing the user's ON/OFF states [28, Sec. III]. To allow for heterogeneity in user behavior, we assume that peer w randomly draws its lifetime cumulative distribution function (CDF) from some finite pool of available distributions and maintains $k_w \ge 1$ *outbound* links to existing peers in the graph. Repair of broken connections along out-links incurs some random delay that is needed to detect the failure and find a replacement user. This process is illustrated in Fig. 1, which shows the status of the first two outgoing links of user w.

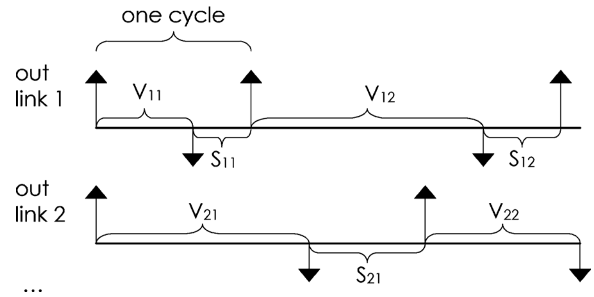

Fig. 1. Active model: connection churn along out-links at user w.

In the figure, the direction of the arrows indicates whether the link is going up (upon creation) or down (upon failure), V_{ij} is the remaining lifetime of the jth selected neighbor along the ith link, and S_{ij} is the corresponding search delay.

Note that *inbound* links are never repaired as this would lead to an explosive (snowball) edge-creation process and eventually a complete graph. As $n \to \infty$, the system described above is fully equivalent to a homogeneous network with $k := E[k_w]$ initial outbound connections and all users having the same lifetime CDF $F_L(x)$, which is a mixture of all possible lifetime distributions weighted by the probability that users select them and the frequency of each user's appearance in the system [28, Theorem 1].

B. Passive Systems

An alternative approach [14] is to never replace the failed links and only restrict neighbor creation to the k_w initial edges during join. This model simplifies operation and reduces overhead at the expense of seemingly poor resilience and low branching factor during search. However, coupling between the diminishing expected out-degree $E[D_{\text{out}}(\tau)]$ and the increasing expected in-degree $E[D_{\text{in}}(\tau)]$ as user age $\tau \to \infty$ creates an intriguing possibility that the average combined degree $E[D(\tau)]$ may remain more or less constant. If so, this allows the user to stay connected with almost no superfluous activity (e.g., keep-alive messages, flooding of the graph to find replacement neighbors). As this idea has not been modeled before, we naturally have to investigate its viability later in the paper.

C. Age-Dependent Neighbor Selection

The rest of this section presents our first contribution—a novel modeling framework for out-link churn that subsumes all previous approaches in this field by allowing arbitrary age-biased neighbor selection. While the results below typically require $n \to \infty$, one should not be discouraged by this assumption since systems with just a few thousand peers match the developed theory very accurately.

At time t, assume a stationary network with $N = N(t) \le n$ live users whose ages $A_1, \dots, A_N$ form a collection of asymptotically independent identically distributed (i.i.d.) random variables with distribution [28]

$$F_A(x) := P(A_i < x) = \frac{1}{E[L]} \int_0^x \bar{F}_L(y)\, dy \qquad (1)$$

where $\bar{F}_L(x) = 1 - F_L(x)$ is the tail CDF of user lifetimes. For Pareto $F_L(x) = 1 - (1 + x/\beta)^{-\alpha}$, it is well known [14] that the shape parameter of both age A and residual R is $\alpha - 1$, i.e., $F_A(x) = 1 - (1 + x/\beta)^{1-\alpha}$. For our later results, define the residual of A (or the *double residual* of L) to be a random variable Z with CDF

$$F_Z(x) := P(Z < x) = \frac{1}{E[A]} \int_0^x \bar{F}_A(y)\, dy \tag{2}$$

where $\alpha > 2$ is assumed if L is Pareto, in which case $F_Z(x)$ is conveniently $1 - (1 + x/\beta)^{2-\alpha}$.

Next, suppose $c_N(v) = P(w \to v \mid A_1, \ldots, A_{N-1})$ is the probability that w connects to v, assuming the latter is alive and conditioning on the ages of $N - 1$ live peers other than w. As $n \to \infty$, selection strategy $c_N(v)$ may obtain users with a distribution of residuals that does not converge. To preclude such cases, we require that $c_N(v)$ asymptotically pick user v proportional to some function of its current age.

Assumption 1: There exists a nonnegative weight function $p(x) \geq 0$ with $p(x) = 0$ for $x < 0$ and $E[p(A)] < \infty$, where $A \sim F_A(x)$, such that

$$\sum_{v=1}^{N} \left| c_N(v) - \frac{p(A_v)}{E[p(A)]N} \right| \to 0 \tag{3}$$

in distribution as $n \to \infty$.

Since the probability space changes with n, the random sum in (3) must converge *in distribution* rather than in probability. We start with two preliminary results that we use throughout the paper. The first one presents a convenient expression for integrals with random upper limits.

Lemma 1: Assume X is a nonnegative random variable with residual Y and $g(x) \geq 0$ is some function. Then

$$E\left[\int_0^X g(y)\, dy\right] = E[X]E[g(Y)]. \tag{4}$$

Proof: As a convention, we write CDFs in uppercase and densities in lowercase, both subscripted with the name of the corresponding variable. Expanding (4)

$$E\left[\int_0^X g(y)\, dy\right] = -\int_0^\infty \left[\int_0^x g(y)\, dy\right] d\bar{F}_X(x). \tag{5}$$

Integrating by parts, noticing that $\bar{F}_X(\infty) = 0$, and using $\bar{F}_X(x) = E[X]f_Y(x)$, observe that (5) becomes

$$\int_0^\infty g(x)\bar{F}_X(x)\, dx = E[X] \int_0^\infty g(x) f_Y(x)\, dx \tag{6}$$

which is the same as (4). ■

The second preliminary result expands the integral of $E[g(X - t)]$ using the residual of X.

Lemma 2: Assume the same as in Lemma 1, but additionally suppose $g(x) = 0$ for $x < 0$. Then, for any $a \geq 0$

$$\int_0^a E[g(X - t)]\, dt = E[X]E[g(Y) - g(Y - a)]. \tag{7}$$

Proof: Rewriting (7)

$$\int_0^a E[g(X - t)]\, dt = E\left[\int_{X-a}^X g(y)\, dy\right]. \tag{8}$$

Since $g(.)$ is zero for negative arguments, we can split the integral into a difference between two expectations

$$E\left[\int_0^X g(y)\, dy\right] - E\left[\int_0^X g(y - a)\, dy\right]. \tag{9}$$

Applying (4) to each term of (9), we get (7). ■

We now fix peer w and deal with the distribution of its out-link lifetime V_{ij}.

Theorem 1: Assuming (3) holds and $n \to \infty$, the collection of variables $\{V_{ij}\}$ is asymptotically i.i.d. with tail distribution

$$\bar{F}_V(x) = \frac{E[p(A - x)]}{E[p(A)]} = \frac{E[p(A - x) | A \geq x]}{E[p(A)]} \bar{F}_A(x) \tag{10}$$

and mean

$$E[V] = \frac{E[p(Z)]}{E[p(A)]} E[A]. \tag{11}$$

Proof: Asymptotic independence of the V_{ij}'s is a consequence of the following facts. For any positive integer j, the chance that peer w selects the same individual more than once in the first j cycles (across all k_w links) in Fig. 1 becomes negligible as $n \to \infty$. Moreover, there are increasingly many users of similar ages, and the choice of one does not substantially change the pool of remaining users of similar ages. Therefore, collection of ages $A_1, \ldots, A_{N-1}$ for live users at any time t is asymptotically i.i.d. with distribution $F_A(x)$. Thus, the users selected by peer w during its first j cycles are asymptotically independent, which holds for all j and thus shows that V_{ij}'s converge in distribution to an i.i.d. set.

Next, we establish (10). Recall that the joint distribution of age A and residual time R from renewal theory implies

$$P(R > x | A = a) = \frac{1 - F_L(a + x)}{1 - F_L(a)} \tag{12}$$

where $a \geq 0$ and $y \geq 0$. Now consider $N - 1$ independent users with ages $A_1, \ldots, A_{N-1}$, and let $R_1, \ldots, R_{N-1}$ be the corresponding residual lifetimes, which are conditionally independent, given all the ages [5]. Thus, the tail distribution of residual lifetime for any live user v is asymptotically

$$P(R_v > x | A_1, \ldots, A_{N-1}) = \frac{1 - F_L(A_v + x)}{1 - F_L(A_v)}. \tag{13}$$

Let R_N^* be the residual of the user that w selects conditional on $A_1, \ldots, A_{N-1}$. Then, its tail CDF is a sum of probabilities that each live user v survives at least x time units and w randomly selects v as a neighbor

$$P(R_N^* > x) = \sum_{v=1}^{N} P(R_v > x \mid A_1, \ldots, A_{N-1}) c_N(v). \tag{14}$$

Applying (13)

$$P(R_N^* > x) = \frac{1}{N} \sum_{v=1}^{N} \frac{1 - F_L(A_v + x)}{1 - F_L(A_v)} N c_N(v). \tag{15}$$

Invoking (3), this reduces in the limit as $n \to \infty$ to

$$\bar{F}_V(x) = E\left[\frac{1 - F_L(A + x)}{1 - F_L(A)} \frac{p(A)}{E[p(A)]}\right]. \tag{16}$$

Noticing that $1 - F_L(y) = f_A(y)E[L]$, where $f_A(x) = F'_A(x)$ is the density of age, and expanding the expectation

$$\bar{F}_V(x) = \frac{1}{E[L]E[p(A)]} \int_0^\infty \bar{F}_L(x+y)p(y)\,dy. \quad (17)$$

Applying $\bar{F}_L(x+y) = E[L]f_A(x+y)$ in (17), we get

$$\bar{F}_V(x) = \frac{1}{E[p(A)]} \int_x^\infty f_A(z)p(z-x)\,dz \quad (18)$$

which is equivalent to the second formula in (10). Recalling that $p(x)$ is zero for negative arguments, we get the first formula of (10) as well.

We now establish (11). Recall that the expectation of nonnegative random variables is the integral of its tail [20]. Using the first expression in (10), this observation produces

$$E[V] = \frac{1}{E[p(A)]} \int_0^\infty E[p(A-x)]\,dx. \quad (19)$$

Applying (7) with $a = \infty$ to this integral, we get (11). ■

Theorem 1 shows that weight $p(x)$ serves as a simple tuning knob for out-neighbor resilience. Specifically, the tail of V in (10) is that of age A scaled by a normalization factor $E[p(A-x)|A \geq x]/E[p(A)]$. Under heavy-tailed L, variable $A - x$ for $A \geq x$ is stochastically larger than A, which indicates that V is stochastically larger than A for nondecreasing $p(x)$ and stochastically smaller for nonincreasing $p(x)$. For light-tailed L, this relationship is reversed.

Similarly, the expected residual $E[V]$ in (11) is that of a random live peer (i.e., $E[A]$) normalized by the ratio of $E[p(Z)]$ to $E[p(A)]$. For heavy-tailed lifetimes, where Z is stochastically larger than A, this leads to $E[V] \geq E[A]$ if $p(x)$ is nondecreasing and $E[V] \leq E[A]$ if $p(x)$ is nonincreasing. For light-tailed distributions, this relationship is again reversed since Z is stochastically smaller than A in those cases.

D. Examples

We next consider selection strategies used in prior literature and explain how to map them into our new model. In the first strategy, suppose w finds neighbors in proportion to some function $h(x)$ applied to peer age

$$c_{1,N}(v) = \frac{h(A_v)}{\sum_{i=1}^{N} h(A_i)} \quad (20)$$

which produces uniform [14] and age-proportional [31] methods using $h(x) = 1$ and $h(x) = x$, respectively.

In the second strategy, w uniformly randomly selects $m \geq 1$ users from the system into a set Γ and then picks the sth-order statistic (e.g., minimum, maximum, median) of the sampled ages to identify the best neighbor, where $s \leq m$. To obtain the corresponding connection probability, denote by r_v the rank order of A_v among the ages of N live users (from the smallest to the largest) and observe that this technique exhibits

$$c_{2,N}(v) = \binom{N}{m}^{-1} \binom{r_v - 1}{s-1} \binom{N - r_v}{m-s} \quad (21)$$

which is the number of ways to select $s - 1$ ages smaller than A_v and $m - s$ ages larger than A_v in a system of N users, normalized by the number of ways to pick m initial peers. Max-age selection [25], [31] falls under (21) with $s = m$. For light-tailed L, usage of $s = 1$ (i.e., min-age) might be more appropriate instead. Note that both m and s could depend on N, but must be bounded as $n \to \infty$.

While the above two strategies are seemingly different, they in fact can be reduced to the same asymptotic model.

Theorem 2: Both (20) and (21) satisfy (3) with respective weights $p_1(x) = h(x)$ and

$$p_2(x) = m\binom{m-1}{s-1} F_A^{s-1}(x)(1 - F_A(x))^{m-s} \quad (22)$$

where $E[p_2(A)] = 1$.

Proof: From the law of large numbers, (20) is directly equivalent to (3). However, analysis of (21) requires more work. To this end, using the properties of rank-order statistics [23], we first obtain that

$$\max_{1 \leq v \leq N} \left| \frac{r_v}{N} - F_A(A_v) \right| \to 0 \quad (23)$$

in distribution as $n \to \infty$. This result says that the fraction of users whose rank is below that of v converges in distribution to the actual probability that the random age of a live user is smaller than A_v, which holds uniformly for all v.

Next, recalling that $m/N \to 0$ as $n \to \infty$ and leveraging (23) in the limit, we get

$$\frac{\binom{r_v - 1}{s-1}}{\binom{N-1}{s-1}} = \prod_{j=1}^{s-1} \frac{r_v - j}{N - j} \to F_A^{s-1}(A_v) \quad (24)$$

and

$$\frac{\binom{N - r_v}{m-s}}{\binom{N-s}{m-s}} = \prod_{j=1}^{m-s} \frac{N - r_v + 1 - j}{N - s + 1 - j} \to \bar{F}_A^{m-s}(A_v). \quad (25)$$

Using (24) and (25) in (21) shows that

$$Nc_{2,N}(v) \to \frac{F_A^{s-1}(A_v)(1 - F_A(A_v))^{m-s}\binom{N-1}{s-1}\binom{N-s}{m-s}}{\binom{N}{m}} = p_2(A_v). \quad (26)$$

We thus get that the maximum deviation of $Nc_{2,N}(v)$ from $p_2(A_v)$ across all v shrinks to zero as $n \to \infty$

$$\max_{1 \leq v \leq N} |Nc_{2,N}(v) - p_2(A_v)| \to 0. \quad (27)$$

In order for (27) to be equivalent to (3), we must next show that $E[p_2(A)] = 1$ since the latter has this term, while the former does not. Letting $B(x, y)$ be the Beta function

$$\begin{aligned} E[p_2(A)] &= m\binom{m-1}{s-1} \int_0^\infty F_A^{s-1}(x)(1 - F_A(x))^{m-s} dF_A(x) \\ &= \frac{m!}{(s-1)!(m-s)!} \int_0^1 y^{s-1}(1-y)^{m-s}\,dy \\ &= \frac{m!}{(s-1)!(m-s)!} B(s, m-s+1) = 1. \end{aligned} \quad (28)$$

Finally, noticing that for any connection function $c_N(v)$ and set of user ages $\{A_1, \dots, A_N\}$

$$\sum_{v=1}^{N} \left| c_N(v) - \frac{p(A_v)}{N} \right| \leq N \max_{1 \leq v \leq N} \left| c_N(v) - \frac{p(A_v)}{N} \right| \quad (29)$$

we get (3) after invoking (27). ■

Besides the two strategies explained above, a much wider class of methods can be covered under the umbrella of (3) as long as w asymptotically selects each live user v with probability proportional to $p(A_v)$. This modeling approach conveniently decouples analysis from complex summations $\sum_{i=1}^{N} h(A_i)$ in the first strategy, sets Γ in the second strategy, and various other details whose contribution in the limit is insignificant. From this point on, we do not dwell on the exact details of neighbor selection, but instead assume that it satisfies (3) and is uniquely described by $p(x)$.

We now return to (10) and (11) to perform a self-check against prior derivations in the literature and common sense. Result (10) becomes [28, Lemma 2] under uniform selection, [31, Theorem 4] under max-age, and [31, Theorem 6] under age-proportional. For exponential lifetimes and any $p(x)$, we trivially get $V \sim F_L(x)$. In (11), uniform selection produces $E[V] = E[A] = E[R]$ and age-proportional $E[V] = E[Z]$, which is consistent with the previously obtained expressions for these special cases [31]. As expected, exponential L leads to $E[V] = E[L]$ for all $p(x)$.

E. Max-Age Discussion

The max-age $p(x) = mF_A^{m-1}(x)$ is an interesting function in the sense that it favors older peers, but without becoming unbounded in x like age-proportional. The main stumbling block to understanding max-age is the obscure impact of m on $E[V]$, even when we have a simple closed-form model for the double-residual Z. In addition, the complex shape of max-age's $p(x)$ makes computation of various metrics developed below very tedious. To overcome this problem, we next propose an approximation to the max-age technique that allows a simple closed-form expression for $E[V]$. Also note that replacing max-age with a *directly evaluated* function $p(x)$ in (20) avoids drawing m initial samples,

Since multiplying $p(x)$ by a constant does not affect V, we need to consider only term $F_A^{m-1}(x)$, which stays near zero for small x, then makes a sharp, almost linear, transition to 1 at some threshold x_0, and finally remains near 1 for larger x. While a three-segment piecewise linear approximation is possible, we find that $p_{\text{step}}(x) = \mathbf{1}_{x \geq x_0}$, where $\mathbf{1}_e$ is an indicator of event e, is sufficient for our examples. To ballpark x_0, one needs to solve $F_A^{m-1}(x_0) = c$, where c is the desired level (such as 0.5) at which the transition from 0 to 1 is considered nonnegligible. Assuming $F_A^{-1}(.)$ denotes the inverse of the CDF $F_A(x)$ and $m \geq 2$, we get

$$x_0 = F_A^{-1}(c^{1/(m-1)}). \quad (30)$$

Using $c = 0.35$, we construct a step-function to approximate max-age with $m = 20$ under Pareto lifetimes with $\alpha = 3$, which results in $x_0 = 3.3$ h. Fig. 2(a) shows the resulting tail distributions of V, together with those of age-proportional V and residual R. As $x \to \infty$, the figure shows that the tails

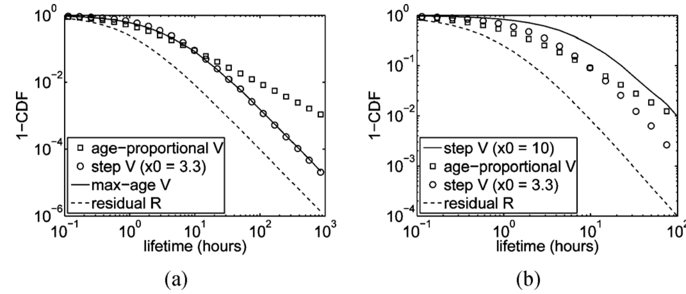

Fig. 2. Tail of V for Pareto L with $\alpha = 3$, $E[L] = 0.5$ h. (a) Comparison to max-age. (b) Different x_0.

$\bar{F}_V(x) = P(V > x)$ under our approximation and max-age are indistinguishable. As both tails exhibit a linear slope matching that of R, it can be conjectured that they are Pareto-like with shape $\alpha - 1$. For max-age, verifying this result is nontrivial, but the step-function readily produces from (10)

$$P(V_{\text{step}} > x) = \frac{\bar{F}_A(x + x_0)}{\bar{F}_A(x_0)} = \left(1 + \frac{x}{x_0 + \beta}\right)^{1-\alpha} \quad (31)$$

which shows that V_{step} is Pareto$(\alpha - 1, \beta + x_0)$.

Going back to Fig. 2(a), observe that the age-proportional tail is much heavier than the other ones since its $V \sim F_Z(x)$ is Pareto with $\alpha - 2$. In fact, the age-proportional scheme more than doubles $E[V]$ of max-age with $m = 20$. While it is possible to blindly tweak parameter m in max-age to obtain the same $E[V]$, we next show that $p_{\text{step}}(x)$ allows a simple closed-form relationship between x_0 and $E[V]$. Applying (11)

$$E[V_{\text{step}}] = \frac{E[A]P(Z > x_0)}{P(A > x_0)} \quad (32)$$

which under Pareto lifetimes becomes

$$E[V_{\text{step}}] = E[A]\left(1 + \frac{x_0}{\beta}\right). \quad (33)$$

For example, with $x_0 = 10$ and the same parameters as in the figure (i.e., $E[A] = 1$ h, $\beta = 1$), we achieve $E[V] = 11$ h, which compares favorably to age-proportional's 8.6. Fig. 2(b) shows the resulting tails. In contrast, age-proportional does not have any tuning knobs as all linear functions $p(x) = ax$ are equivalent when used in (20).

F. Power-Law Bias

Since power functions logically generalize uniform and age-proportional, we investigate them next.

Theorem 3: Suppose L is Pareto(α, β) with $\alpha > 1$ (i.e., $E[L] < \infty$) and $p(x) = x^\rho$, where $-1 < \rho < \alpha - 1$. Then, V is Pareto$(\alpha - \rho - 1, \beta)$.

Proof: Rewriting (17)

$$\bar{F}_V(x) = \frac{1}{E[L]E[p(A)]} \int_0^\infty y^\rho \bar{F}_L(x + y)\, dy$$
$$= \frac{\beta^\alpha}{E[L]E[p(A)]} \int_0^\infty \frac{y^\rho}{(\beta + x + y)^\alpha}\, dy. \quad (34)$$

Define $r(x, y) = y^\rho(\beta + x + y)^{-\alpha}$. Noticing that $\bar{F}_V(0) = 1$, we immediately get

$$\frac{\beta^\alpha}{E[L]E[p(A)]} \int_0^\infty r(0, y)\, dy = 1 \quad (35)$$

and thus

$$\bar{F}_V(x) = \frac{\int_0^\infty r(x,y)\,dy}{\int_0^\infty r(0,y)\,dy}. \tag{36}$$

Applying [10, Eq. 3.194.3] yields

$$\int_0^\infty r(x,y)\,dy = \frac{\Gamma(\rho+1)\Gamma(\alpha-\rho-1)}{(\beta+x)^{\alpha-\rho-1}\Gamma(\alpha)} \tag{37}$$

and therefore

$$\bar{F}_V(x) = \left(1+\frac{x}{\beta}\right)^{-(\alpha-\rho-1)} \tag{38}$$

which is the tail of the Pareto$(\alpha-\rho-1,\beta)$ distribution. ■

Theorem 3 is a striking result indicating that power-law neighbor selection inherits Pareto tails in out-link lifetimes for both positive and negative ρ, with a clear stochastic ordering between them. Values of $\rho \le -1$ result in undefined V since $E[p(A)] = \infty$ and Assumption 1 fails to hold. Therefore, the lightest achievable tails arise when $\rho \to -1$, in which case $F_V(x) \to F_L(x)$ in distribution, i.e., both new arrivals and neighbors behave probabilistically the same. As ρ increases, the tails get heavier and V goes through the well-known uniform selection ($\rho = 0$) and age-proportional ($\rho = 1$). For $\rho > 1$, it is possible to increase $E[V]$ further. However, this may severely impact the in-degree of long-lived peers and overload them with a massive amount of inbound connection requests, some of which they may have to reject. We later develop the necessary mechanisms for analyzing this tradeoff.

III. Message Overhead

Our second contribution is to analyze the edge-replacement process and obtain the rate λ at which neighbors are sought in the system, which provides a platform for understanding resilience and message overhead of the system. Since each search may require substantial network resources (e.g., flooding and/or random walks), minimization of λ might be one of the possible objectives in practice.

A. Edge-Creation Process

In light of Theorem 1, the rest of the paper uses a single variable $V \sim F_V(x)$ to represent the remaining uptime of an out-neighbor. Similarly, we assume that search delays are i.i.d. and replace them with $S \sim F_S(x)$, where $F_S(x)$ is some CDF of a nonnegative random variable.

Define $\delta := V + S$ to be the length of one up/down cycle in Fig. 1, and let $F_\delta(x) = (F_V * F_S)(x)$ be its CDF, where $*$ denotes convolution. Focusing on a single link, define t_j to be the instance when this link gets its jth out-neighbor, where $t_1 = 0$ and $t_j = t_{j-1} + \delta_j$ for $j \ge 2$. Then, suppose that $\{U(t)\}$ is a renewal process whose interrenewal delays are distributed according to $F_\delta(x)$

$$U(t) := \begin{cases} \sum_{j=1}^{\infty} \mathbf{1}_{t_j \in [0,t]}, & t \ge 0 \\ 0, & \text{otherwise.} \end{cases} \tag{39}$$

Note that $U(t)$ counts the number of replacements in $[0,t]$, where the first renewal always occurs at 0 (i.e., $U(0) = 1$). Then, the expected number of outbound connections generated along a single out-link of w in the interval $[0,t]$ is the *renewal function* $u(t) := E[U(t)]$, which can be expressed as [20]

$$u(t) = \begin{cases} 1 + \sum_{r=1}^{\infty} F_\delta^{*r}(t), & t \ge 0 \\ 0, & \text{otherwise} \end{cases} \tag{40}$$

where $F_\delta^{*r}(t)$ is the r-fold convolution of $F_\delta(t)$.

In passive systems, where the failed neighbors are not replaced, the counting process in (39) reduces to $U(t) = \mathbf{1}_{t \ge 0}$.

B. Cost of Active Replacement

First notice that the number of edges generated along each out-link during the lifetime of a user (i.e., in $[0,L]$) is a random variable $U(L)$. Thus, the average number of connections created per join is simply $kE[u(L)]$, which is the only contributor to the churn-related overhead of the system. Informally speaking, this term depends on the number of out-links k, search delay S, and the rate of churn $1/E[V]$ in the out-neighbors.

To understand this better, observe that connections generated by w during its presence in the system can be either *initial* (i.e., during join) or *replacement* (i.e., during out-link repair). This difference can be seen in Fig. 1, which shows two initial and three replacement edges. Denote by

$$\theta := kE[u(L) - 1] \tag{41}$$

the expected number of *replacement* edges thrown by a peer during its lifetime. Then, the average rate at which out-links are created by a live user in the system is

$$\lambda = \frac{k+\theta}{E[L]} = \frac{kE[u(L)]}{E[L]}. \tag{42}$$

The first term $k/E[L]$ is responsible for the initial edges and cannot be minimized unless k is reduced. The second term, i.e., $\theta/E[L]$, is determined by the resilience of out-links and may be controlled by either increasing the tail weight of the lifetime distribution or changing function $p(x)$ to be more aggressively biased towards older peers.

C. Examples

While the expected search delay $E[S]$ plays a major role in out-degree resilience models [14], [31] regardless of its magnitude, it has only a mild impact on link lifetimes and their churn rate, unless it becomes comparable to $E[V]$. Since measurement studies have shown [26] that $E[V]$ is at least 1 h and considering that finding a neighbor should not take longer than 30–60 s, examples below often assume that S is negligible.

Theorem 4: For exponential lifetimes, $\theta = k$ holds for all $p(x)$ and $E[L]$. For heavy-tailed lifetimes and uniform selection, θ is always smaller than k, eventually reducing to 0 as $R \to \infty$. For light-tailed lifetimes and uniform selection, θ is always larger than k.

Proof: For exponential $F_L(x)$, the number of replacements in $[0,t]$ is simply a Poisson random variable with rate $t/E[L]$. Therefore

$$u(t) = 1 + \frac{t}{E[L]} \quad (43)$$

which leads to $E[u(L)] = 2$ and thus θ in (41) reduces to k.

Next, under uniform selection and zero search delay, cycle length δ is the residual lifetime $R \sim F_A(x)$. Then, the derivative of $u(t)$ is

$$\begin{aligned} u'(t) &= \int_0^t \frac{1-F_L(t-y)}{E[L]}\,du(y) \\ &= \frac{1}{E[L]}\left[1+\int_0^t (F_A(t-y) - F_L(t-y))\,du(y)\right]. \end{aligned} \quad (44)$$

Since under heavy-tailed lifetimes $F_A(x) < F_L(x)$ for all $x \geq 0$, it is easy to see from (44) that $u'(t) < 1/E[L]$, indicating that

$$u(t) < 1 + \frac{t}{E[L]}. \quad (45)$$

This directly leads to $E[u(L)] < 2$ and thus $\theta < k$ for all $\alpha > 1$. As the lifetime tail becomes heavier (e.g., Pareto α gets smaller), the number of renewals of size R in $[0,L]$ monotonically reduces. Eventually, as $\alpha \to 1$, residual R converges to infinity almost surely and becomes larger than L with probability 1, i.e., $P(R > L) \to 1$. In these limiting cases, no replacements happen in $[0,L]$ and thus $\theta \to 0$.

Similarly, given $F_A(x) > F_L(x)$ for lighted-tailed lifetimes (e.g., the uniform distribution), we get from (44) that $(u(t))' > 1/E[L]$, which establishes $E[u(L)] > 2$ and thus $\theta > k$. ■

This result shows that by providing users with neighbors whose remaining lifetimes R are stochastically larger than L, Pareto systems exhibit smaller link-related churn and thus lower overhead compared to the exponential case. In the best scenario of $\alpha \to 1$ (i.e., $R \to \infty$ almost surely), the amount of replacement traffic can be reduced to zero, while the total number of neighbor searches shrinks by half compared to exponential lifetimes, i.e., from $k + \theta = 2k$ per user join to k. As measurement studies show $\alpha \approx 1.1$ [1], [26] in real P2P networks, this effect might be achievable in practice.

Under nonuniform selection, conclusions similar to those in Theorem 4 hold, except the reasoning replaces residuals R with out-link lifetimes V, i.e., the more heavy-tailed V, the smaller θ. We can also state that the fraction of all connection requests that come from initial edges is $\pi = k/(k+\theta)$. If this metric is above $1/2$ (i.e., heavy-tailed lifetimes), the system is driven by join overhead. If it equals $1/2$, then we have the exponential case where both types of edges are equally likely. Finally, if π is smaller than $1/2$ (i.e., light-tailed L), then the system is driven by edge failure.

We next compare θ against simulations, which we perform throughout the paper by emulating full graphs with n heterogeneous users, each with its own ON/OFF renewal process. The number of live peers at any time is approximately half the system, i.e., $E[N(t)] \approx n/2$. Simulations run for a sufficient amount of time to make the system stationary and achieve convergence of the metric being measured. Fig. 3 shows that (41) matches simulation results very well for both uniform and age-proportional neighbor selection, remains bounded by k as predicted in Theorem 4, and decreases as shape α becomes smaller. Age-proportional maintains lower θ compared to uniform, achieving $\theta \to 0$ as $\alpha \to 2$, which can be explained by its $V = Z = \infty$ (almost surely) for $\alpha \leq 2$.

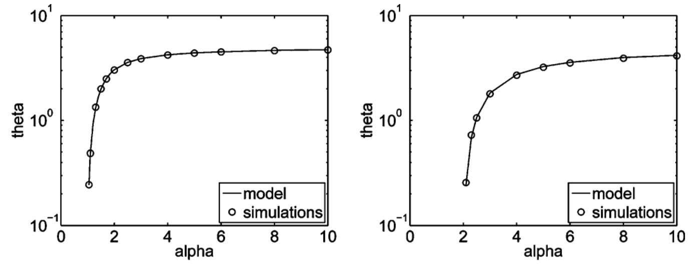

Fig. 3. Model (41) and simulations of θ under Pareto L with $E[L] = 0.5$ h and $k = 5$ (active system). (a) Uniform ($n = 2$ K). (b) Age-proportional ($n = 2$ K).

D. Passive Systems

When out-neighbor failure is ignored, we have $\theta = 0$ and thus $\lambda = k/E[L]$, which represents the optimal case from the overhead standpoint. On the flipside, passive systems throw fewer edges for the same value of k and thus grow their in-degree at a lower rate than active systems. To examine if $E[D_{\text{in}}(\tau) + D_{\text{out}}(\tau)]$ can in fact stay bounded away from zero, we first need to analyze in-link lifetime W, whose distribution, combined with $p(x)$, will eventually determine $E[D_{\text{in}}(\tau)]$.

IV. In-Link Churn

Our third contribution is to derive the distribution and mean of in-link lifetime W, shedding light on its relationship to residuals R of live peers and lifetime L of fresh arrivals. We now focus on node v receiving edges from a random live peer w. Unlike earlier analysis, link (w,v) is considered failed when user w departs, not v.

A. Distribution and Mean

We start with the distribution of W and the average lifetime of in-neighbors $E[W]$, the latter of which also allows us to determine in-link failure rate $\mu = 1/E[W]$.

Theorem 5: The complementary CDF (CCDF) of in-link lifetime is asymptotically

$$\bar{F}_W(x) = \frac{E[u(L-x)]}{E[u(L)]} = \frac{E[u(L-x)|L \geq x]}{E[u(L)]}\bar{F}_L(x) \quad (46)$$

and its mean is

$$E[W] = \frac{E[u(A)]}{E[u(L)]}E[L]. \quad (47)$$

Proof: We first establish (46). Our goal is to determine the fraction of edges created (i.e., selections made) in $[0,t]$ with residual lifetimes of the originating nodes larger than x as $t \to \infty$. Assume fixed n and examine interval $[0,t]$. Place all user sessions entirely contained in this interval (i.e., both created and terminated before t) into set $\mathcal{D}$ (dead) and those still ongoing at t into set $\mathcal{L}$ (live).

For each user w, define $\{A_w(t)\}$ to be the age process of the user (i.e., current age if the user is alive at t, and 0 otherwise) and $\{U_{wi}(t)\}$ to be its renewal process (39) of neighbor selection for the ith link. It then follows that w makes $U_{wi}(L_i - x)$ selections for link i at instances when its own residual is larger than x, which makes the total number of such edges thrown in $[0, t]$ equal to

$$J(x) = \sum_{w\in\mathcal{D}} \sum_{i=1}^{k_w} U_{wi}(L_w - x) + \sum_{w\in\mathcal{L}} \sum_{i=1}^{k_w} r_{wi}(t) \qquad (48)$$

where $r_{wi}(t) = U_{wi}(\min(A_w(t), L_w - x))$ accounts for the edges created by live users within their current age $A_w(t)$.

Note that the bias applied by $p(x)$ determines *how many* of the created edges arrive to v, not *which ones*. Therefore, v is equally likely to receive any of the generated links and the probability that an incoming edge in $[0, t]$ has a lifetime larger than x is simply $J(x)/J(0)$. This shows that among all edges created in $[0, \infty)$, which is equivalent to placing v into a stationary system, the tail distribution of W is

$$\bar{F}_W(x) = P(W > x) = \lim_{t\to\infty} \frac{J(x)}{J(0)}. \qquad (49)$$

Next, since $|\mathcal{D}| \to \infty$ in probability as $t \to \infty$, which is a consequence of the churn model in [28], and $|\mathcal{L}| \le n$, it follows that the second summation term in (48) is upper-bounded by a constant as $t \to \infty$. Then, from the law of large numbers and independence of k_w from L_w

$$\lim_{t\to\infty} \frac{J(x)}{M(t)} = kE[u(L - x)] \qquad (50)$$

where $L \sim F_L(x)$. Combining (49) and (50), we get (46), which immediately leads to (47) following the proof of (11) in Theorem 1. ∎

Interestingly, (46) is very similar to (10), except the tail of W now depends on that of L instead of A and the normalization factor is determined by a monotonically increasing function $u(t)$ instead of $p(x)$. The various cases considered following Theorem 1 apply here as well, i.e., W is stochastically larger than L for heavy-tailed lifetimes and smaller for light-tailed. To perform a sanity check, notice that in passive systems $E[u(L - x)|L \ge x] = E[u(L)] = 1$, which converts (46) to $F_W(x) = F_L(x)$. We get the same result for exponential L in active systems since $L - x$ for $L \ge x$ has the same distribution as L from the memoryless property.

Consistency between model (46) and simulation results for Pareto lifetimes is illustrated in Fig. 4. For uniform selection, Fig. 4(a) shows that W is stochastically smaller than residuals $R \sim F_A(x)$, but larger than $L \sim F_L(x)$. This indicates that peers throwing in-links are less reliable than random live users, but more reliable than fresh arrivals. More interestingly, Fig. 4(b) shows that the tail of W under age-proportional selection is *lighter* than that under uniform selection. This occurs because of the lower churn rate $\theta/E[L]$ in the replacement links and thus a higher fraction of inbound connections coming from

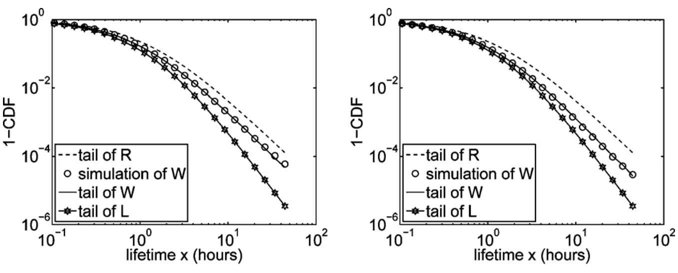

Fig. 4. Comparison of (46) to simulations for Pareto L with $\alpha = 3.5$, mean 0.5 h, and $k = 5$ (active system). (a) Uniform ($n = 2$ K). (b) Age-proportional ($n = 2$ K).

newly joining peers. *Hence, more aggressive functions $p(x)$ reduce message overhead and increase resilience of out-links at the expense of lowering resilience of in-links.* In the worst case, F_W may "deteriorate" down to F_L, which is reminiscent of the situation occurring in DHTs [29], where it happens due to the arrival of new users who take over the zones of existing neighbors.

Turning attention to (47), exponential lifetimes or passive systems produce $E[u(A)] = E[u(L)]$ and therefore $E[W] = E[L]$. For all other cases, (47) is still quite easy to interpret as it shows that $E[W]$ is determined by the ratio of $E[u(A)]$ to $E[u(L)]$, i.e., how many renewal cycles of size δ can fit into $[0, A]$ vs $[0, L]$. For heavy-tailed lifetimes, A is stochastically larger than L and this ratio is above 1, which means that $E[W] > E[L]$ (for light-tailed distributions, the opposite is true). Furthermore, $E[W] = \infty$ whenever $E[A] = \infty$.

B. Discussion

The distribution of W is rather complex because in-links are a combination of initial edges (with lifetime Q) thrown by existing users. It was conjectured in [30] that $Q \sim F_A(x)$ is simply the residual lifetime of w. The rationale for this was that a failed out-edge occurred equally likely within the lifetime of w, and thus w's remaining uptime Q had to follow $F_A(x)$.

Since Q is conditioned on the fact that w's out-link has failed at least once, we easily obtain that the distribution of Q is more heavy-tailed than that of W. However, its relationship to $F_A(x)$ is far from obvious. In our next result, we aim to address this question.

Theorem 6: As $n \to \infty$, the tail distribution of replacement in-link lifetime Q converges to

$$\bar{F}_Q(x) = \frac{E[u(L - x)] - P(L \ge x)}{E[u(L)] - 1}. \qquad (51)$$

Proof: Derivations here are mostly identical to those in Theorem 5, with the only caveat being subtraction of initial edges from (48). Recall that $u(0) = E[U(0)] = 1$ due to the initial edge being thrown when w joins the system. This initial edge needs to be subtracted from $U_{wi}(L_w - x)$, but only for those users whose $L_w \ge x$. This leads to

$$\lim_{t\to\infty} \frac{J(x)}{M(t)} = k(E[u(L - x)] - P(L \ge x)) \qquad (52)$$

which reduces to (51) after canceling k in $J(x)/J(0)$ and replacing $P(L \ge 0)$ with 1 in the denominator. ∎

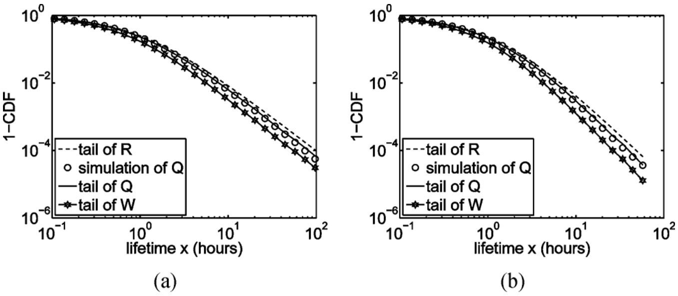

Fig. 5. Comparison of (51) to simulations for Pareto L with mean 0.5 h and $k = 5$ (active system). (a) Uniform ($\alpha = 3$, $n = 2$ K). (b) Age-proportional ($\alpha = 3.5$, $n = 2$ K).

Note that this result is meaningful only for active systems since Q is undefined for networks that do not replace neighbors. It is easy to verify that for exponential L, (51) produces the usual $Q \sim F_L(x)$. However, for lifetimes that exhibit memory, we have yet another distribution that does not equal any of $F_L(x)$, $F_W(x)$, or $F_A(x)$. Fig. 5 shows that the tail of Q is "sandwiched" right between $\bar{F}_W(x)$ and $\bar{F}_A(x)$, i.e., $W <_{st} Q <_{st} R$, where $<_{st}$ means stochastically smaller.

While much of related work [12]–[14], [25], [31] has focused on the lifetime of out-links, it turns out that in-links have a much more interesting and complex behavior. Armed with the distribution of W, we next obtain the in-degree of live users.

V. In-Degree

In our fourth contribution, we examine the aggregate edge-arrival process to a live user v from the rest of the system and obtain the distribution of its in-degree at different ages τ. Recall that outbound connections from w increase the in-degree of other peers in the network. However, this increase is only temporary as all of the established out-links are terminated when w fails at the end of its lifetime. Both active and passive neighbor-replacement models [14], [30] do not impose any limits on the in-degree (i.e., all inbound connections are accepted) and rely on the system to be *self-balancing*, i.e., higher in-degree means faster combined failure of in-neighbors, which should lead to eventual stabilization of in-degree at some finite value. The models developed later in this section help us answer whether this is indeed true.

A. In-Link Arrival Process

Recall that $\lambda = (k+\theta)/E[L]$ is the rate at which users generate outgoing edges. Now, fix a node v and define $\{A_v(t)\}_{t\geq 0}$ to be its *age process*, which is the time elapsed since v's last join into the system if it is alive at t, and 0 otherwise.

Theorem 7: Under Assumption 1 and $n \to \infty$, the arrival process of in-links to v converges in distribution to a nonhomogeneous Poisson process with local rate $\lambda(A_v(t))$, where

$$\lambda(x) := \lambda \frac{p(x)}{E[p(A)]} = \frac{(k+\theta)p(x)}{E[L]E[p(A)]}. \tag{53}$$

Furthermore, the corresponding in-link lifetimes converge in distribution to i.i.d. random variables with CCDF (46).

Proof: At time t, all users in the system are generating out-edges at an overall rate $\lambda N(t)$, which equals the rate at which edges are being absorbed by the live users. If at a given time t the age of user v is $A_v(t) = x$, the selection process and (3) imply that v receives edges at a rate asymptotically proportional to $p(x)$. Thus, the instantaneous rate of in-links arriving to v at t is

$$\lambda(x) = \lambda N(t) \frac{p(x)}{N(t)E[p(A)]} = \frac{\lambda p(x)}{E[p(A)]} \tag{54}$$

in the limit.

These edges are coming from asymptotically independent users, as it is increasingly unlikely that v gets more than one edge from the same user as $n \to \infty$. The in-link arrival process is thus, asymptotically, a superposition of i.i.d. simple point processes, which must be a Poisson process by the traditional point process theory [4]. Finally, asymptotic independence of the users generating these edges implies that of their residual lifetimes, which we know from Theorem 5 follow tail distribution $\bar{F}_W(x)$ in (46). ■

Note that Theorem 7 applies to both passive and active systems, where the only difference arises in parameter θ (i.e., 0 for passive and (41) for active). We next focus on understanding whether users can achieve a balance between arrival of new in-edges and failure of existing ones.

B. In-Degree Distribution

As node connectivity, isolation probability, and routing performance (e.g., coverage during flooding) rely on *transient* properties of node in-degree, we specifically target small age τ in our analysis.

Theorem 8: For a fixed age $\tau \geq 0$, in-degree $D_{\text{in}}(\tau)$ of a live peer v converges in distribution as $n \to \infty$ to a Poisson random variable with mean

$$\nu(\tau) = \int_0^\tau \bar{F}_W(x)\lambda(\tau - x)\, dx. \tag{55}$$

Proof: Notice that we can view arrival of in-edges to each live user v as an $M_t/G/\infty$ queuing system. Specifically, edges find v according to a nonhomogeneous Poisson process with rate $\lambda(x)$ shown in Theorem 7. Furthermore, each in-link can be modeled by a virtual queue whose service time is W. Then, it is well known [6] that the number of busy queues (i.e., live in-neighbors) at fixed time τ has a Poisson distribution with mean (55). ■

While this result shows a clear dependence of $E[D_{\text{in}}(\tau)]$ on the tail of W, an alternative form will be useful later. Substituting (46) and (53) into (55), then expanding λ using (42), yields

$$\begin{aligned}\nu(\tau) &= \int_0^\tau \frac{E[u(L-x)]}{E[u(L)]} \cdot \frac{kE[u(L)]}{E[L]} \cdot \frac{p(\tau - x)}{E[p(A)]}\, dx \\ &= \frac{k}{E[L]E[p(A)]} \int_0^\tau E[u(L-x)]p(\tau - x)\, dx. \end{aligned} \tag{56}$$

In Fig. 6, we plot the distribution of in-degree $D_{\text{in}}(\tau)$ for max-age and age-proportional selection along with a Poisson

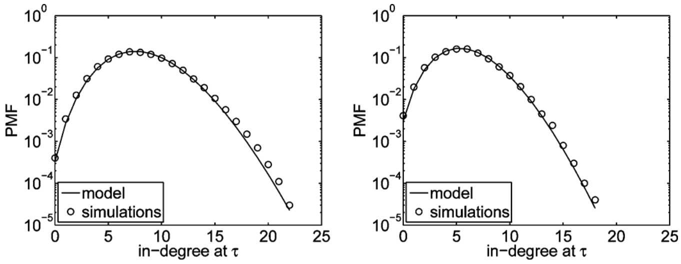

(a) (b)

Fig. 6. Poisson result in Theorem 8 and simulations at $\tau = 1$ h under Pareto lifetimes with $\alpha = 3$, $E[L] = 0.5$, and $k = 8$ (active system). (a) Max-age $m = 5$ ($n = 2$ K). (b) Age-proportional ($n = 2$ K).

distribution with the mean in (56). As the figure shows, the in-degree at given age τ follows the model very well.

C. Examples With Active Systems

The next question relates to our ability to simplify $\nu(\tau)$ under active neighbor replacement (we cover the passive case in Section VI-B). For exponential lifetimes, the CDF of in-neighbor residuals remains the same, i.e., $F_W(x) = F_L(x) = 1 - e^{-x/E[L]}$. From Theorem 4, we have $\lambda = 2k/E[L]$ and thus (55) becomes

$$\nu(\tau) = \frac{2k}{E[L]E[p(L)]}\int_0^\tau \bar{F}_L(x)p(\tau - x)\,dx. \quad (57)$$

Under uniform selection, $p(x) = 1$ yields a scaled CDF of the original lifetime distribution

$$\nu_{\text{unif}}(\tau) = 2k(1 - e^{-\tau/E[L]}) = 2kF_L(\tau) \quad (58)$$

while age-proportional with $p(x) = x$ results in the lifetime CDF being subtracted from a linear function of τ

$$\nu_{age}(\tau) = 2k\left(\frac{\tau}{E[L]} - F_L(\tau)\right). \quad (59)$$

A more captivating case arises when L is not exponential. To expand $\nu(\tau)$ for general lifetimes, we need the next result, which treats $p(x)$ as a signed measure (i.e., difference between two nondecreasing right-continuous functions). This allows integrals to be taken with respect to $dp(x)$, without forcing $p(x)$ to be differentiable or even continuous.

Theorem 9: For $n \to \infty$, the mean in-degree of a live user v at fixed age $\tau \geq 0$ is given by the Lebesgue–Stieltjes integral

$$\nu(\tau) = \frac{k}{E[p(A)]}\int_0^\tau E[u(A) - u(A - \tau + t)]\,dp(t). \quad (60)$$

Proof: Moving the expectation outside in (56)

$$\nu(\tau) = \frac{k}{E[L]E[p(A)]}E\left[\int_0^\tau p(\tau - x)u(L - x)\,dx\right]. \quad (61)$$

Setting ϕ_z to represent the integral in (61) conditional on $L = z$ and recalling that $p(.)$ is a signed measure, observe

$$\phi_z := \int_0^\tau p(\tau - x)u(z - x)\,dx = \int_0^\tau \int_0^{\tau - x} dp(t)u(z - x)\,dx$$
$$= \int_0^\tau \int_0^{\tau - t} u(z - x)\,dx\,dp(t) = \int_0^\tau \int_{z-\tau+t}^{z} u(y)\,dy\,dp(t)$$
$$= \int_0^\tau \int_0^z (u(y) - u(y - \tau + t))\,dy\,dp(t). \quad (62)$$

Using (62) in (61) leads to

$$\nu(\tau) = \frac{k}{E[L]E[p(A)]}E[\phi_L] \quad (63)$$

where

$$E[\phi_L] = \int_0^\tau E\left[\int_0^L (u(y) - u(y - \tau + t))\,dy\right]dp(t). \quad (64)$$

Applying (4) to (64) twice, first with $g(y) = u(y)$ and then with $g(y) = u(y - \tau + t)$, leads to

$$E[\phi_L] = E[L]\int_0^\tau E[u(A) - u(A - \tau + t)]\,dp(t). \quad (65)$$

Substituting (65) into (63) produces (60). ■

Not surprisingly, max-age does not admit closed-form simplification from any of (55), (56), or (60). However, invoking Theorem 9 for the other three methods does lead to rather interesting expressions. *Note that renewal functions $u(t)$ below depend on $p(x)$ and are thus unique to each formula.*

Theorem 10: The step-function produces in (60)

$$\nu_{\text{step}}(\tau) = \frac{kE[u(A) - u(A - \tau + x_0)]}{1 - F_A(x_0)} \cdot \mathbf{1}_{\tau \geq x_0} \quad (66)$$

uniform selection exhibits $\nu_{\text{unif}} = kE[u(A) - u(A - \tau)]$, and age-proportional yields

$$\nu_{\text{age}}(\tau) = k\left(\tau\frac{E[u(A)]}{E[A]} - E[u(Z) - u(Z - \tau)]\right). \quad (67)$$

Proof: Step function $p(x) = \mathbf{1}_{x \geq x_0}$ is right-continuous and thus satisfies the requirement of the theorem. From the definition of Lebesgue–Stieltjes integrals

$$\int_0^\tau g(t)\,dp(t) = \mathbf{1}_{\tau \geq x_0} \cdot g(x_0) \quad (68)$$

for any integrable function $g(.)$. Combining this observation with $E[p(A)] = 1 - F_A(x_0)$ leads to (66), which also produces $\nu_{\text{unif}}(\tau)$ as a special case with $x_0 = 0$.

For age-proportional, setting $z = t - \tau$ makes (60) into

$$\nu_{age}(\tau) = \frac{k}{E[p(A)]}\int_0^\tau E[u(A) - u(A - z)]\,dz$$
$$= \frac{k}{E[A]}\left(\tau E[u(A)] - \int_0^\tau E[u(A - z)]\,dz\right). \quad (69)$$

Invoking (7) with $a = \tau$ for the integral in (69) and moving $1/E[A]$ inside the parentheses, we get (67). ■

Fig. 7 shows simulations of $\nu(\tau)$, leveraging the simplest available model for each case. The figure demonstrates that the considered models are indeed very accurate, albeit somewhat sensitive to the size of the graph. Specifically, uniform and max-age are accurate for n as small as 2 K. However, the age-proportional case (with its $E[V] = \infty$ for $\alpha = 3$) requires $n = 15$ K to maintain a large-enough pool of long-lived peers for a sufficiently randomized neighbor selection.

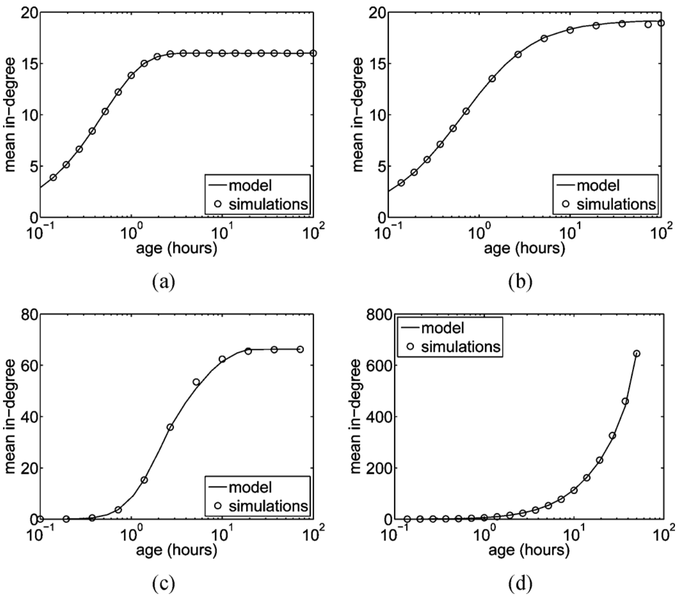

Fig. 7. Comparison of models of $\nu(\tau)$ to simulation results for $E[L] = 0.5$ h, $\alpha = 3$, and $k = 8$ (active system). (a) Uniform/exponential ($n = 2$ K). (b) Uniform/Pareto ($n = 2$ K). (c) Max-age/Pareto ($m = 5$, $n = 2$ K). (d) Age-proportional/Pareto ($n = 15$ K).

Comparing the exponential and Pareto cases in Fig. 7(a) and (b), the latter saturates at a higher value (i.e., 19.1 instead of $2k = 16$) and delivers more edges to long-lived users. This explains its smaller θ and lower overhead discussed earlier. Interestingly, under uniform selection in Fig. 7(a) and (b), saturation point $\nu(\infty) = kE[u(A)] = \lambda/\mu$ is simply the ratio of the rates at which in-links arrive to a user (i.e., λ) and at which they fail (i.e., $\mu = 1/E[W]$). This can be seen using $x_0 = 0$ in (66) and letting $\tau \to \infty$, followed by substitutions from (42) and (47).

The max-age strategy in Fig. 7(c) almost completely ignores small-age peers, but then starts accumulating in-degree at a more healthy pace, surpassing uniform selection by $\tau \approx 1.2$ h and reaching 62 neighbors in 10 h. A similarly interesting case is age-proportional in Fig. 7(d), whose expected degree also starts slow, gaining just 5.7 neighbors in the first hour, but then becomes wildly aggressive, hitting 35 neighbors in 3.7 h and 112 in 10 h. Eventually, $\nu(\tau)$ transitions to a linear function proportional to $k\tau E[u(A)]/E[A]$, reaching the final point in the figure with 645 neighbors in 2 days.

D. Discussion

From (55) and assuming $E[W] < \infty$, uniformly bounded preference functions, i.e., $p(x) \leq M$ for some constant M and all x, lead to finite mean degree $\nu(\infty)$. Likewise, if $p(x)$ is allowed to grow in x to infinity, it follows that $\nu(\tau) \to \infty$ as $\tau \to \infty$. Since the number of connections at each host must be bounded (e.g., due to shortage of sockets, bandwidth, and/or processing power), we arrive at a surprising discovery that age-proportional may lead to peer overload with traffic, rejected connections, and possibly unbounded join delays. In fact, our analysis shows that if selection is made using flooding or random walks, which find nodes in proportion to their degree and thus age, *these strategies may also experience overload and be unsuitable for real networks.*

VI. Combined Degree

Our fifth and final contribution is to analyze the behavior of joint in/out degree, study resilience of the system, and examine various ways to select preference function $p(x)$. In this section, it suffices to consider a single user with k out-links.

A. Active Systems

It is not difficult to see that out-degree $D_{\text{out}}(\tau)$ is a binomial random variable with parameters k and $q(\tau)$, where the latter is the probability that an out-link is live at age τ. Recalling from Fig. 1 and (39) that t_j for $j = 1, 2, \ldots$ is the time instance when the jth neighbor is selected along a single out-link and $V_j \sim F_V(x)$ is the remaining lifetime of the corresponding out-neighbor, we get using integration by parts

$$\begin{aligned} q(\tau) &= E\left[\sum_{j=1}^{U(\tau)} \mathbf{1}_{V_j \geq \tau - t_j}\right] = \int_0^{\tau} \bar{F}_V(\tau - t)\, du(t) \\ &= u(\tau) - E[u(\tau - V)]. \end{aligned} \tag{70}$$

Since the out-degree starts at k, the initial likelihood that the link is active is higher than its stationary equivalent from renewal theory

$$q(\tau) \geq \frac{E[V]}{E[S] + E[V]}. \tag{71}$$

However, as $\tau \to \infty$, $q(\tau)$ monotonically converges to this lower bound. For small mean search delays $E[S] \ll E[V]$, the out-degree may be considered virtually constant and equal to k for all τ, which means that the combined expected degree $E[D(\tau)]$ in active systems is that shown in Fig. 7 with k added to each point.

A simple way to ballpark the operating range of degree is to notice that $D_* = \inf_{\tau \geq 0}\{E[D(\tau)]\}$ is lower-bounded by $kq(\infty)$ and $E[D(\infty)] = kq(\infty) + \nu(\infty)$ represents the maximum achievable number of neighbors. Under uniform selection, exponential L produces $E[D(\infty)] = kq(\infty) + 2k \approx 3k$, while for Pareto lifetimes, this is typically only slightly higher as seen from Fig. 7(b).

Making $p(x)$ more aggressive (e.g., by shifting x_0 in the step-function to larger values) makes V more heavy-tailed, increases resilience of out-links, and reduces their failure rate $\theta/E[L]$, but at the expense of also lowering the resilience of in-links and increasing the degree of high-age peers. Assuming the design calls for lower/upper bounds B_{L} and B_{U} on the expected degree, parameters (k, x_0) of the step-function may be determined by solving $kq(\infty) \geq B_{\text{L}}$ and $\nu(\infty) \leq B_{\text{U}} - B_{\text{L}}$ using the smallest suitable k. In this case, the system is guaranteed to have the maximum resilience and lowest message overhead among all solutions that keep $E[D(\tau)] \in [B_{\text{L}}, B_{\text{U}}]$.

B. Passive Systems

In this case, the network admits closed-form results that do not depend on renewal function $u(x)$. Recall that in passive systems, failure rate $\theta = 0$, in-link lifetimes $W \sim F_L(x)$, $u(X - a) = \mathbf{1}_{X \geq a}$, and $E[u(X - a)] = P(X \geq a)$ for any random

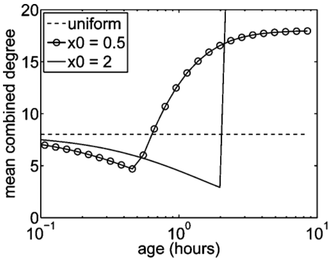

(a)

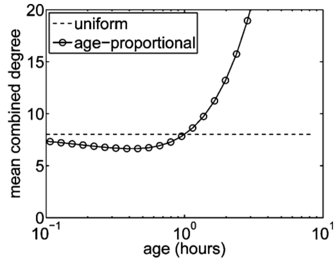

(b)

Fig. 8. Combined expected degree under Pareto lifetimes with $\alpha = 3$, $E[L] = 0.5$, and $k = 8$ (passive system). (a) Step-function. (b) Age-proportional.

variable X. Rewriting (56), recalling that $p(x) = 0$ for $x < 0$, and using Lemma 1, we get

$$\nu(\tau) = \frac{k}{E[L]E[p(A)]} E\left[\int_0^\tau \mathbf{1}_{L \geq x} \cdot p(\tau - x)\, dx\right] = \frac{k}{E[L]E[p(A)]} E\left[\int_0^L p(\tau - x)\, dx\right] = \frac{kE[p(\tau - A)]}{E[p(A)]}$$

which saturates at $\nu(\infty) = kp(\infty)/E[p(A)]$. This shows that *unbounded functions* $p(x)$ *may be unsuitable in practice not just for active, but also passive, systems*.

Simplifying (66), we get for the step-function

$$\nu_{\text{step}}(\tau) = \frac{kF_A(\tau - x_0)}{1 - F_A(x_0)} \quad (72)$$

which leads to the uniform case $\nu_{\text{unif}} = kF_A(\tau)$ via $x_0 = 0$. Expanding (67) results in

$$\nu_{\text{age}}(\tau) = k\left(\frac{\tau}{E[A]} - F_Z(\tau)\right). \quad (73)$$

For exponential L, active P2P networks generate outbound links at double the rate of passive systems, i.e., at $(k + \theta)/E[L] = 2k/E[L]$ compared to $k/E[L]$. It is then reasonable that their expected in-degree in (58) and (59) is similarly larger (i.e., exactly by a factor of two) than in passive systems (72) and (73). Unfortunately, the same conversion factor $(k + \theta)/k$ does not hold for nonexponential L, where link dynamics are significantly more complex.

The expected *out-degree* in passive networks is also very simple and equals the mean number of original neighbors whose residual V is at least τ, i.e., $E[D_{\text{out}}(\tau)] = k\bar{F}_V(\tau)$. Uniform selection combined with its $V \sim F_A(x)$ produces $D(\tau) = k\bar{F}_A(\tau) + kF_A(\tau) = k$ for all τ. The other two cases allow the combined degree to dip below k, but then recover and eventually exhibit $D(\infty) = 0 + D_{\text{in}}(\infty)$. This translates into a limit equal to $k/\bar{F}_A(x_0)$ for the step-function in (72) and ∞ for age-proportional in (73).

Fig. 8 shows this effect in comparison to the uniform case (drawn as a dashed line). Observe that the combined degree of the step-function monotonically decays until $\tau = x_0$ and only then begins to recover. The lowest point of the curve is determined by $F_V(x_0)$ since $D_* = k\bar{F}_V(x_0)$, which is 4.5 and 2.88 for the two cases in Fig. 8(a). The two saturations points are $D(\infty) = k/\bar{F}_A(x_0) = 18$ and 72 neighbors, respectively. The age-proportional case in Fig. 8(b) does not allow the average degree to drop below 6.63, but its $E[D(\tau)]$ increases very aggressively after 1 h and eventually tends to infinity as a linear function $k\tau/E[A]$ in (73). Interestingly, this rate is exactly $E[u(A)]$ times smaller than in the active case (67).

As age-proportional again fails to bound user degree, we next analyze how to use the step-function to achieve $E[D(\tau)] \in [B_{\text{L}}, B_{\text{U}}]$ for all τ. Observe that this can be satisfied with any combination of (k, x_0) such that

$$\begin{cases} D_* = k\bar{F}_V(x_0) \geq B_{\text{L}} \\ D(\infty) = k/\bar{F}_A(x_0) \leq B_{\text{U}}. \end{cases} \quad (74)$$

Additionally, recalling that the message overhead of the system is proportional to $\lambda = k/E[L]$, it makes sense to minimize k among the pairs that conform to (74). In that case, a unique optimal solution emerges as $k = B_{\text{L}}$ and $x_0 = 0$. This shows that *uniform selection minimizes the overhead among all methods that satisfy* $E[D(\tau)] \geq B_{\text{L}}$, where parameter B_{U} actually becomes irrelevant as long as it is no smaller than B_{L}.

The reason why enforcing the lower bound B_{L} is so important *specifically in passive systems* is that their variance in $D(\tau)$ is much higher than in active systems, where $E[S] \ll E[V]$ keeps the degree bounded away from zero. Approximating $D(\tau)$ as a sum of two Poisson variables, it follows that

$$P(D(\tau) = 0) \approx e^{-E[D_{\text{out}}(\tau) + D_{\text{in}}(\tau)]} \leq e^{-D_*} \quad (75)$$

meaning that the higher the mean of $D(\tau)$ at every point τ, the less likely the system is to disconnect. Similarly, maximizing D_* helps improve resilience. In particular, under uniform selection, we get $P(D(\tau) = 0) \approx e^{-k}$, which suggests that 30 neighbors commonly seen in Gnutella are excessive *even for passive systems*, which Gnutella is not. A more reasonable k would be 16 or even 12, with (75) contained below 10^{-7} and 10^{-5}, respectively.

VII. Conclusion

We introduced a novel stochastic framework for tackling link lifetimes and degree evolution in random graphs under churn, covering both passive and active systems under the same umbrella. This work has shown that neighbor-selection mechanisms and the lifetime distribution have a significant impact on the properties of the system, including its message overhead, node resilience to disconnection, and their ability to function as part of the system. We also offered practical guidelines for selecting system parameters and balancing the various tradeoffs.

References

[1] F. E. Bustamante and Y. Qiao, "Friendships that last: Peer lifespan and its role in P2P protocols," in *Proc. Web Content Caching Distrib.*, Sep. 2003.

[2] M. Castro, M. Costa, and A. Rowstron, "Performance and dependability of structured peer-to-peer overlays," in *Proc. IEEE DSN*, Jun. 2004, pp. 9–18.

[3] 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.

[4] D. Daley and D. Vere-Jones, *An Introduction to the Theory of Point Processes, Volume II: General Theory and Structure*. New York, NY, USA: Springer, 2008.

[5] H. David and H. Nagaraja, *Order Statistics*. New York, NY, USA: Wiley, 2003.
[6] S. G. Eick, W. A. Massey, and W. Whitt, "The physics of the $M_t/G/\infty$ queue," *Oper. Res.*, vol. 41, no. 4, pp. 731–742, 1993.
[7] A. Ganesh and L. Massoulié, "Failure resilience in balanced overlay networks," in *Proc. Allerton Conf. Commun., Control, Comput.*, Oct. 2003.
[8] "Gnutella," 2014 [Online]. Available: http://en.wikipedia.org/wiki/Gnutella
[9] B. Godfrey, S. Shenker, and I. Stoica, "Minimizing churn in distributed systems," in *Proc. ACM SIGCOMM*, Sep. 2006, pp. 147–158.
[10] I. Gradshteyn and I. Ryzhik, *Table of Integrals, Series, and Products*, 7th ed. New York, NY, USA: Academic, 2007.
[11] 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.
[12] M. F. Kaashoek and D. Karger, "Koorde: A simple degree-optimal distributed Hash table," in *Proc. IPTPS*, Feb. 2003, pp. 98–107.
[13] S. Krishnamurthy, S. El-Ansary, E. Aurell, and S. Haridi, "A statistical theory of Chord under Churn," in *Proc. IPTPS*, Feb. 2005, pp. 93–103.
[14] D. Leonard, Z. Yao, V. Rai, and D. Loguinov, "On lifetime-based node failure and stochastic resilience of decentralized peer-to-peer networks," *IEEE/ACM Trans. Netw.*, vol. 15, no. 3, pp. 644–656, Jun. 2007.
[15] D. Leonard, Z. Yao, X. Wang, and D. Loguinov, "On static and dynamic partitioning behavior of large-scale P2P networks," *IEEE/ACM Trans. Netw.*, vol. 16, no. 6, pp. 1475–1488, Dec. 2008.
[16] 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.
[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] P. Maymounkov and D. Mazieres, "Kademlia: A peer-to-peer information system based on the XOR metric," in *Proc. IPTPS*, Mar. 2002, pp. 53–65.
[19] 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.
[20] S. Resnick, *Adventures in Stochastic Processes*. Cambridge, MA, USA: Birkhäuser, 2002.
[21] S. Rhea, D. Geels, T. Roscoe, and J. Kubiatowicz, "Handling churn in a DHT," in *Proc. USENIX ATC*, Jun. 2004, pp. 127–140.
[22] S. Saroiu, P. K. Gummadi, and S. D. Gribble, "A measurement study of peer-to-peer file sharing systems," in *Proc. SPIE/ACM Multimedia Comput. Netw.*, Jan. 2002, vol. 4673, pp. 156–170.
[23] R. Serfling, *Approximation Theorems of Mathematical Statistics*. New York, NY, USA: Wiley, 2001.
[24] I. Stoica, R. Morris, D. Liben-Nowell, D. R. Karger, M. F. Kaashoek, F. Dabek, and H. Balakrishnan, "Chord: A scalable peer-to-peer lookup protocol for Internet applications," *IEEE/ACM Trans. Netw.*, vol. 11, no. 1, pp. 17–32, Feb. 2003.
[25] G. Tan and S. Jarvis, "Stochastic analysis and improvement of the reliability of DHT-based multicast," in *Proc. IEEE INFOCOM*, May 2007, pp. 2198–2206.
[26] X. Wang, Z. Yao, and D. Loguinov, "Residual-based estimation of peer and link lifetimes in P2P networks," *IEEE/ACM Trans. Netw.*, vol. 17, no. 3, pp. 726–739, Jun. 2009.
[27] Z. Yao, D. Cline, and D. Loguinov, "In-degree dynamics of large-scale P2P systems," in *Proc. ACM HotMetrics*, Jun. 2010, pp. 37–42.
[28] Z. Yao, D. Cline, and D. Loguinov, "On superposition of heterogeneous edge processes in dynamic random graphs," in *Proc. IEEE INFOCOM Mini-Conf.*, Mar. 2012, pp. 2991–2995.
[29] Z. Yao and D. Loguinov, "Analysis of link lifetimes and neighbor selection in switching DHTs," *IEEE Trans. Parallel Distrib. Syst.*, vol. 22, no. 11, pp. 1834–1841, Nov. 2011.
[30] Z. Yao, D. Leonard, X. Wang, and D. Loguinov, "Modeling heterogeneous user Churn and local resilience of unstructured P2P networks," in *Proc. IEEE ICNP*, Nov. 2006, pp. 32–41.
[31] Z. Yao, X. Wang, D. Leonard, and D. Loguinov, "Node isolation model and age-based neighbor selection in unstructured P2P networks," *IEEE/ACM Trans. Netw.*, vol. 17, no. 1, pp. 144–157, Feb. 2009.

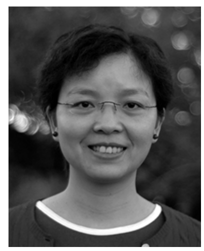

Zhongmei Yao (S'06) received the B.S. degree in engineering from Donghua University, Shanghai, China, in 1997, the M.S. degree in computer science from Louisiana Tech University, Ruston, LA, USA, in 2004, and the Ph.D. degree in computer science from Texas A&M University, College Station, TX, USA, in 2009.

She is currently an Assistant Professor of computer science with the University of Dayton, Dayton, OH, USA. Her research interests include P2P networks, stochastic modeling, and performance analysis.

Dr. Yao has been a student member of the Association for Computing Machinery since 2006.

Daren B. H. Cline received the B.S. degree in mathematics from Harvey Mudd College, Claremont, CA, USA, in 1978, and the M.S. and Ph.D. degrees in statistics from Colorado State University, Fort Collins, CO, USA, in 1980 and 1983, respectively.

He joined the Statistics Department with Texas A&M University, College Station, TX, USA, in 1984 and is now a tenured Professor. His current research interests include applied stochastic processes, Markov chain theory, heavy-tailed distributions, and nonlinear time series models.

Dmitri Loguinov (S'99–M'03–SM'08) received the B.S. degree (with honors) from Moscow State University, Moscow, Russia, in 1995, and the Ph.D. degree from the City University of New York, New York, NY, USA, in 2002, both in computer science.

He is currently a Professor with the Department of Computer Science and Engineering, Texas A&M University, College Station, TX, USA. His research interests include P2P networks, information retrieval, congestion control, Internet measurement, and modeling.