

On the Tradeoff Between Resilience and Degree Overload in Dynamic P2P Graphs

Zhongmei Yao
The University of Dayton
Dayton, OH 45469
Email: zyao@udayton.edu

Daren B.H. Cline
Texas A&M University
College Station, TX 77843
Email: dcline@stat.tamu.edu

Dmitri Loguinov*
Texas A&M University
College Station, TX 77843
Email: dmitri@cs.tamu.edu

Abstract—Prior work has shown that resilience of random P2P graphs under dynamic node failure can be improved by age-biased neighbor selection that delivers outbound edges to more reliable users with higher probability. However, making the bias too aggressive may cause the group of nodes that receive connections to become small in comparison to graph size. As they become overloaded, these peers are forced to reject future connection requests, which leads to potentially unbounded join delays and high traffic overhead. To investigate these issues, we propose several analytical models for understanding the interplay between resilience and degree. We formulate a Pareto-optimal objective for this tradeoff, introduce new metrics of resilience and degree, analyze them under Pareto lifetimes, and discover that traditional techniques can be highly suboptimal in this setting. We then show evidence that optimization can be solved by a family of step-functions, which connect outgoing edges to uniformly random users whose age exceeds some threshold.

I. INTRODUCTION

P2P networks map communication nodes onto random graphs and scale extremely well in terms of bandwidth, storage, processing power, resilience, and resource availability. The P2P paradigm has resulted in diverse applications including online storage systems [20], [24], video streaming [19], [21], conferencing and instant messaging [26], content delivery networks [2], [7], [9], anonymous routing [32], peer trading and e-commerce [1], [33], and mobile ad-hoc networks [6], [16]. Despite being ubiquitous in the current Internet, these networks are still poorly understood, especially under node churn and non-uniform neighbor selection.

A. Motivation

Consider an unstructured P2P graph with n nodes. Upon arrival into the system, suppose users create k outgoing edges using some preference function $w(\tau)$ that selects live peers proportional to their current age τ . When the lifetime distribution $F_L(x)$ allows prediction of residuals from the observed ages (i.e., in all cases except memoryless), the rationale for using age-bias is to allow connectivity to concentrate on more reliable nodes. However, as the target population of users preferred by $w(\tau)$ shrinks in size, they must sustain higher connection-request rates from the rest of the graph. It is therefore possible that an $O(1)$ fraction of the system is eventually forced to handle $\Theta(n)$ load in the network.

This situation not only presents local problems (i.e., overload for certain peers), but also leads to larger global issues. Suppose each user has an upper bound on its operational degree. After this bound is reached, any surplus connection load is rejected and joining users are required to seek other peers that can accept their connections. As $n \rightarrow \infty$ and the degree at each reliable node becomes maxed out with probability $1 - o(1)$, the join delay approaches infinity and the network ceases to function. This is because none of the new users can join and the existing network collapses under an infinite load of redundant join messages.

While special arrangements can be made that alter weight $w(\tau)$ (e.g., revert to uniform) after a number of repeated connection failures, the resulting system is not only difficult for analysis and less insightful, but it may in fact be equivalent to the original system with some other function $w'(\tau)$. Instead, we are interested in the behavior of the pure age-biasing algorithm, i.e., the one driven by the same $w(\tau)$ at all times. Our goal is not to simply prevent scenarios described above, but rather to find a selection strategy that achieves the best resilience for a given degree. Unlike prior work [39] that examined *all* live peers (some of which may never receive connection requests), we focus only on the reliable core, which is a probabilistic artifact of $w(\tau)$.

B. Contributions

Age-biased neighbor preference can be implemented in two general frameworks – active [5], [15], [23], [27], [34] and passive [14], [22]. The former relies on k out-neighbors for routing/resilience and uses edge rewiring upon detecting out-neighbor departure. The latter uses all links bidirectionally, benefitting from both incoming and outgoing edges, and never replaces those that fail. Recent work [39] showed initial evidence suggesting that passive networks were highly appealing as they exhibited good levels of disconnection resilience, in addition to simpler operation and lower overhead compared to active networks [22]. We therefore continue investigation in this direction and limit analysis to these types of P2P graphs.

Our first contribution is to articulate a novel tradeoff between resilience and degree overload, which gives rise to an optimization problem that aims to obtain a Pareto-optimal curve based on the points swept by the two parameters in question. Traditionally [14], resilience was measured by the

*Supported by NSF grants CNS-1017766 and CNS-1319984.

probability that a node lost all of its neighbors before departing from the system; however, this metric is virtually impossible to obtain in closed-form, even for the simplest $w(\tau)$. Instead, we introduce a replacement measure φ that reduces the probability of disconnection to the amount of time a user spends with zero neighbors during its lifetime, which we call the *idle fraction*. We also formally define the age Y of users to whom incoming edges arrive and focus analysis on their expected degree $E[D(Y)]$.

Our second contribution is to study the idle fraction under general $w(\tau)$, where we prove that φ decreases exponentially fast in the number of initial neighbors k . While aggressive bias functions (e.g., step-weight [39], max-age [30], [40]) skew selection towards more reliable peers, they surprisingly fail to *always* achieve better φ than uniform selection (i.e., certain parameter choices make them worse). Our third contribution is to derive a model for $E[D(Y)]$ and examine its growth as age-bias of $w(\tau)$ increases. We show existence of a fundamental tradeoff in the design of P2P networks – achieving $\varphi \rightarrow 0$ always comes at the expense of $E[D(Y)] \rightarrow \infty$. However, the difference between the various preference functions $w(\tau)$ lies in the rate at which $E[D(Y)]$ increases, with some of them being vastly suboptimal compared to the others.

Our last contribution is to analyze the issue of finding the best pair $(k, w(\tau))$ that minimizes φ subject to $E[D(Y)] = d$, where $d > 0$ is a given constant. We use the obtained models to formulate a numerical optimization for this problem and discover that among the strategies considered in this paper step-function $w(\tau) = \mathbf{1}_{\tau \geq x_0}$ comes out as the clear winner. While proving a similar result for an unrestricted class of $w(\tau)$ is a difficult problem, we provide intuition for why we believe this is generally true. Since any nondecreasing and bounded $w(\tau)$ can be represented as a (possibly infinite) sum of step-functions $\mathbf{1}_{\tau \geq x_i}$, Monte Carlo simulations can be used to shed light on the existence of a random mixture that achieves better performance than the step-weight. Our results with over 10K random iterations show that a single step-function indeed forms a Pareto-optimal boundary of the feasible $(E[D(Y)], \varphi)$ space, at least to the extent considered here.

II. RELATED WORK

Analysis of P2P networks under node failure has become a well-studied area [4], [5], [10], [11], [12], [13], [14], [15], [17], [18], [23], [27], [30], [35], [36], [40], [38], [39]. While initial papers built upon Poisson arrivals and exponential lifetimes [4], [13], [17], [27], [35], later work took into account measurements from real P2P systems [3], [20], [28], [29], [36] and examined their performance under non-Poisson arrivals and heavy-tailed lifetimes [10], [14], [30], [20], [39], [40].

Close to our work are [38] and [39], which studied node in/out degree evolution in P2P systems under user churn. In [38], a heterogeneous user ON/OFF model was considered. When each user returned into the network, it connected to a number of out-degree neighbors using uniform selection. Upon detecting neighbor failure, the user sought a replacement in

order to remain connected to the system. This *active* neighbor-management model was used in a number of P2P designs for handling node failure [5], [15], [23], [27], [34]. It was shown in [38] that the edge-arrival process to any live user was asymptotically Poisson with a fixed rate λ .

More recent work [39] also covered *passive* systems, where users searched for neighbors only when joining the network, and introduced generic age preferences, under which any live node with age $\tau \geq 0$ was selected by other peers with a probability asymptotically proportional to its current weight $w(\tau)$. Results showed that passive systems were surprisingly resilient due to the constant arrival of in-edges from other users in the system. While [39] derived node degree $D(\tau)$ at fixed age τ , it did not address resilience, degree of the users that received edges, or the tradeoff between the two.

Other objectives have been considered in prior work – optimizing churn rate [10], node isolation probability [40], connectivity [15], data delivery ratio [30], retrieval latency in P2P storage systems [20], and content search [8], [25], [31] – but they are orthogonal to our results.

III. UNDERSTANDING THE TRADEOFF

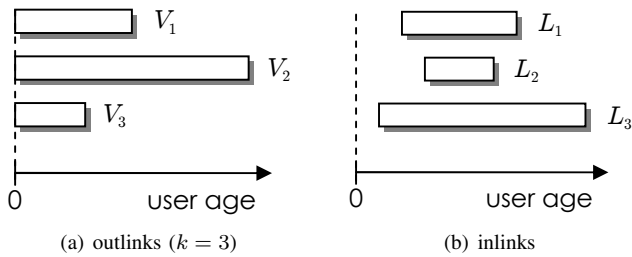
In this section, we first explain our model of decentralized P2P systems and then define our objectives.

A. Dynamic Graph Model

User join/departures are modeled by stationary ON/OFF renewal processes $\{Z_i(t)\}_{i=1}^n$, where n is the number of peers participating in the system. Each user randomly transitions between ON (online) and OFF (offline) states, where duration of ON cycles are called *lifetimes*. Even though each peer may have its own lifetime distribution, [38] shows that such heterogeneous systems can in fact be reduced to homogenous, where all users have the same lifetime CDF (cumulative distribution function) $F_L(x)$.

Upon join, each node v performs k neighbor searches under the general age-preference model of [39]. Specifically, each live peer i with age A_i is selected proportionally to $w(A_i)$, where $w(\tau)$ is some non-negative weight function. Nodes chosen using this approach are called *out-neighbors* with respect to v . Similarly, as v stays in the system, it receives incoming edges from other continuously joining users, which are *in-neighbors* from v 's perspective. Active systems replace only out-links, ignoring failure of in-links, while passive systems replace neither. Since passive networks have many benefits (i.e., low traffic overhead, simplicity of design, predictable performance), we model them below.

Define $V \sim F_V(x)$ to the random lifetime of an out-neighbor. As illustrated in Fig. 1(a), each user starts with k initial out-links, whose lifetimes V_1, \dots, V_k are the remaining delays until departure of the corresponding peers. Define $L \sim F_L(x)$ to be the lifetime of a random peer and observe in Fig. 1(b) that incoming links of duration L_1, L_2, \dots randomly appear and disappear throughout v 's online session. The number of these edges at any time may exceed k and even become unbounded as v 's age increases.

Fig. 1. In/out-link lifetimes at user v in the passive model.

Define $A \sim F_A(x)$ to be the age of L , whose distribution is well-known from renewal theory [37]:

$$F_A(x) := P(A < x) = \frac{1}{E[L]} \int_0^x (1 - F_L(y)) dy. \quad (1)$$

If we examine live users at some time t in a stationary system, the *age* (i.e., elapsed duration of the current ON cycle up to time t) and *residual* (i.e., remainder if their lifetime) are both distributed according to $F_A(x)$. As given by [39, Theorem 1], the distribution of out-link lifetime can be expressed as a function of $w(\tau)$ and A :

$$F_V(x) := P(V < x) = 1 - \frac{E[w(A - x)]}{E[w(A)]}, \quad (2)$$

where we assume that $E[w(A)] < \infty$. Equipped with (2), notice that the out-degree $D_{out}(\tau)$ of node v at given age $\tau \geq 0$ is a binomial variable with mean:

$$E[D_{out}(\tau)] = k\bar{F}_V(\tau) = \frac{kE[w(A - \tau)]}{E[w(A)]}, \quad (3)$$

where $\bar{F}_V(x) = 1 - F_V(x)$ refers to the complementary CDF of V . Since the arrival process of incoming edges towards v tends to non-homogeneous Poisson [39, Theorem 6] as $n \rightarrow \infty$, we assume the system is sufficiently large and the in-degree $D_{in}(\tau)$ at age τ is a Poisson variable with mean:

$$E[D_{in}(\tau)] = \frac{kE[w(\tau - A)]}{E[w(A)]}. \quad (4)$$

The combined degree at age τ is then defined as $D(\tau) := D_{out}(\tau) + D_{in}(\tau)$.

B. Weight Functions

It has been shown that user lifetimes in real P2P systems are either Pareto [36] or Weibull with shape $a < 1$ [20], both from the class of NWU distributions¹. In such cases, non-decreasing functions $w(\tau)$ place more bias on users with higher age and thus produce stochastically larger V [39]. In passive networks, (3) monotonically decreases in user age τ , which is compensated by a monotonic increase in (4). As a result, users are kept connected though out-links long enough

¹A random variable X is called NWU (new worse than used) if its remaining lifetime at age t is stochastically larger than X , i.e., $P(X - t > x | X > t) \geq P(X > x)$ for all $x, t \geq 0$. If the inequality is reversed, X is called NBU (new better than used). If both inequalities hold simultaneously, X is said to be memoryless (i.e., exponential).

for in-degree to take over, resulting in a well-functioning and efficient system.

For performance-comparison purposes, we examine four categories of non-decreasing weight functions. The first is *uniform* [14], which is given by $w(\tau) = 1$. It is simple to implement and not in danger of causing degree overload, although optimality of its resilience is currently unknown. The second is *max-age* [30], [40], where joining users sample m uniformly random peers and then select the one with the largest age. As shown in [39], this is equivalent to using $w(\tau) = m(F_A(\tau))^{m-1}$. The third is the *step-weight* $w(\tau) = \mathbf{1}_{\tau \geq x_0}$ [39], which selects uniformly among the users whose age is at least τ_0 . This can be viewed as creating a virtual “waiting room,” where nodes with age $\tau < x_0$ are not eligible to receive any in-links.

The fourth category consists of truncated-power functions $w(\tau) = \min((\tau/x_0)^\rho, 1)$, where $x_0 > 0$ and $\rho \geq 0$. These generalize two methods discussed above, as well as one additional approach. Specifically, truncated-power becomes uniform using $\rho = 0$, step-weight using $\rho \rightarrow \infty$, and age-proportional $w(\tau) = \tau$ [40] using $\rho = 1$ and $x_0 = \infty$.

C. Performance Measures

Building decentralized graphs requires taking into account resilience against node failure. Assigning larger weights $w(\tau)$ to long-lived users forms a stable core in the graph; however, this also creates a possibility of severely overloading a small cluster of peers with too much traffic and collapsing the rest of the network that depends on them. We thus aim to create an algorithm that achieves a good level of protection against node failure and partitioning, but with proper load-balancing of edges across the graph.

Suppose T is the random delay before v loses all of its neighbors, which is simply the first hitting time of $D(\tau)$ to zero. Then, resilience is usually assessed [14] by the probability that v can remain online without disconnecting, i.e., $\phi = P(T > L)$. In general, however, ϕ is difficult to analyze. This is because $D(\tau)$ is a complex birth-death process, whose first-hitting time to zero is unavailable in closed-form even under uniform selection and exponential lifetimes [14].

To overcome this setback, we treat $D(\tau)$ as a non-absorbing process and approximate ϕ by the expected fraction of time

$$\varphi = \frac{1}{E[L]} E \left[\int_0^L P(D(\tau) = 0) d\tau \right] \quad (5)$$

it spends in state 0 within a user’s lifetime L . As shown in [14], the probability of disconnection is closely related to the recurrence time between visits of $D(\tau)$ to state 0, which in turn is related to φ (i.e., larger ϕ implies larger φ and vice versa). Given this monotonic relationship between the two metrics, we can interchangeably use them for our tradeoff analysis.

As for load-balancing, denote by Y the age of a random user *at the time it is selected* by v and by $D(Y)$ its degree when an incoming request arrives. With the degree capped at some constant $d > 0$ (e.g., 30 in Gnutella [36]), the fraction of rejected connections is $P(D(Y) \geq d)$. In the interest of

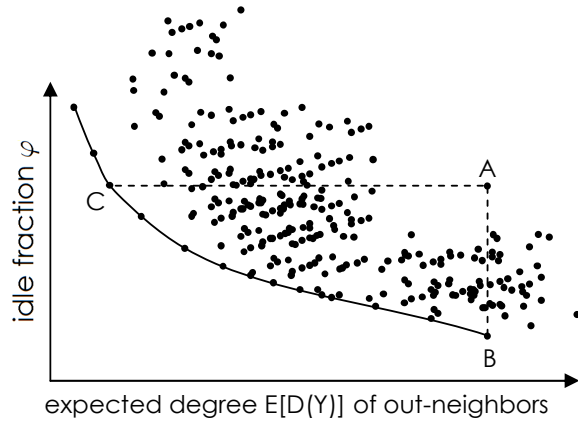


Fig. 2. Pareto-optimal curves in two dimensions.

simplifying presentation, we replace this probability with the expected node degree $E[D(Y)]$, which in many cases also has a monotonic relationship with the metric in question.

D. Optimization

To build intuition, consider Fig. 2, where each point signifies a system driven by some combination $(k, w(\tau))$. We allow fractional k since it represents the mean out-degree upon join [38], which is not restricted to integers. In this 2D space, assume that point (x_1, y_1) dominates (x_2, y_2) if either $x_1 < x_2, y_1 \leq y_2$ or $x_1 \leq x_2, y_1 < y_2$. For example, point A in the figure is dominated by both B and C . Then, define a point to be *Pareto-optimal* if it is dominated by no other point.

As the number of points tends to infinity and assuming the function mapping $(k, w(\tau))$ to the (degree, resilience) pair is sufficiently smooth, Pareto-optimal points merge into a curve, whose piece-wise linear approximation is shown in the figure. The goal of our optimization problem is to obtain the best preference function $w(\tau)$ that places points *only* along the Pareto-optimal curve. This can be rephrased as finding the lowest φ for a given $E[D(Y)] = d$.

IV. IDLE FRACTION

In this section, we derive the idle fraction and analyze the impact of various system parameters.

A. General Case

The next result shows that φ can be expressed as the probability that a random live node is currently idle, i.e., its degree is zero. While we consider non-resilience uses of φ beyond the scope of the paper, it should be noted that this probability may be useful to the analysis of routing performance as it determines the fraction of the graph that is unable to forward messages at any given time. As byproduct of that analysis, one could also obtain the load on the remaining users that must support the full overhead of the system.

Theorem 1: The idle fraction can be written as:

$$\varphi = P(D(A) = 0), \quad (6)$$

where A is the age of a random live node in the graph.

Proof: Denote by $z(y)$ the expected idle duration conditioned on $L = y$:

$$z(y) := \frac{1}{E[L]} \int_0^y P(D(\tau) = 0) d\tau. \quad (7)$$

Using integration by parts, this leads to:

$$\begin{aligned} \varphi &= - \int_0^\infty z(y) d\bar{F}_L(y) \\ &= -z(y)\bar{F}_L(y) \Big|_0^\infty + \int_0^\infty \bar{F}_L(y) dz(y) \\ &= \frac{1}{E[L]} \left(\int_0^\infty \bar{F}_L(\tau) P(D(\tau) = 0) d\tau \right). \end{aligned} \quad (8)$$

Recalling $\bar{F}_L(x) = E[L]f_A(x)$, (8) becomes (6). ■

Suppose $\nu(\tau)$ is the expected in-degree of a node with age τ in a system with $k = 1$:

$$\nu(\tau) := \frac{E[D_{in}(\tau)]}{k} = \frac{E[w(\tau - A)]}{E[w(A)]}. \quad (9)$$

Then, expanding (6) we obtain the next result.

Theorem 2: The idle fraction φ can be written as:

$$\varphi = E \left[\left(F_V(A) e^{-\nu(A)} \right)^k \right]. \quad (10)$$

Proof: Given a fixed age τ , the in/out degrees of a node are independent. As mentioned earlier, the out-degree is binomial with parameters k and $\bar{F}_V(x)$, while the in-degree is Poisson with mean $k\nu(\tau)$. This leads to:

$$\begin{aligned} \varphi &= \int_0^\infty P(D(\tau) = 0) dF_A(\tau) \\ &= \int_0^\infty P(D_{out}(\tau) = 0) P(D_{in}(\tau) = 0) dF_A(\tau) \\ &= \int_0^\infty (F_V(\tau))^k e^{-k\nu(\tau)} dF_A(\tau), \end{aligned} \quad (11)$$

which is the same as (10). ■

While (10) is fairly complex, we can at least extract the impact of k on its decay rate, as shown next.

Theorem 3: For any $w(\tau)$, the following is a tight bound:

$$\varphi \leq \left(\sup_{\tau \geq 0} F_V(\tau) e^{-\nu(\tau)} \right)^k. \quad (12)$$

Proof: Define $H(x) = F_V(x) e^{-\nu(x)}$ and observe that $E[H(A)^k]^{1/k}$ is monotonically non-decreasing as $k \rightarrow \infty$. Therefore:

$$\varphi^{1/k} = E[H(A)^k]^{1/k} \leq \lim_{k \rightarrow \infty} E[H(A)^k]^{1/k}, \quad (13)$$

where the limit is the essential supremum of $H(A)$, which can be shown to equal $\sup_{\tau \geq 0} H(\tau)$. ■

Observe that (12) does not impose any conditions on $w(\tau)$ or $F_A(x)$. Setting

$$c := \sup_{\tau \geq 0} F_V(\tau) e^{-\nu(\tau)} < 1, \quad (14)$$

it follows that the idle fraction $\varphi \leq c^k$ must decrease at least exponentially fast in k . Furthermore, this upper bound is the best possible (i.e., among those that hold for all k). This follows from the fact that $\varphi^{1/k}$ converges from below to c as $k \rightarrow \infty$. This result will come in handy shortly.

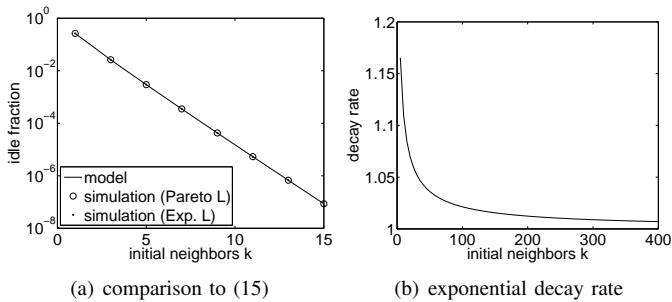


Fig. 3. Idle fraction φ under uniform selection ($n = 5K$ and $E[L] = 0.5$ hours). Pareto lifetimes use $\alpha = 3$.

B. Uniform Weight

We now examine (10) under uniform selection in hopes of obtaining more clarity about the role of k and $F_L(x)$.

Theorem 4: For uniform weights $w(\tau) = 1$:

$$\varphi = \int_0^1 (ye^{-y})^k dy. \quad (15)$$

Proof: Under uniform selection, $\nu(x) = F_A(x)$ and out-link lifetime $V \sim F_A(x)$ [39]. Re-writing (10):

$$\varphi = \int_0^\infty (F_A(x)e^{-F_A(x)})^k dF_A(x), \quad (16)$$

which is the same as (15). ■

While the new expression for φ can be computed more easily than the general expectation (10), it still does not provide much qualitative information on the decay rate towards zero. We obtain that next.

Theorem 5: For uniform selection and all lifetime distributions, φ is tightly upper-bounded by:

$$\varphi \leq e^{-k}. \quad (17)$$

Proof: Since ye^{-y} monotonically increases in $y \in [0, 1]$ and $\nu(\tau) = F_A(\tau)$, we get from (12) that

$$\varphi \leq (F_A(\infty)e^{-F_A(\infty)})^k = e^{-k}, \quad (18)$$

which is the desired conclusion. ■

While the results in [38], [39] upon which we rely are of asymptotic nature (i.e., $n \rightarrow \infty$), they hold in relatively small graphs too. To demonstrate this effect, we use simulations with n in the thousands and two lifetime distributions – exponential:

$$F_L(x) = 1 - e^{-\lambda x} \quad (19)$$

and Pareto:

$$F_L(x) = 1 - (1 + x/\beta)^{-\alpha}, \quad (20)$$

where $\alpha > 1$ and $E[L] = \beta/(\alpha - 1)$. The result is shown in Fig. 3(a). Notice from this plot that the exact model (15) is accurate and that φ remains the same for both exponential and Pareto L . As illustrated in Fig. 3(b), the exponential decay rate $-\ln \varphi/k$ starts at 1.17 and eventually converges to 1 as $k \rightarrow \infty$, which was predicted by the tightness of the upper bound e^{-k}

in Theorem 4. Therefore, in cases when certain approximation liberty is tolerated, and especially when k is large, one may use $\varphi \approx e^{-k}$.

In systems that utilize only the out-degree for resilience [14], it is well-known that Pareto $F_L(x)$ improves protection against node failure compared to exponential L . In fact, shape α has a direct impact on V and the resulting disconnection likelihood (i.e., smaller α produces stochastically larger V). Furthermore, results show [14] that passive usage of out-neighbors is insufficient for maintaining a highly fault-tolerant environment (e.g., $\varphi = \phi = 1/(k + 1)$ under exponential lifetimes). This is why these networks must use active neighbor replacement, especially when L is memoryless or light-tailed.

In contrast, (15) shows that uniform selection keeps φ the same in systems with *different* user lifetime distributions. This allows any L to reach the same level of resilience as heavy-tailed lifetimes, which might be important to systems where $F_L(x)$ is unknown, light-tailed, or changing. In addition, the improved resilience achieved through in-degree neighbors is quite non-trivial and in many cases sufficient for keeping the network connected. For instance, $\varphi = 5 \times 10^{-10}$ for $k = 20$ and $\varphi = 2 \times 10^{-14}$ for $k = 30$, which for exponential L are respectively 8 and 12 orders of magnitude better than in networks that utilize only the out-degree (see above).

Results in this subsection confirm previous intuition [39] that passive neighbor management belongs in the realm of feasibility and show that uniform selection is a good baseline algorithm, whose mean combined degree $E[D(\tau)] = k\bar{F}_A(\tau) + kF_A(\tau) = k$ remains fixed throughout a node's lifetime. Next, we examine other alternatives.

C. Non-Uniform Weights

Since we now consider age-biased preference functions, more discussion is required on the properties of typical lifetimes found in real networks, i.e., Pareto [36] and Weibull [20]. For the former case, the distribution of age A remains Pareto, but its shape decreases to $\alpha - 1$:

$$F_A(x) = 1 - (1 + x/\beta)^{1-\alpha}. \quad (21)$$

For Weibull lifetimes L , the CDF is:

$$F_L(x) = 1 - e^{-(x/b)^a}, \quad (22)$$

where $b > 0$ is the scale parameter, $a > 0$ is the shape parameter, mean $E[L] = b\Gamma(1 + 1/a)$, and $\Gamma(x)$ is the gamma function. Note that (22) is an exponential distribution if $a = 1$, heavy-tailed (NWU) if $a < 1$, and light-tailed (NBU) when $a > 1$. Weibull's age is distributed according to:

$$\begin{aligned} F_A(x) &= \frac{1}{E[L]} \int_0^x (1 - F_L(y)) dy = \frac{1}{E[L]} \int_0^x e^{-(y/b)^a} dy \\ &= \frac{1}{\Gamma(1/a)} \int_0^{(x/b)^a} e^{-u} u^{\frac{1}{a}-1} du, \end{aligned} \quad (23)$$

which can be computed using the incomplete gamma function in numerical software packages (e.g., Matlab). Its density is

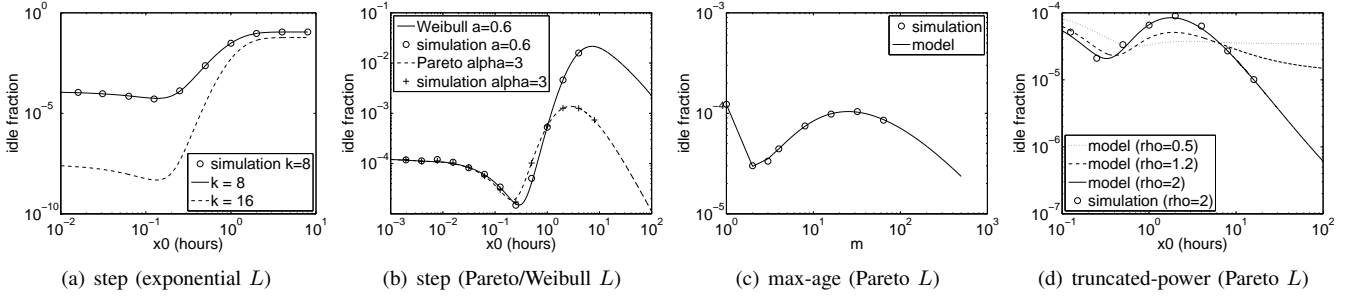


Fig. 4. Examination of idle fraction φ in (10) under various weight functions ($n = 10K$, $\alpha = 3$, $k = 8$, and $E[L] = 0.5$ hours).

available directly as $\bar{F}_L(x)/E[L]$:

$$f_A(x) = \frac{e^{-(x/b)^a}}{b\Gamma(1+1/a)}. \quad (24)$$

We start with the step-weight, where we have from [39]:

$$F_V(x) = 1 - \frac{\bar{F}_A(x+x_0)}{\bar{F}_A(x_0)} \quad (25)$$

and

$$\nu(\tau) = \begin{cases} 0 & \tau \leq x_0 \\ \frac{F_A(\tau-x_0)}{1-F_A(x_0)} & \tau > x_0 \end{cases}. \quad (26)$$

While max-age and truncated power do not allow similar closed-form expansion, the corresponding metrics are easily computed numerically. Recalling both age distributions (21) and (23) yields φ in (10) for all four methods. Fig. 4 shows that simulations under different lifetime distributions and weight functions match model (10) quite well. We next examine each of the subfigures individually.

Starting with Fig. 4(a), observe that φ for exponential L initially decreases in x_0 , but then rebounds towards much higher values. This shows existence of some optimal value, which we call x_0^* , where the step-weight outperforms uniform selection (i.e., $x_0 = 0$) *even in systems with memoryless L* . This intuitively suggests that live nodes with age smaller than x_0^* are sufficiently connected through their out-links and hence do not need improved resilience from in-degree neighbors. Instead, the network benefits by delivering edges to users with larger age whose out-links are more likely to have gone offline. For $x_0 \rightarrow \infty$, users rely entirely on the out-degree and thus:

$$\varphi \rightarrow \int_0^\infty (F_L(x))^k dF_L(x) = \frac{1}{k+1}, \quad (27)$$

which is consistent with our previous discussion of this special case from [14] and the saturation point in Fig. 4(a).

The Pareto and Weibull cases are covered in Fig. 4(b). Notice that φ here exhibits a more pronounced dip compared to exponential L . For $x_0 \in [0, x_0^*]$, the system monotonically improves by obtaining neighbors with progressively larger V , which outweighs the opposite effect applied to in-degree. After crossing the optimal point, the idle fraction reverses direction and then peaks at some x_0^+ . Region $[x_0^*, x_0^+]$ shows that better out-neighbors can no longer compensate for fewer incoming

edges. This relationship is again switched after $x_0 > x_0^+$, where φ monotonically decays towards zero.

Fig. 4(c) illustrates the max-age case. Observe that φ also encounters a dip and a peak, confirming that this behavior is not limited to the step-function. The optimal value here is $m^* = 2$, but this is not always the case (i.e., m^* gets larger for smaller α). Lastly, plot (d) displays φ under truncated power. Interestingly, this method outperforms uniform selection for *all* x_0 , despite having a wobble like the other methods. Smaller values of ρ produce less fluctuation, but their φ either stays constant or decays really slow as $x_0 \rightarrow \infty$.

Complex behavior of φ is peculiar in another aspect – careless choice of parameters (x_0, m) may lead to *lower* resilience than under uniform selection, which underscores the importance of seeking insight into what makes $w(\tau)$ better. While age-biased methods are successful at driving φ to arbitrarily low values (e.g., using $x_0 \rightarrow \infty$), this comes at the expense of severe suboptimality and high degree. For example, using the Pareto case in Fig. 4(b), idle fraction 10^{-5} can be achieved using $x_0 = 0.2$ and $x_0 = 100$. What differentiates the two is the corresponding degree $E[D(Y)]$, which equals approximately 9 in the former case and 100K in the latter. Our tradeoff analysis later in the paper takes exactly these situations into consideration.

D. Asymptotic Decay Rate

Figs. 4(b)-(d) show that NWU lifetimes allow $\varphi \rightarrow 0$ with a sufficiently biased age-preference function. The next result examines this decay rate.

Theorem 6: For Pareto L , step weight $w(\tau)$, and $k > \alpha - 1$, the following ratio converges to a constant:

$$\lim_{x_0 \rightarrow \infty} \frac{\varphi}{\bar{F}_A(x_0)} = (\alpha - 1) \int_0^1 (1 - (1+y)^{1-\alpha})^k y^{-\alpha} dy. \quad (28)$$

Proof: If $k > \alpha - 1$, we have:

$$\begin{aligned} \frac{\varphi}{\bar{F}_A(x_0)} &= \int_0^\infty (F_V(x))^k e^{-k\nu(x)} \frac{dF_A(x)}{\bar{F}_A(x_0)} \\ &= \int_0^\infty (F_V(x_0 y))^k e^{-k\nu(x_0 y)} \frac{x_0 f_A(x_0 y)}{\bar{F}_A(x_0)} dy. \end{aligned} \quad (29)$$

Recalling (25)-(26), observe that:

$$\lim_{x_0 \rightarrow \infty} F_V(x_0 y) = 1 - (1+y)^{1-\alpha} \quad (30)$$

and

$$\lim_{x_0 \rightarrow \infty} \nu(x_0 y) = \begin{cases} 0 & y \leq 1 \\ \infty & y > 1 \end{cases}. \quad (31)$$

Since $\nu(x_0 y) \rightarrow \infty$ for $y > 1$, it suffices to consider the integral in (29) only on the interval $[0, 1]$. To deduce the behavior of the density term in (29), notice that for $y \leq 1$:

$$\lim_{x_0 \rightarrow \infty} \frac{x_0 f_A(x_0 y)}{\bar{F}_A(x_0)} = (\alpha - 1)y^{-\alpha}. \quad (32)$$

Using the fact that all functions inside the integral of (29) are uniformly bounded in x_0 and applying dominated convergence, we get that:

$$\lim_{x_0 \rightarrow \infty} \frac{\varphi}{\bar{F}_A(x_0)} = (\alpha - 1) \int_0^1 (1 - (1 + y)^{1-\alpha})^k y^{-\alpha} dy.$$

The final note is that if $k \leq \alpha - 1$, then the last limit is infinite by Fatou's lemma. ■

This result demonstrates that the idle fraction of the step-weight follows the tail of the age distribution, i.e., $\Theta(x_0^{1-\alpha})$, under Pareto L . The power-law decay rate of φ is quite slow compared to the exponential e^{-k} observed earlier. This suggests that using large $x_0 \rightarrow \infty$ is never beneficial; instead, the optimal technique might be to set x_0 close to x_0^* and then increase k until the desired φ is reached. However, the question of what $w(\tau)$ to use in this algorithm still remains open.

V. DEGREE OF SELECTED USERS

In what follows, we first obtain $E[D(Y)]$ and then analyze its growth rate in more detail.

A. General Case

Suppose Y is the random age of users that receive inbound connections. Note that its distribution depends on $w(x)$ and $F_L(x)$. With that in mind, consider the next result.

Theorem 7: The mean degree of a selected users is:

$$E[D(Y)] = k \frac{E[w(|A_1 - A_2|)w(A_1)]}{(E[w(A)])^2}, \quad (33)$$

where A , A_1 , and A_2 are iid (independent and identically distributed) with distribution $F_A(x)$.

Proof: Given N live users with ages A_1, \dots, A_N , the probability of selecting one with age smaller than τ is:

$$\frac{\sum_{i=1}^N w(A_i) \mathbf{1}_{A_i \leq \tau}}{\sum_{i=1}^N w(A_i)} \quad (34)$$

As $n \rightarrow \infty$, this converges to the CDF of Y :

$$P(Y < \tau) = \frac{E[w(A) \mathbf{1}_{A \leq \tau}]}{E[w(A)]} = \int_0^\tau \frac{w(y) f_A(y)}{E[w(A)]} dy. \quad (35)$$

Differentiating with respect to τ , we get the density of Y :

$$f_Y(\tau) = \frac{w(\tau)}{E[w(A)]} f_A(\tau). \quad (36)$$

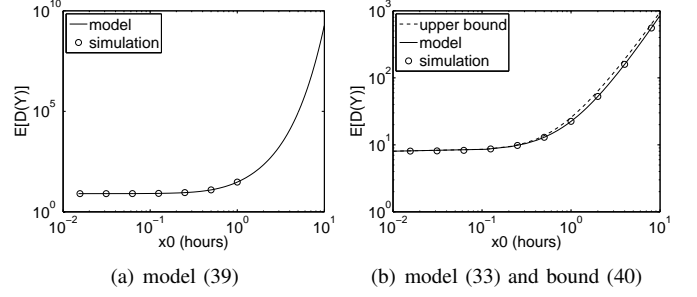


Fig. 5. Degree $E[D(Y)]$ under the step-weight ($n = 10K$, $\alpha = 3$, $k = 8$, and $E[L] = 0.5$ hours).

Invoking (25)-(26) yields:

$$\begin{aligned} E[D(Y)] &= E[D_{out}(Y)] + E[D_{in}(Y)] \\ &= \frac{k \int_0^\infty (E[w(A - y)] + E[w(y - A)]) w(y) f_A(y) dy}{(E[w(A)])^2} \\ &= k \frac{E[w(|A_1 - A_2|)w(A_1)]}{(E[w(A)])^2}, \end{aligned} \quad (37)$$

where we use the fact that $w(x - y) + w(y - x) = w(|x - y|)$ since $w(x) = 0$ for $x < 0$. ■

Theorem 7 indicates that $E[D(Y)]$ is finite when $E[w^2(A)]$ is, which can be seen from:

$$E[D(Y)] \leq k \left(1 + \frac{E[w^2(A)]}{(E[w(A)])^2} \right). \quad (38)$$

B. Asymptotic Growth

Since max-age and truncated power are difficult to simplify, we limit analytical insight in this section to the step function. For exponential L , it is easy to get from (33) that

$$E[D(Y)] = k \cosh\left(\frac{x_0}{E[L]}\right), \quad (39)$$

which increases exponentially in x_0 . One example is shown in Fig. 5(a), where the degree skyrockets to 1 billion neighbors with x_0 just 10 hours. Note that some of the plots in the paper miss simulations for large x_0 , which is a consequence of them being too slow (since we simulate an entire graph) or requiring impossibly large n for the system to be random enough.

We next obtain an upper bound on $E[D(Y)]$ for Pareto L to understand the impact of x_0 in non-memoryless cases.

Theorem 8: For the step-weight and Pareto lifetimes:

$$E[D(Y)] \leq k \left[\left(1 + \frac{x_0}{\beta} \right)^{\alpha-1} - \frac{2(\alpha-1)^2 x_0}{(2\alpha-1)(x_0 + \beta)} \right]. \quad (40)$$

Proof: Given the step function $w(\tau) = \mathbf{1}_{\tau \geq x_0}$, recall that $E[w(A - x)] = \bar{F}_A(x + x_0)$ and $E[w(x - A)] = F_A(x - x_0)$ for $x \geq x_0$. Using the middle equation in (37), we get:

$$\begin{aligned} \frac{E[D(Y)]}{k} &= \int_{x_0}^\infty \frac{1 - F_A(z + x_0) + F_A(z - x_0)}{(1 - F_A(x_0))^2} dF_A(z) \\ &= \frac{1}{1 - F_A(x_0)} - \frac{\int_{x_0}^\infty (F_A(z + x_0) - F_A(z - x_0)) dF_A(z)}{(1 - F_A(x_0))^2}. \end{aligned}$$

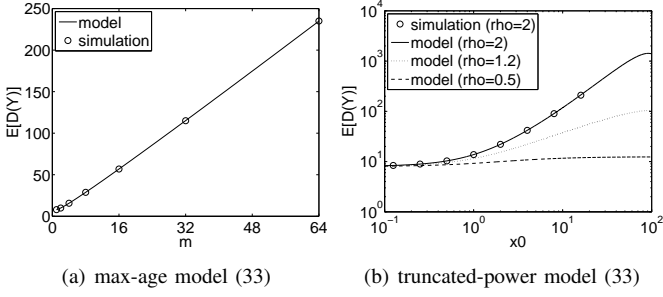


Fig. 6. Degree $E[D(Y)]$ for Pareto lifetimes ($n = 10K$, $\alpha = 3$, $k = 8$, and $E[L] = 0.5$ hours).

Since $f_A(x)$ is convex for Pareto L , we get:

$$F_A(z + x_0) - F_A(z - x_0) = \int_{z-x_0}^{z+x_0} f_A(y) dy \geq 2x_0 f_A(z).$$

Therefore, combining the last two equations:

$$\frac{E[D(Y)]}{k} \leq \frac{1}{1 - F_A(x_0)} - \frac{2x_0 \int_{x_0}^{\infty} f_A^2(z) dz}{(1 - F_A(x_0))^2}. \quad (41)$$

Since $A \sim \text{Pareto}(\alpha - 1, \beta)$, this works out to:

$$\frac{E[D(Y)]}{k} \leq \left(1 + \frac{x_0}{\beta}\right)^{\alpha-1} - \frac{2(\alpha-1)^2 x_0}{(2\alpha-1)(x_0 + \beta)}, \quad (42)$$

establishing this theorem. \blacksquare

As shown in Fig. 5(b), the exact model (33) is very accurate for Pareto L , but the upper bound in (40) is also very close. The increase in degree is not as rapid as for exponential L , but still $E[D(Y)]$ in this case reaches close to 1 thousand neighbors with $x_0 = 10$ hours. It is not difficult to see that after some initial period, the curve becomes a straight line, which indicates a power-law function. From (40), we know it must be quadratic since $\alpha - 1 = 2$. Hence, assuming large x_0 and ignoring the second term in (40), we get:

$$E[D(Y)] \approx k \left(1 + \frac{x_0}{(\alpha-1)E[L]}\right)^{\alpha-1} = O(x_0^{\alpha-1}). \quad (43)$$

The next result is a companion lower bound on $E[D(Y)]$.

Theorem 9: For the step-weight and all lifetimes:

$$E[D(Y)] \geq k F_A(x_0) \frac{1 - F_A(2x_0)}{(1 - F_A(x_0))^2}. \quad (44)$$

Proof: To get a lower bound on (33), write:

$$\begin{aligned} E[w(|A_1 - A_2|)w(A_1)] &= P(|A_1 - A_2| \geq x_0, A_1 \geq x_0) \\ &\geq P(A_1 \geq 2x_0, A_2 \leq x_0) \\ &= F_A(x_0)(1 - F_A(2x_0)), \end{aligned} \quad (45)$$

which leads to the desired result. \blacksquare

For Pareto L with $\alpha > 1$, (44) reduces to:

$$E[D(Y)] \geq k F_A(x_0) \left(1 + \frac{2x_0}{\beta}\right)^{1-\alpha} \left(1 + \frac{x_0}{\beta}\right)^{-2+2\alpha}.$$

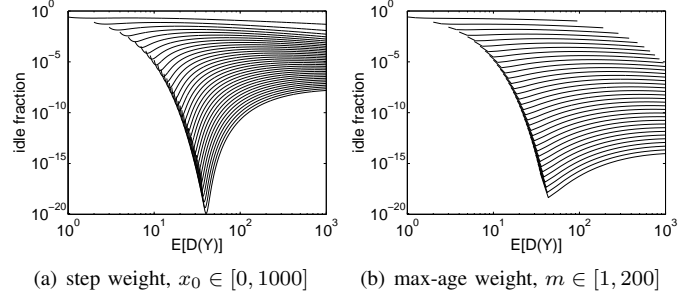


Fig. 7. Resilience-degree tradeoff for Pareto lifetimes ($\alpha = 3$ and $E[L] = 0.5$ hours).

Combining with (43), we reach a conclusion that Pareto L scales $E[D(Y)]$ proportional to the inverse of the tail of the age distribution:

$$E[D(Y)] = \Theta(1/\bar{F}_A(x_0)) = \Theta(x_0^{\alpha-1}). \quad (46)$$

Note that $E[D(Y)] \propto 1/\bar{F}_A(x_0)$ holds for exponential L in (39) and is likely to apply to other families of distributions. In any case, these results show that increasing x_0 beyond some value (e.g., 0.5 hours in Fig. 5) may lead to explosion in $E[D(Y)]$, with the aggressiveness depending on the tail of the lifetime distribution. Therefore, x_0 can be safely tuned within some interval near zero, after which the remainder of resilience must come from an increase in k , whose impact on $E[D(Y)]$ is milder (i.e., linear). This conclusion works well with our earlier analysis of φ , where the optimal x_0^* was in the ballpark of 0.1 – 0.2 hours.

The expected degree under max-age and truncated power is illustrated in Fig. 6. Notice in (a) that $E[D(Y)]$ is linear in m , which is a slower increase than discovered earlier in the section. However, it is still unclear whether the combination $(E[D(Y)], \varphi)$ of this scenario can beat those of the step-function. Subfigure (b) shows that under truncated-power weight, larger ρ leads to higher $E[D(Y)]$, where $E[D(Y)]$ remains bounded in x_0 if $\rho < (\alpha - 1)/2$.

VI. RESILIENCE AND DEGREE-OVERLOAD TRADEOFF

It should be fairly obvious now that minimizing φ and keeping $E[D(Y)]$ to a minimum are conflicting requirements. In this section, we investigate this tradeoff in more depth and offer avenues for finding optimal neighbor-selection strategies.

A. Objective Function

We first discuss two examples to better motivate the problem. Fig. 7(a) uses Pareto L and the step-weight to plot a number of tradeoff curves, each of which is obtained by fixing k and varying x_0 from 0 to 100. As k increases, these curves move in the south-east direction, with the location of the dips (i.e., optimal points x_0^*) shifting right as k increases. Note that the majority of feasible points in this 2D space are vastly suboptimal and observe the emergence of a Pareto-optimal boundary, which is continuous if we treat k as such. Similar results appear under max-age weights shown in Fig. 7(b);

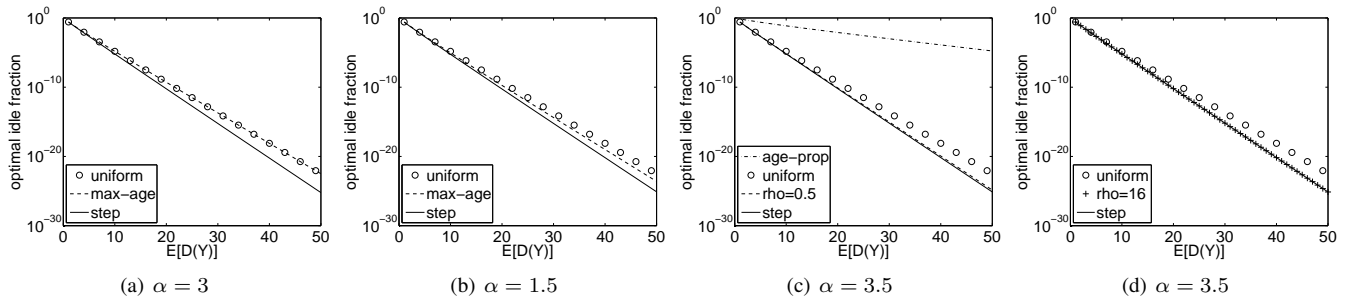
Fig. 8. Pareto-optimal curves for Pareto lifetimes ($E[L] = 0.5$ hours).

TABLE I
OPTIMAL PAIR (x_0^*, k^*) THAT MINIMIZES (49) UNDER THE STEP WEIGHT
AND PARETO LIFETIMES ($E[L] = 0.5$ HOURS)

d	$\alpha = 3$			$\alpha = 2$		
	x_0^*	k^*	$\min \varphi$	x_0^*	k^*	$\min \varphi$
10	0.13	9.2	6.2×10^{-16}	0.17	8.6	5.7×10^{-6}
20	0.14	18.2	5.8×10^{-11}	0.18	16.9	4.8×10^{-11}
30	0.14	27.3	5.7×10^{-16}	0.19	25.2	4.1×10^{-16}
40	0.14	36.4	5.8×10^{-21}	0.20	33.5	3.5×10^{-21}
50	0.14	45.5	6.0×10^{-26}	0.20	41.8	3.1×10^{-26}

however, comparison of its Pareto-optimal curve against that in (a) shows slightly worse results.

Our objective is to determine $w(\tau)$ that achieves the best φ for a given $E[D(Y)] = d$. Unfortunately, neither numerical integration nor simulations can exhaust all non-decreasing functions $w(\tau)$. We therefore first offer techniques for finding the best parameters within a given class of functions and then argue why the found solutions may be optimal across all $w(\tau)$. Suppose θ is some parameter of $w(\tau)$ that we aim to optimize (e.g., x_0 in the step-weight or m in max-age). Then, define:

$$Q(\theta) := \frac{E[D(Y)]}{k} = \frac{E[w(A_1)w(|A_1 - A_2|)]}{(E[w(A)])^2} \quad (47)$$

to be the expected degree under parameter θ and $k = 1$. Since:

$$k = \frac{E[D(Y)]}{Q(\theta)} = \frac{d}{Q(\theta)}, \quad (48)$$

the objective function to minimize is then reduced to:

$$T(\theta) := \int_0^\infty (F_V(x))^{d/Q(\theta)} e^{-dv(x)/Q(\theta)} dF_A(x), \quad (49)$$

which is the value of φ , given d and θ . Running minimization of $T(\theta)$ produces the optimal parameter θ for $w(\tau)$ and yields the best k from (48) as well.

B. Results

Table I lists the optimal pair (x_0^*, k^*) that minimizes (49) under the step weight and Pareto L . Observe that x_0^* increases in d and then converges to some value smaller than $E[L]$. Condition $x_0 < E[L]$ makes sense as it ensures that the reliable core does not contain too few nodes. Further notice in the table that the optimal k^* is slightly smaller than d , with the heavier-tail case (i.e., $\alpha = 2$) needing fewer initial edges due

to its higher resilience of out-neighbors. As a result, it sustains lower join overhead, while achieving better resilience.

We next plot in Fig. 8 the optimal φ subject to $E[D(Y)] = d$ under additional weight functions. Somewhat unexpectedly, part (a) shows that uniform and max-age have the same Pareto-optimal boundaries in this case (i.e., max-age with $m = 1$ performs best). Nevertheless, both methods are worse than the step-weight and the difference grows as d gets larger. By the time d reaches 50, the step-weight offers φ that is 3 orders of magnitude lower. As displayed in part (b), heavier tails of L allow max-age to outperform uniform selection, but the method is still inferior compared to the step-weight, although the gap between the curves is closer.

In Fig. 8(c), we examine age-proportional $w(\tau) = \tau$ (i.e., truncated power function with $\rho = 1$ and $x_0 = \infty$). This approach was used in [40] as a way to stochastically increase V . Later analysis [39], however, found that this method was likely infeasible due to the unbounded degree for users with age $\tau \rightarrow \infty$. The figure confirms this finding and shows that its aggressive bias towards older peers produces an extremely suboptimal result. More interestingly, notice that truncated power with $\rho = 0.5$ performs better than uniform, but still slightly worse than the step-weight.

As shown in plot (d), when ρ becomes large (e.g., $\rho = 16$), truncated power merges with the step-weight. Further increase in ρ yields no improvement, suggesting that the family of truncated power functions can offer no benefit over the step-weight, even though it has an extra parameter ρ . Is it possible that the same conclusion holds for a wider class of $w(\tau)$?

C. General Optimality

Any non-decreasing function $w(\tau)$ can be represented as a (possibly infinite) sum of step-functions $\sum_i \mathbf{1}_{\tau \geq x_i}$. To examine whether there exist such combinations that beat a single step-weight, we generate mixtures of 20 random step-weights and examine the result in Fig. 9. Surprisingly, in over 10K iterations, we are unable to find a combination that works better than a single step-weight. Experiments with random piece-wise linear continuous (PLC) functions produce no counter-examples either. Based on this evidence, we conjecture that there is some inherent optimality to using a single step-weight in solving the degree-resilience tradeoff. Future work will address this issue further, but our current conclusion is that

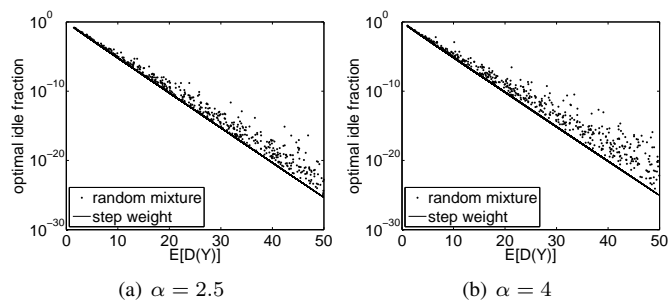


Fig. 9. Performance comparison under the step weight and random mixture functions for Pareto L ($\beta = 1$).

this strategy not only achieves simplicity of implementation and at least partial tractability in modeling, but also optimality within the four classes of considered functions.

VII. CONCLUSION

This paper introduced models of resilience and degree overload in decentralized P2P graphs under node failure. We examined a number of link-creation methods and proved that maximizing resilience alone (i.e., sending $\varphi \rightarrow 0$) without regard to degree provably overloads the reliable users, leading to a non-functional network. Instead, we argued that the problem should be addressed from the perspective of minimizing disconnection under constraints on the mean degree of chosen users. Our findings suggest that a single step-weight may be optimal not just within the class of preference functions considered here, but also outside of it. Further investigation is needed to address this in future work.

REFERENCES

- [1] Bitcoin, "P2P Digital Currency." [Online]. Available: <http://bitcoin.org>.
- [2] BitTorrent. [Online]. Available: <http://www.bittorrent.com>.
- [3] F. E. Bustamante and Y. Qiao, "Friendships that Last: Peer Lifespan and its Role in P2P Protocols," in *Proc. Web Content Caching and Distribution*, Sep. 2003.
- [4] 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.
- [5] 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.
- [6] O. Dousse, P. Thiran, and M. Stoica, "Connectivity in Ad-Hoc and Hybrid Networks," in *Proc. IEEE INFOCOM*, Jun. 2002, pp. 1079–1088.
- [7] eMule. [Online]. Available: <http://www.emule-project.net>.
- [8] R. A. Ferreira, M. K. Ramanathan, A. Awan, A. Grama, and S. Jagannathan, "Search with Probabilistic Guarantees in Unstructured Peer-to-Peer Networks," in *Proc. IEEE P2P*, Aug. 2005, pp. 165–172.
- [9] Gnutella. [Online]. Available: <http://en.wikipedia.org/wiki/Gnutella>.
- [10] B. Godfrey, S. Shenker, and I. Stoica, "Minimizing Churn in Distributed Systems," in *Proc. ACM SIGCOMM*, Sep. 2006, pp. 147–158.
- [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. Networking*, vol. 15, no. 3, 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. Networking*, vol. 16, no. 6, pp. 1475–1488, Dec. 2008.
- [16] J. Li, C. Blake, D. S. J. D. Couto, H. I. Lee, and R. Morris, "Capacity of Ad Hoc Wireless Networks," in *Proc. ACM MOBICOM*, Jul. 2001, pp. 61–69.
- [17] D. Liben-Nowell, H. Balakrishnan, and D. Karger, "Analysis of the Evolution of the Peer-to-Peer Systems," in *Proc. ACM PODC*, Jul. 2002, pp. 233–242.
- [18] P. Maymounkov and D. Mazières, "Kademlia: A Peer-to-Peer Information System Based on the XOR Metric," in *Proc. IPTPS*, Mar. 2002, pp. 53–65.
- [19] P2PTV. [Online]. Available: <http://en.wikipedia.org/wiki/P2PTV>.
- [20] L. Pamies-Juarez, P. Garcia-Lopez, and M. Sanchez-Artigas, "Availability and Redundancy in Harmony: Measuring Retrieval Times in P2P Storage Systems," in *Proc. IEEE P2P*, Aug. 2010.
- [21] PPLive. [Online]. Available: <http://www.pplive.com/>.
- [22] C. Reddy, D. Leonard, and D. Loguinov, "Optimizing Capacity-Heterogeneous Unstructured P2P Networks for Random-Walk Traffic," in *Proc. IEEE P2P*, Sep. 2009, pp. 41–50.
- [23] S. Rhea, D. Geels, T. Roscoe, and J. Kubiatowicz, "Handling Churn in a DHT," in *Proc. USENIX ATC*, Jun. 2004, pp. 127–140.
- [24] 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.
- [25] N. Sarshar, P. O. Boykin, and V. P. Roychowdhury, "Percolation Search in Power Law Networks: Making Unstructured Peer-to-Peer Networks Scalable," in *Proc. IEEE P2P*, Aug. 2004, pp. 2–9.
- [26] Skype. [Online]. Available: <http://www.skype.com>.
- [27] 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.
- [28] D. Stutzbach and R. Rejaie, "Understanding Churn in Peer-to-Peer Networks," in *Proc. ACM IMC*, Oct. 2006, pp. 189–202.
- [29] D. Stutzbach, R. Rejaie, and S. Sen, "Characterizing Unstructured Overlay Topologies in Modern P2P File-Sharing Systems," *IEEE/ACM Trans. on Networking*, vol. 16, no. 2, pp. 267–280, 2008.
- [30] 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.
- [31] W. W. Terpstra, J. Kangasharju, C. Leng, and A. P. Buchmann, "Bubblestorm: Resilient, Probabilistic, and Exhaustive Peer-to-Peer Search," in *Proc. ACM SIGCOMM*, Aug. 2007, pp. 49–60.
- [32] Tor, "Anonymity Online." [Online]. Available: <https://www.torproject.org/>.
- [33] TradePal, "People-Centric Marketplace." [Online]. Available: <http://www.tradePal.com>.
- [34] V. Vishnumurthy and P. Francis, "On Heterogeneous Overlay Construction and Random Node Selection in Unstructured P2P Networks," in *Proc. IEEE INFOCOM*, Apr. 2006.
- [35] M. Vojnovic and A. Proutiere, "Hop Limited Flooding over Dynamic Networks," in *Proc. IEEE INFOCOM*, Apr. 2011, pp. 685–693.
- [36] X. Wang, Z. Yao, and D. Loguinov, "Residual-Based Estimation of Peer and Link Lifetimes in P2P Networks," *IEEE/ACM Trans. Networking*, vol. 17, no. 3, pp. 726–739, Jun. 2009.
- [37] R. W. Wolff, *Stochastic Modeling and the Theory of Queues*. Prentice Hall, 1989.
- [38] Z. Yao, D. Cline, and D. Loguinov, "On Superposition of Heterogeneous Edge Processes in Dynamic Random Graphs," in *Proc. IEEE INFOCOM Mini-Conference*, Mar. 2012, pp. 2991–2995.
- [39] Z. Yao, D. Cline, and D. Loguinov, "Unstructured P2P Link Lifetimes Redux," in *Proc. IEEE INFOCOM*, Apr. 2013, pp. 1762–1770.
- [40] Z. Yao, X. Wang, D. Leonard, and D. Loguinov, "Node Isolation Model and Age-Based Neighbor Selection in Unstructured P2P Networks," *IEEE/ACM Trans. Networking*, vol. 17, no. 1, pp. 144–157, Feb. 2009.