

Understanding Disconnection and Stabilization of Chord

Zhongmei Yao and Dmitri Loguinov*

Department of Computer Science

Texas A&M University, College Station, TX 77843 USA

Email: {mayyao, dmitri}@cs.tamu.edu

Abstract—Previous analytical work [15], [16] on the resilience of P2P networks has been restricted to disconnection arising from simultaneous failure of all neighbors in routing tables of participating users. In this paper, we focus on a different technique for maintaining consistent graphs – Chord’s successor sets and periodic stabilizations – under both static and dynamic node failure. We derive closed-form models for the probability that Chord remains connected under both types of node failure and show the effect of using different stabilization interval lengths (i.e., exponential, uniform, and constant) on the probability of partitioning in Chord.

I. INTRODUCTION

Peer-to-peer (P2P) networks have received tremendous interest in recent years among both Internet users and computer networking professionals. One of fundamental problems in the study of these systems is the ability of the network to stay connected under node failure [1], [3], [7], [9], [10], [14], [15], [19], [21], [23], [25]. While previous analytical work [15], [16] on disconnection of P2P networks has focused on neighbor tables and partitioning arising from failure of entire routing tables, structured P2P networks usually maintain auxiliary sets called *successor lists* [24], [25], whose sole purpose is to recover the system from inconsistent states and provide resilience [25]. In this paper, we focus on partitioning of one particular Distributed Hash Table (DHT) called Chord [25] and note that similar results can be obtained for other types of successor/leaf sets.

Recall that each node v in Chord maintains a list consisting of its $r = \Theta(\log n)$ successors and a routing table containing $k = \Theta(\log n)$ neighbor pointers, where n is the system size. Note that routing tables are used to reduce lookup latency, while successors ensure resilience during churn. Even if all routing tables are in the failed state, Chord is still able to function by forwarding queries, repairing failures, and finding new neighbors via successor lists. When all r successors of any node fail simultaneously, the system becomes partitioned and is potentially unable to recover without a bootstrap.¹ We generally call the event of a user losing all of its successors *node isolation* and note that it determines the likelihood of

graph partitioning:

$$P(\text{graph disconnects}) = P(X > 0), \quad (1)$$

where X is the number of users that are isolated in the system. Due to the strong dependency among successor lists of consecutive users along the circle and entirely different stabilization strategies studied in this paper, previous neighbor churn models [15] cannot be applied to obtain the probability in (1). We perform this task below for both static and dynamic node failure.

A. Static Failure

Many prior studies have been interested in the resilience of structured P2P networks against *static* node failure [9], [10], [25], i.e., when each node independently fails with a certain probability p . We apply the Erdős-Rényi theorem to show that under p -fraction node failure, the probability that Chord with size $n \rightarrow \infty$ remains connected is asymptotically:

$$\lim_{n \rightarrow \infty} \frac{P(X = 0)}{e^{-n(1-p)p^r}} = 1, \quad (2)$$

where $r = \Theta(\log n)$ is the number of immediate successors a user monitors. It is rather surprising to find from (2) that although the dependency among successor lists of consecutive users is very strong, Chord enjoys the same level of static resilience as networks where connectivity is determined using routing tables consisting of largely independent neighbors [16]. Setting $r = c \log_2 n$, where $c > 0$ is a constant, (2) shows that as $n \rightarrow \infty$ the probability that Chord remains connected approaches 1 if $p < 2^{-1/c}$ and 0 if $p > 2^{-1/c}$.

B. Dynamic Failure

As observed in deployed structured P2P file-sharing systems [22], [26], users join and fail at a high rate of churn. The second part of this paper focuses on the connectivity of Chord under dynamic node failure. We assume that each joining user v obtains r clockwise closest peers as its successor list and then stays in the system for L time units, where L is drawn from some user lifetime distribution $F(x)$. User v then stabilizes its successor list every S time units, where S can be random or constant, and brings the number of successors back to r after each stabilization. For a particular stabilization to be successful, at least one user among r successors must stay alive for the entire interval S .

*Supported by NSF grants CCR-0306246, ANI-0312461, CNS-0434940, CNS-0519442, and CNS-0720571.

¹Although neighbors in some routing tables may still be alive, there is no guarantee that the system can return to a consistent state after partitioning.

Assuming exponential user lifetimes L and exponential intervals S , we show that probability ϕ that node v is isolated due to simultaneous failure of its r successors within v 's lifetime is upper bounded by:

$$\phi \leq \frac{\rho \rho! r!}{(\rho + r)!}, \quad (3)$$

where $\rho = E[L]/E[S]$. Furthermore, we prove that as $\rho \rightarrow \infty$, the above upper bound becomes exact.

We then examine how individual node isolation affect partitioning of the system as nodes continuously join and leave. Using the Chen-Stein method [2], we establish that when $r \rightarrow \infty$ the probability that Chord stays connected after experiencing N user joins is asymptotically:

$$\lim_{N \rightarrow \infty} \frac{P(X = 0)}{(1 - \phi)^N} = 1, \quad (4)$$

where ϕ is the node isolation probability given in (3). This result shows that isolations of individual users in Chord can be treated as *independent* when system size and successor lists become large. While a similar phenomenon has been observed in [16] without proof for independent neighbor behavior in routing tables, our result in (4) is again for dependent node isolations and is formally proven.

As (4) indicates that the task of studying global connectivity can be reduced to that of local connectivity, we next focus on isolation probability ϕ under different stabilization strategies. We derive closed-form models of ϕ for uniform and constant S , both of which have been suggested for use in Chord [25]. Our results show that both stabilization strategies are much better than the exponential S suggested in [13], often reducing ϕ by several orders of magnitude. We further show that constant stabilization delays S are *optimal* and keep Chord's isolation probability as $E[S] \rightarrow 0$ approximately equal to:

$$\phi \approx \frac{\rho \rho!}{(\rho + r)!}, \quad (5)$$

where $\rho = E[L]/E[S]$. The amount of improvement over the exponential version (3) of this metric is by a factor of $r!$, which is significant in most cases.

We finish the paper by studying non-exponential lifetimes observed in real P2P graphs [28]. Even though models of ϕ for heavy-tailed user lifetimes are currently intractable, we show that ϕ in such systems is upper bounded by the exponential metric (3). We confirm this effect and demonstrate the distance to the upper bound in simulations.

The rest of the paper is organized as follows. Section (II) provides the basics of Chord and related work on static and dynamic node failure. In Section III, we derive closed-form results on static resilience of Chord. In Section IV, we formalize the successor list model and derive node isolation probability and graph disconnection probability in Chord under dynamic node failure. Section V deals with node isolation probability under different stabilization strategies and finds the optimal method to keep Chord connected with highest probability. In Section VI we show simulation results on isolation probability




Fig. 1. User v 's successors and neighbors in Chord.

for heavy-tailed user lifetimes. Section VII concludes this paper.

II. BACKGROUND

A. Chord Basics

Chord [25] maps each node and key using a uniform hashing function into the identifier (ID) space $\{0, 1, \dots, 2^m - 1\}$, where m is some sufficiently large number that can accommodate all nodes without conflict. Each key is assigned to the *successor* node, i.e., the first peer whose identifier is larger than the key in the clockwise direction along the ring. As illustrated in Fig. 1, each user v in Chord builds a *successor list* and a *finger table*. Assuming n users in the system, the former set contains $r = \Theta(\log n)$ peers immediately following user v along the ring and the latter set consists of $k = \Theta(\log n)$ neighbor pointers where the i -th neighbor is the owner of the key $id(v) + 2^i$.

Finger tables are used during key lookup where the originating node performs jumps of exponentially decreasing length until it finds the node responsible for the key or encounters an inconsistent state (e.g., stale pointer, dead successor) at one of the intermediate nodes. Inconsistent states in finger tables and successor lists are periodically repaired using a *stabilization technique*, which allows Chord to fix links broken during user departure, detect new peer arrival, and ensure lookup success during churn. When any node v leaves the system, its predecessor u notices v 's departure during its periodic stabilization. Peer u then replaces v with the next alive user along the circle and adjusts its successor list accordingly. This process tolerates multiple nodes failing simultaneously and only requires that no successor list sustain a failure of all r nodes within a given stabilization interval. Similarly, node v learns of new arrivals during its stabilization process and properly adjusts its successor list to include the new peers.

Successor lists are generally used in routing only during the last step of a lookup or when all finger pointers corresponding to desired jump lengths have failed. As long as each node has at least one alive peer in its successor list, the system is able to correct (after some delay) all stale finger pointers and re-populate each successor list with r correct entries, thus ensuring consistency and efficiency of subsequent lookups. However, when the entire successor list of any user v fails, that user is considered *isolated* and Chord becomes *partitioned* [25]. Recovery from such disconnection is not guaranteed in the general case.

B. Resilience under Node Failure

Performance of DHTs under p -fraction node failure [9], [10], [25] and churn [6], [13], [17], [18], [19], [20], [23] have received significant attention since the advent of structured P2P networks. While the problem of connectivity under failure for general graphs remains NP-complete [8], [11], [27], recent work [16] shows that several types of deterministic and random networks remain connected if and only if they do not develop isolated nodes after the failure. Despite its importance, the methodology in [16] only considers the resilience of *neighbor tables* rather than that of successors and does not model stabilization. The issues studied in this paper are analytically different due to the much stronger dependency between successor lists of neighboring nodes than between their finger tables and the fact that stabilization requires an entirely different model than the one in [16].

Another modeling work by Krishnamurthy *et al.* [13] studies the probability of finding a neighbor or successor in one of its three states (alive, failed or incorrect) and uses this model to predict lookup consistency and latency for exponential user lifetimes and exponential stabilization intervals.

III. STATIC NODE FAILURE

In this section, we tackle resilience of Chord under *static node failure*, which means that the system sustains a one-time simultaneous failure event where each user becomes dead with an independent probability p . This analysis introduces a new model of handling dependent random events in Chord and can be applied to systems of non-human entities (e.g., file systems) where failures can in fact be synchronized. The next section covers the more typical case of user churn observed in human-based P2P systems.

A. Basic Asymptotic Model

Suppose that Chord is in a consistent state such that each node correctly links to its r closest successors. Under static node failure, p fraction of nodes in the system fail simultaneously, where $0 \leq p \leq 1$ is a given number [9], [10], [16], [25]. Define a Bernoulli random variable X_i indicating whether node i is *isolated* due to the fact that its r successors all fail while i survives:

$$X_i = \begin{cases} 1 & \text{user } i \text{ is alive and its } r \text{ successors failed} \\ 0 & \text{otherwise} \end{cases}. \quad (6)$$

Note that unlike [16], our definition does not involve finger tables since we are only interested in disconnection/isolation arising from disrupted successor lists. Then, the number of isolated nodes X in the system is the sum of a large number of *dependent* random variables X_i :

$$X = \sum_{i=1}^n X_i, \quad (7)$$

where n is the number of nodes in Chord. It is then clear from (1) that the probability that Chord remains connected (i.e., is not partitioned) is equal to $P(X = 0)$. The next theorem

provides an asymptotic closed-form expression of $P(X = 0)$; however, we should note that this result is very different from similar analysis in [16] for two reasons: 1) the model in [16] only considers variables X_i with diminishing dependency as $r \rightarrow \infty$, which is not the case here; 2) the final result on the behavior of X is given in [16] without a formal proof due to a much wider variety of neighbor sets covered by [16].

Theorem 1: The probability that each user in Chord remains connected to at least one successor under p -fraction node failure is asymptotically:

$$\lim_{n \rightarrow \infty} \frac{P(X = 0)}{e^{-n(1-p)p^r}} = 1, \quad (8)$$

where r is the number of successors at each node.

Proof: Denote by a Bernoulli random variable Y_i the event that node i has failed. Then, we have:

$$p = P(Y_i = 1) = 1 - P(Y_i = 0). \quad (9)$$

Define L_n to be the length of the longest consecutive run of 1s in sequence $\{Y_1, \dots, Y_n\}$:

$$L_n = \max_{1 \leq i \leq n-k+1} \{k : Y_i = Y_{i+1} = \dots = Y_{i+k-1} = 1\}. \quad (10)$$

Now notice that computing $P(X = 0)$ can be reduced to finding the distribution of L_n and ensuring that no run longer than $r - 1$ peers exists:

$$P(X = 0) = P(L_n < r). \quad (11)$$

Given that $r = \Theta(\log n)$ so that $r \rightarrow \infty$ as $n \rightarrow \infty$, the distribution of L_n converges to the following based on the Erdős and Rényi law [4]:

$$\frac{P(L_n < r)}{e^{-n(1-p)p^r}} \rightarrow 1, \quad (12)$$

as $n \rightarrow \infty$, which immediately leads to (8). \blacksquare

The asymptotic result in (8) allows us to utilize a very accurate approximation:

$$P(\text{Chord is connected}) = P(X = 0) \approx e^{-n(1-p)p^r}, \quad (13)$$

which we verify next in finite-size graphs. Simulation results of $P(X = 0)$ in Chord under static node failure are presented in Table I. In simulations, each node selects its node ID according to a uniform hashing function and connects to its r successors. After p fraction of users are uniformly randomly chosen and removed, the graph is checked to see how many users X are isolated. Notice from the first three columns in Table I that simulation results with $r = \lceil 2 \log_2 n \rceil$ and $p = 2^{-1/2} = 0.993$ show that as n increases from 1,000 to 10,000, the discrepancy between model (13) and simulation results reduces fast. The rest of the table shows additional examples of model's accuracy for several choices of p and r .

TABLE I
COMPARISON OF SIMULATION RESULTS OF $P(X = 0)$ UNDER STATIC NODE FAILURE TO MODEL (13) IN CHORD

| $p = .933, r = \lceil 2 \log_2 n \rceil$ | | | $n = 50,000, r = \lceil 2 \log_2 n \rceil$ | | | $n = 50,000, r = \lceil 10 \log_2 n \rceil$ | | | $n = 50,000, r = \lceil \sqrt{n} \rceil$ | | |
|--|-------------|-------|--|-------------|--------|---|-------------|-------|--|-------------|--------|
| n | Simulations | (13) | p | Simulations | (13) | p | Simulations | (13) | p | Simulations | (13) |
| 1,000 | .9417 | .9369 | .5 | 1.0000 | 1.0000 | .89 | .9999 | .9999 | .92 | 1.0000 | 1.0000 |
| 5,000 | .9373 | .9360 | .55 | .9999 | .9999 | .9 | .9997 | .9997 | .93 | .9997 | .9997 |
| 10,000 | .9367 | .9360 | .6 | .9983 | .9984 | .91 | .9983 | .9983 | .94 | .9971 | .9971 |
| 20,000 | .9365 | .9360 | .65 | .9821 | .9821 | .92 | .9919 | .9918 | .95 | .9747 | .9747 |
| 30,000 | .9368 | .9367 | .70 | .8472 | .8473 | .93 | .9614 | .9613 | .96 | .8077 | .8076 |
| 40,000 | .9363 | .9361 | .71 | .7771 | .7771 | .94 | .8344 | .8343 | .97 | .1950 | .1954 |
| 50,000 | .9393 | .9393 | .75 | .2850 | .2849 | .95 | .4514 | .4514 | .98 | .0000 | .0000 |
| 100,000 | .9395 | .9394 | .79 | .0038 | .0038 | .96 | .0368 | .0371 | .99 | .0000 | .0000 |

B. Discussion

We next relate our results in Theorem 1 to those in [16, Proposition 3]. Recall that [16] defines isolation as an event of a user losing all of its neighbors in Fig. 1(b). Their results show that all users have at least one alive neighbor with probability:

$$P(X = 0) \approx e^{-n(1-p)p^k}, \quad (14)$$

where n is the system size, p is the independent node failure probability, and k is the number of neighbors in each node's table. Note that we have obtained an almost identical result (13) for successor lists in Chord, which is rather surprising since the dependency among isolation of nodes in Chord is much more significant than assumed in [16] (e.g., node i and node $i + 1$ in Chord share $r - 1$ common successors).

In fact, observe that the probability that node i is isolated due to the failures of its r successors is simply:

$$\phi = P(X_i = 1) = (1 - p)p^r, \quad 1 \leq i \leq n \quad (15)$$

where X_i is the Bernoulli variable defined in (6). Note that given that $r \rightarrow \infty$ as $n \rightarrow \infty$, it is readily seen from (15) that $\phi \rightarrow 0$ as $n \rightarrow \infty$. Using (15), the approximation in (13) can be transformed into:

$$P(X = 0) \approx e^{-n\phi} \approx (1 - \phi)^n, \quad (16)$$

where Taylor expansion $e^{-x} = 1 - x$ holds for small enough x as $n \rightarrow \infty$. Thus, (16) indicates that

$$P(X = 0) = P\left(\bigcap_{i=1}^n [X_i = 0]\right) \approx \prod_{i=1}^n P(X_i = 0) \quad (17)$$

as $n \rightarrow \infty$, which shows that variables X_i in Chord behave *as if* they are completely independent. Note that when $r \rightarrow \infty$ as $n \rightarrow \infty$, node isolations become rare events. Then (17) can be explained by the Chen-Stein theorem [2], which proves that the number of occurrences of dependent rare events X_i is approximately a Poisson random variable under certain conditions (this method will be explicitly used in the next section when we discuss these conditions). Therefore, as $n \rightarrow \infty$, Chord asymptotically exhibits the *same* static resilience using its successor lists composed of largely dependent users as other P2P networks using mostly independent peers in their neighbor sets [16]. However, the rate of convergence of $P(X = 0)$ in (13) and (14) is different.

IV. DYNAMIC NODE FAILURE: GENERAL RESULTS

Recent measurements of P2P networks [5], [22], [26] show that peers continuously join and depart the system, which is often called *churn*. Thus, unlike static node failures which happen simultaneously, node failures in human-based P2P networks often occur dynamically as the system evolves over time. In this section, we first introduce the successor list model under churn, examine probability ϕ that all successors of node v 's fail within its lifetime, and then derive the probability that Chord remains connected when stabilization intervals are exponentially distributed. We leave derivations for non-exponential intervals for the next section.

A. Successor List Model

When each user v joins the system, it acquires a successor list with r nearest nodes and then maintains it through periodic stabilizations (i.e., checks for consistency and dead users). We assume that v does not attempt to track failure of individual users as soon as they occur, but rather performs stabilization every S time units on the entire successor list (i.e., as done in Chord). At each stabilization interval, v corrects its successor list by skipping over failed nodes and appropriately adding to the list new arrivals (if any) [19], which always brings the number of successors at the end of stabilization back to r as long as the system has not been disconnected at some earlier time. For stabilization to be successful, at least one user among r successors must survive the entire stabilization interval. The interval S between two successive stabilizations reflects the duration needed to complete network-related activity to detect failure, exchange neighbor information, and any stabilization rate-limiting applied by the nodes.

Fig. 2 illustrates the evolution of user v 's successor list in our simple model. As shown in this figure, the number of successors is r in the beginning of each stabilization interval of size S . This number then monotonically decreases over time until the next interval starts. If all r successors fail within any interval S before v departs, v is isolated and Chord is disconnected.

In general, as users continuously join and leave the system, the evolution of a node's successor list is rather complicated. It involves not only newly arriving users that replace existing successors, but remaining lifetimes of existing successors at




Fig. 2. Evolution of a node's successor list over time.

the start of each stabilization interval. For exponential user lifetimes, however, user disconnection under this successor-list model becomes tractable as we show next.

Before we proceed with derivations, we introduce the rules for running simulations that verify our theoretical results. In simulations, user arrivals occur according to a Poisson process derived in [29] for the heterogenous churn model proposed therein. The rate of this arrival process is given by $E[N]/E[L]$, where $E[N]$ is the mean system size in equilibrium and $E[L]$ is the mean user lifetime. When a new user joins the system, it is assigned a uniformly random ID in the set $\{0, 1, \dots, 2^{32}-1\}$ and given r immediate successors. Each user then monitors its r successors, stabilizes them every S -interval, and departs from the system after L time units, where L is drawn from some user lifetime distribution $F(x)$.

B. Node Isolation

Denote by $Z(t)$ the number of successors of node v at time t , where $t = 0$ is the time when v joins the system. Note that $Z(0) = r$ and $Z(t) \leq r$ at any age t . In the following, we show that $\{Z(t)\}$ is a Markov chain for exponential user lifetimes and exponential stabilization intervals, which is followed by the derivation of the exact model of node isolation probability ϕ . This exact model is necessary for verifying the accuracy of our later closed-form bounds on ϕ .

Observe from Fig. 2 that state transitions of process $\{Z(t)\}$ are triggered by either failure of existing successors or stabilizations that occur at rate of $\theta = 1/E[S]$. Due to the memoryless property of exponential lifetime distributions, the failure rate of each existing successor (no matter old or new) is $\mu = 1/E[L]$, which is the key reason that makes the successor list tractable for exponential L . This leads to the following lemma.

Lemma 1: For exponential lifetimes $L \sim \exp(\mu)$ and exponential stabilization intervals $S \sim \exp(\theta)$, the process $\{Z(t)\}$ is a continuous-time Markov chain with the state space $\{0, 1, \dots, r\}$ and transition rate matrix $Q = (Q_{jj'})$:

$$Q_{jj'} = \begin{cases} \theta & j \neq r, j' = r \\ j\mu & 1 \leq j \leq r, j' = j - 1 \\ -\theta - j\mu & j' = j < r \\ -j\mu & j = j' = r \\ 0 & \text{otherwise} \end{cases}, \quad (18)$$




Fig. 3. Markov chain $\{Z(t)\}$ modeling a node's successor list.




Fig. 4. Comparison of model (21) to simulation results on node isolation probability ϕ for exponential lifetimes with $E[L] = 0.5$ hours and exponential stabilization intervals with $E[S] = E[L]/\rho$.

where $\theta = 1/E[S]$ and $\mu = 1/E[L]$.

The state diagram and transition rates of process $\{Z(t)\}$ are illustrated in Fig. 3, where each state models the number of alive successors and absorbing state 0 corresponds to user isolation. We usually write matrix Q in (18) in the canonical form:

$$Q = \begin{pmatrix} 0 & 0 \\ \mathbf{r} & Q_0 \end{pmatrix}, \quad (19)$$

where $\mathbf{r} = (q_{j0})^T$ for $j \neq 0$ is a column vector representing the transition rates to the absorbing state 0 and Q_0 is the rate matrix obtained by removing the rows and columns corresponding to state 0 from Q .

Define the first-hitting time T onto state 0 as:

$$T = \inf(t > 0 : Z(t) = 0 | Z(0) = r). \quad (20)$$

Then, the isolation probability $\phi = P(T < L)$ can be reduced to [29, Theorem 11]:

$$\phi = \pi(0)VBV^{-1}\mathbf{r}, \quad (21)$$

where $\pi(0) = (0, \dots, 1)_{1 \times r}$ is the initial state distribution, V is a matrix of eigenvectors of Q_0 , $B = \text{diag}(b_j)$ is a diagonal matrix with:

$$b_j = 1/(\mu - \xi_j), \quad (22)$$

$\mu = 1/E[L]$, $\xi_j \leq 0$ is the j -th eigenvalue of Q_0 , and Q_0 and \mathbf{r} are in (19).

Simulation results of isolation probability ϕ are shown in Fig. 4. Notice from this figure that model (21) is very accurate compared to simulations. Also observe that as ρ or r increase, node isolation probability sharply decreases. While (21) allows easy numerical computation, it provides little

TABLE II
COMPARISON OF THE ASYMPTOTIC MODEL (24) TO THE EXACT MODEL (21) OF NODE ISOLATION PROBABILITY ϕ WITH $E[L] = 0.5$ HOURS, $\rho = E[L]/E[S]$, AND $r = 8$

| ρ | $E[S]$ s | (21) | (24) | Relative Error |
|--------|----------|------------------------|------------------------|----------------|
| 10 | 180 | 1.46×10^{-4} | 2.29×10^{-4} | 57.05% |
| 50 | 36 | 2.30×10^{-8} | 2.61×10^{-8} | 13.41% |
| 100 | 18 | 2.66×10^{-10} | 2.84×10^{-10} | 6.85% |
| 200 | 9 | 2.55×10^{-12} | 2.64×10^{-12} | 3.46% |
| 500 | 3.6 | 4.74×10^{-15} | 4.80×10^{-15} | 1.29% |
| 1,000 | 1.8 | 3.86×10^{-17} | 3.89×10^{-17} | 0.69% |

qualitative information about how ϕ behaves as a function of ρ and r . It is further difficult to compare the various stabilization strategies (studied later in the paper) if an explicit model of ϕ is not derived. We perform this task next.

C. Closed-Form Bounds on ϕ

Note from Fig. 2 that the sequence of stabilization intervals forms a renewal process with cycle length S . It then follows that isolation probability ϕ is equal to the probability that r successors simultaneously fail in *any* interval S before user v 's lifetime expires. Note that the probability that all r successors fail in a particular interval S is given by:

$$f = P(\max\{L_1, \dots, L_r\} < S), \quad (23)$$

where $L_i \sim \exp(\mu)$ is the remaining lifetime of the i -th successor at the beginning of a particular interval. Then, from Jensen's inequality [12, page 118], it is not hard to obtain the following closed-form upper bound on ϕ and prove that it becomes exact as the ratio $E[L]/E[S] \rightarrow \infty$.

Theorem 2: For $L \sim \exp(\mu)$ and $S \sim \exp(\theta)$, isolation probability ϕ is upper-bounded by:

$$\phi < \rho f, \quad (24)$$

where $f = \rho!r! / (\rho + r)!$ and $\rho = E[L]/E[S] = \theta/\mu$. Moreover, the bound becomes tight as stabilization intervals become negligible compared to user lifetimes:

$$\lim_{\rho \rightarrow \infty} \frac{\phi}{\rho f} = 1. \quad (25)$$

The result in (25) indicates that for $\rho \rightarrow \infty$, probability ϕ for any user v to become isolated within its lifetime L can be approximated as the summation of probabilities that v is isolated in each individual interval. Indeed, an average user has approximately $\rho = E[L]/E[S]$ intervals in its lifetime and it gets isolated in any interval with probability f . Thus, since ϕ is asymptotically equal to ρf , isolation events in different intervals behave as if they were independent.

Table II illustrates the relative distance between the upper bound in (24) and the exact result (21) for $E[L] = 0.5$ hours and $r = 8$. It is clear from the table that as ρ increases, the two models converge and that the upper bound is never violated. Also note that other comparisons for different values of $E[L]$ and r exhibit similar results and are omitted for brevity.

We finish this section by examining how individual node isolations affect the connectivity of Chord as users continuously join and depart the system.

D. Graph Disconnection

Notice that Bernoulli variable X_i in (6) can be used to indicate whether user i is isolated due to the failure of its successor list under churn as well. Then node isolation probability can be expressed as:

$$\phi = P(X_i = 1) = 1 - P(X_i = 0), \quad (26)$$

where ϕ is given by (21) or approximated by the upper bound in (24). If user i is isolated during its lifetime, we consider the system disconnected during that user's presence in the system; otherwise, the network is said to *survive* the join of peer i .

Supposing that N users have joined the system, we have that:

$$X_N = \sum_{i=1}^N X_i, \quad (27)$$

is the number of isolations among N join events. In the following, we use the Chen-Stein method [2] to study the probability that Chord survives N user joins without disconnection, i.e., $P(X_N = 0)$. Note that again this result is stronger than that in [16] since it applies to successor lists that exhibit much higher dependency during failure than neighbor lists studied in prior work and relies on more rigorous derivations.

Theorem 3: Given that $N\phi r \rightarrow 0$ as $N \rightarrow \infty$, the probability that Chord survives N user joins without disconnection approaches:

$$\lim_{N \rightarrow \infty} \frac{P(X_N = 0)}{(1 - \phi)^N} = 1, \quad (28)$$

where X_N is defined in (27) and ϕ is given in (21).

Proof: The basic idea of the Chen-Stein method is that the distance between the distribution of X_N , i.e., a sum of N dependent Bernoulli variables, and that of a Poisson random variable of the same mean can be upper-bounded by [2]:

$$|P(X_N = 0) - P(V_N = 0)| \leq \alpha(b_1 + b_2 + b_3), \quad (29)$$

where V_N is a Poisson random variable with mean $E[V_N] = E[X_N] = N\phi$, $\alpha = \min(1, 1/E[X_N])$, and constants b_1 , b_2 and b_3 are defined in [2]. Convergence to the Poisson distribution happens when all of $b_1 - b_3$ tend to zero as $N \rightarrow \infty$. Our main task is to compute these metrics and observe under what condition they become negligibly small.

Define B_i to be a set of users who share at least one successor of user i in Chord:

$$B_i = \{i - r + 1, \dots, i, \dots, i + r - 1\} \quad (30)$$

with $i \in B_i$ and size $|B_i| = 2r - 1$. It follows that $b_3 = 0$ since Bernoulli variable X_i is independent of X_j for $j \notin B_i$. To calculate b_1 , note that:

$$\begin{aligned} b_1 &= \sum_{i=1}^N \sum_{j \in B_i} P(X_i = 1)P(X_j = 1) = \sum_{i=1}^N \sum_{j \in B_i} \phi^2 \\ &= N(2r - 1)\phi^2. \end{aligned} \quad (31)$$

TABLE III

COMPARISON OF MODEL (35) OF $P(X = 0)$ TO SIMULATION RESULTS FOR $r = 8$, MEAN SYSTEM SIZE 2,500, EXPONENTIAL L WITH $E[L] = 0.5$ HOURS, AND EXPONENTIAL S WITH $E[S] = E[L]/\rho$.

| $\rho = 40$ ($E[S] = 45$ s) | | | $N = 50,000$ | | | |
|------------------------------|--------|-------|--------------|----------|--------|-------|
| N | Simul. | (35) | ρ | $E[S]$ s | Simul. | (35) |
| 1,000 | 1.000 | .9999 | 16 | 112.5 | .4831 | .4557 |
| 5,000 | .9996 | .9995 | 24 | 75.0 | .9176 | .9139 |
| 8,000 | .9993 | .9993 | 32 | 56.3 | .9833 | .9829 |
| 10,000 | .9992 | .9991 | 40 | 45.0 | .9954 | .9955 |
| 50,000 | .9954 | .9955 | 48 | 37.5 | .9985 | .9985 |
| 100,000 | .9910 | .9910 | 56 | 32.1 | .9995 | .9994 |
| 500,000 | .9555 | .9556 | 64 | 28.1 | .9998 | .9998 |
| 1,000,000 | .9129 | .9131 | 80 | 25.7 | 1.000 | .9999 |

Likewise, we obtain:

$$\begin{aligned}
b_2 &= \sum_{i=1}^N \sum_{j \neq i, j \in B_i} P(X_i = X_j = 1) \\
&= \sum_{i=1}^N \phi \sum_{j \neq i, j \in B_i} P(X_j = 1 | X_i = 1) \\
&\leq N\phi(2r - 2).
\end{aligned} \tag{32}$$

The last step is to observe that $b_1 = N\phi^2(2r - 1) \rightarrow 0$ and $b_2 \leq N\phi(2r - 2) \rightarrow 0$ as $N \rightarrow \infty$. Finally, given $b_1 + b_2 \rightarrow 0$, it is shown in (29) that X approaches a Poisson random variable with mean $E[X_N]$. This directly leads to:

$$\lim_{N \rightarrow \infty} \frac{P(X_N = 0)}{e^{-E[X_N]}} = \lim_{N \rightarrow \infty} \frac{P(X_N = 0)}{e^{-N\phi}} = 1. \tag{33}$$

Recalling that $\phi \rightarrow 0$ as $N \rightarrow \infty$ given the assumption of this theorem and using Taylor expansion $e^{-\phi} = 1 - \phi$ for $\phi \rightarrow 0$, (33) yields:

$$\lim_{N \rightarrow \infty} \frac{P(X_N = 0)}{(1 - \phi)^N} = 1, \tag{34}$$

which establishes the desired result. ■

Theorem 3 indicates that as long as ϕ is sufficiently small, probability $P(X_N = 0)$ that Chord accommodates N joining users without partitioning simply converges to the product of probabilities that individual nodes remain non-isolated. Note that (28) holds under a wider set of conditions on ϕ that do not necessarily require $N\phi r \rightarrow 0$, but derivations in those cases are more tedious. Also note that a typical way of accomplishing $N\phi r \rightarrow 0$ is to scale r with N so as to converge ϕ to zero faster than product Nr converges to infinity.

Armed with (28), we propose the following approximation to $P(X_N = 0)$ for finite N :

$$P(X_N = 0) \approx (1 - \phi)^N, \tag{35}$$

where the exact model of ϕ is given by (21) and its asymptotic approximation is shown in (24).

Comparison of simulation results of $P(X_N = 0)$ to (35) is presented in Table III where model ϕ is computed based on (21). Notice from the first three columns in this table that simulation results are very close to (35) from $N = 10^3$

TABLE IV

CONVERGENCE OF SIMULATION RESULTS TO MODEL $\phi_u/\phi = .0127$ FROM (37) FOR $E[L] = 0.5$ HOURS, $r = 6$, AND $\rho = E[L]/E[S]$

| ρ | $E[S]$ s | Simulations of ϕ_u | Simulations of ϕ | ϕ_u/ϕ |
|--------|----------|-------------------------|-----------------------|---------------|
| 20 | 90 | 2.15×10^{-6} | 7.10×10^{-5} | .0303 |
| 40 | 45 | 7.59×10^{-8} | 3.86×10^{-6} | .0197 |
| 60 | 30 | 9.98×10^{-9} | 6.10×10^{-7} | .0164 |
| 80 | 22.5 | 2.28×10^{-9} | 1.62×10^{-7} | .0141 |
| 100 | 18 | 7.18×10^{-10} | 5.59×10^{-8} | .0128 |

to 10^6 for $\rho = 40$. The rest of this table shows that as ρ increases (i.e., ϕ gets closer to zero), the model becomes more accurate as expected. Simulations for different r show similar results that are omitted for brevity. As an example of applying (35), assume that Chord has a mean size 5,000 users, $r = \lceil \log_2 5000 \rceil = 13$ successors, $E[L] = 0.5$ hours and $E[S] = 21$ seconds. We then obtain from (35) that the probability that Chord survives $N = 1$ billion user joins without disconnection is 0.999987. If we assume that each user joins and departs the network once per hour, this duration corresponds to 228 years. Furthermore, the system survives for $N = 100$ billion joins (i.e., 22,831 years) with probability 0.998558.

V. DYNAMIC NODE FAILURE: EFFECT OF STABILIZATION INTERVALS

Results in the previous section only apply to exponential intervals S between two consecutive stabilizations. Though many modeling studies assume exponential stabilization intervals [13], [15] to obtain Markovian models, Chord by default uses uniform intervals [25]. In this section, we study isolation probability ϕ for uniform S , deal with ϕ for constant S , and then find the optimal method for stabilizing successors.

A. Uniform Stabilization Delays

Denote by f_u the probability that all r successors of node v fail within interval S where S is uniformly distributed in $[0, 2E[S]]$. Based on the renewal process with cycle length S , it is not hard to show that for uniform S , node isolation probability ϕ_u converges to:

$$\frac{\phi_u}{\rho f_u} \rightarrow 1, \tag{36}$$

as $E[S] \rightarrow 0$, which is similar to the result shown in (25). Then, the ratio of isolation probability ϕ_u for uniform S to ϕ for exponential S is $\phi_u/\phi = f_u/f$, where f is given in (24). Deriving f_u , we obtain the next theorem.

Theorem 4: For fixed r and $E[L]$, and uniform $S \in [0, 2E[S]]$, the ratio of isolation probability ϕ_u for uniform S to ϕ for exponential S converges to the following constant:

$$\lim_{E[S] \rightarrow 0} \frac{\phi_u}{\phi} = \frac{2^r}{(r+1)!}. \tag{37}$$

Simulation results of ϕ_u for uniform S are shown in Table IV. Notice from this table that the ratio ϕ_u/ϕ indeed approaches that given by our model (37) as $E[S]$ becomes small.

TABLE V

CONVERGENCE OF SIMULATION RESULTS TO MODEL $\phi_c/\phi = .0014$ FROM (38) FOR $E[L] = 0.5$ HOURS, $r = 6$, AND $\rho = E[L]/E[S]$

| ρ | $E[S]$ s | Simulations of ϕ_c | Simulations of ϕ | ϕ_c/ϕ |
|--------|----------|-------------------------|-----------------------|---------------|
| 20 | 90 | 2.72×10^{-7} | 7.10×10^{-5} | .0038 |
| 40 | 45 | 8.51×10^{-9} | 3.86×10^{-6} | .0022 |
| 60 | 30 | 9.82×10^{-10} | 6.10×10^{-7} | .0016 |
| 80 | 22.5 | 2.35×10^{-10} | 1.62×10^{-7} | .0015 |
| 100 | 18 | 7.61×10^{-11} | 5.59×10^{-8} | .0014 |

Since $\phi_u \leq \phi$ for all r , the above result demonstrates that using uniform S is a better strategy than using exponential S and that the amount of improvement becomes more significant when r increases, e.g., $\phi_u/\phi = 7.055 \times 10^{-4}$ for $r = 8$ and $\phi_u/\phi = 6.578 \times 10^{-7}$ for $r = 12$.

B. Constant Stabilization Delays

Next, following the derivations of ϕ_u/ϕ in Theorem 4, we easily obtain isolation probability ϕ_c for constant S .

Theorem 5: For fixed r and $E[L]$, and constant S , the ratio of isolation probability ϕ_c to ϕ approaches:

$$\lim_{E[S] \rightarrow 0} \frac{\phi_c}{\phi} = \frac{1}{r!}. \quad (38)$$

Table V presents simulation results on ϕ_c when stabilization intervals are constant. Notice that ratio ϕ_c/ϕ obtained from simulations is very close to that predicted by model (38) even for $\rho = 60$ and that it converges to (38) as ρ increases further. Model (38) indicates that simply stabilizing successors at constant intervals can reduce isolation probability ϕ_c by a factor of $r!$ compared to ϕ as $E[S] \rightarrow 0$. To show the exact improvement over exponential S , we have $\phi_c/\phi = 2.480 \times 10^{-5}$ for $r = 8$ and 2.088×10^{-9} for $r = 12$. In addition, it is easy to notice from (37) and (38) that $\phi_c \leq \phi_u$ and the ratio ϕ_c/ϕ_u approaches $(r+1)/2^r \leq 1$ as $E[S] \rightarrow 0$. This ratio is 0.035 for $r = 8$ and 0.003 for $r = 12$.

C. Optimal Strategy

The above analysis shows that for exponential lifetimes, the ratio of ϕ_c under constant S to ϕ_o under any other S can be transformed into:

$$\lim_{E[S] \rightarrow 0} \frac{\phi_c}{\phi_o} = \frac{P(\max\{L_1, \dots, L_r\} < E[S])}{P(\max\{L_1, \dots, L_r\} < S)}, \quad (39)$$

where $L_i \sim \exp(\mu)$ is the residual lifetime of the i -th successor of node v at the beginning of a particular interval. While we already established that the above ratio is asymptotically less than 1 for both exponential and uniform S , the next theorem indicates that the same result holds for all other distributions as well.

Theorem 6: For exponential user lifetimes with fixed $E[L] > 0$ and the same mean stabilization interval $E[S] \rightarrow 0$, node isolation probability ϕ_c under constant S is no greater than that under any random S .

Proof: For exponential user lifetimes with mean $E[L] = 1/\mu$, recall that the probability that all r successors of node v fail within a particular interval S is:

$$P(\max\{L_1, \dots, L_r\} < S) = \int_0^\infty G(x) f_S(x) dx, \quad (40)$$

where $G(x) = P(\max\{L_1, \dots, L_r\} < x) = (1 - e^{-\mu x})^r$. The second derivative of $G(x)$ is thus:

$$G''(x) = r\mu^2 e^{-\mu x} (1 - e^{-\mu x})^{r-2} (r e^{-\mu x} - 1), \quad (41)$$

for $r \geq 3$. Then, it is easy to see that for $r \geq 3$:

$$\begin{cases} G''(x) > 0 & x < E[L] \ln r \\ G''(x) \leq 0 & \text{otherwise} \end{cases}, \quad (42)$$

which indicates that $G(x)$ is a convex function for $x < E[L] \ln r$ and concave for $x > E[L] \ln r$.

For $E[S] \rightarrow 0$, notice that $S \leq E[L] \ln r$ holds with probability approaching 1. This immediately transforms (40) into:

$$P(\max\{L_1, \dots, L_r\} < S) = \int_0^{E[L] \ln r} G(x) f_S(x) dx, \quad (43)$$

showing that the convex part of $G(x)$ determines the above metric. Then, for $E[S] \rightarrow 0$ we obtain from Jensen's inequality [12] that:

$$P(\max\{L_1, \dots, L_r\} < S) \geq P(\max\{L_1, \dots, L_r\} < E[S]),$$

since $G(x)$ is strictly convex for $x < E[L] \ln r$. This directly leads to:

$$\lim_{E[S] \rightarrow 0} \frac{\phi_c}{\phi_o} = \frac{P(\max\{L_1, \dots, L_r\} < E[S])}{P(\max\{L_1, \dots, L_r\} < S)} \leq 1, \quad (44)$$

for any random S , which completes the proof. ■

Theorem 6 shows that using constant S is not only a simple but *optimal* method to stabilize successors in Chord.

VI. HEAVY-TAILED LIFETIMES

Without the memoryless property on lifetime L , derivation of probability f that all r successors fail within interval S is simply intractable. However, for systems with heavy-tailed lifetimes [5], [28] where old users are more likely to remain alive for a longer time in the system, a mixture of old and new users within a given successor list leads to a smaller f compared to that for exponential lifetimes. Thus, the probability of node isolation due to failure of the entire successor list in Chord is *smaller* when the distribution of user lifetimes is heavy-tailed compared to the exponential case studied earlier in this paper, which we next confirm in simulations.

We examine four different distributions of interval S , including exponential with rate $1/E[S]$, Pareto with CDF $F(x) = 1 - (1 + x/\beta)^{-\alpha}$ where $\alpha = 3$ and $\beta = (\alpha - 1)E[S]$, uniform in $[0, 2E[S]]$, and constant equal to $E[S]$. Simulation results of isolation probability ϕ for exponential and Pareto lifetimes under the four stabilization strategies are plotted in Fig. 5. Notice in the figure that S with the highest variance

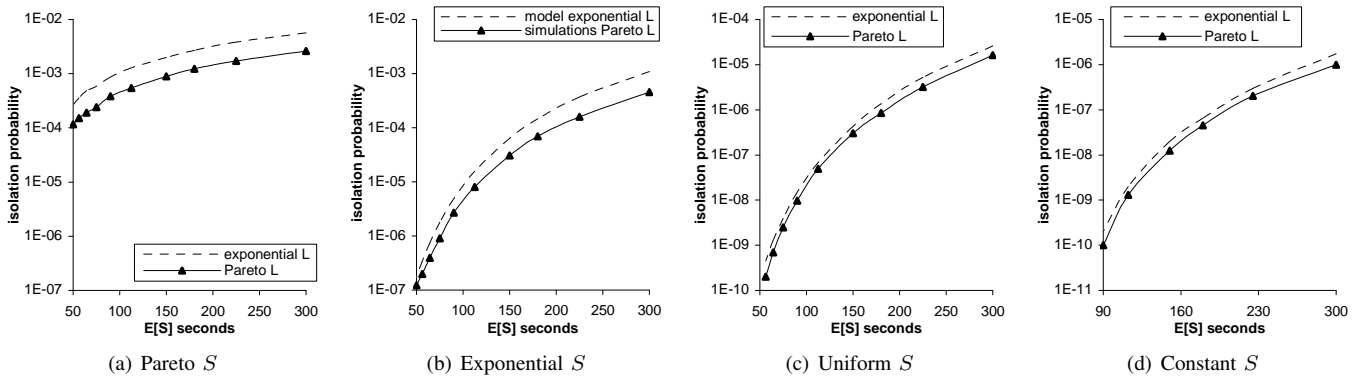


Fig. 5. Comparison of simulation results on node isolation probability ϕ under different stabilization strategies for exponential and Pareto lifetimes with $\alpha = 3$ and $E[L] = 0.5$ hours, mean system size 2,500, and $r = 8$ in Chord.

(i.e., Pareto S) performs the worst, followed by exponential and uniform cases, while constant S is the best. Further observe that ϕ for Pareto lifetimes is smaller than that for exponential lifetimes under all four stabilization strategies and that the difference becomes smaller as $E[S]$ decreases. In fact, the model is a very close match to the Pareto case in Fig. 5(c)-(d). These observations confirm that our exponential model of ϕ provides an upper bound for systems with heavy-tailed lifetimes over a wide range of stabilization delays S .

VII. CONCLUSION

This paper tackled the problem of deriving formulas for the resilience of Chord's successor list under both static and dynamic node failure. We found that under static node failure, Chord exhibited the same resilience through the successor list as that many other DHTs and unstructured P2P networks [16] through their randomized neighbor tables. We also demonstrated that when Chord experienced continuous node joins/departures, stabilization with constant intervals was optimal and kept Chord connected with the highest probability.

REFERENCES

- [1] R. Albert and A. Barabási, "Topology of Evolving Networks: Local Events and Universality," *Phys. Rev. Lett.*, vol. 85, no. 24, pp. 5234–5237, Dec. 2000.
- [2] R. Arratia, L. Goldstein, and L. Gordon, "Two Moments Suffice for Poisson Approximations: The Chen-Stein Method," *The Annals of Probability*, vol. 17, no. 1, pp. 9–25, Jan. 1989.
- [3] J. Aspnes, Z. Diamadi, and G. Shah, "Fault-Tolerant Routing in Peer-to-Peer Systems," in *Proc. ACM PODC*, Jul. 2002, pp. 223–232.
- [4] N. Balakrishnan and M. V. Koutras, *Runs and Scans with Applications*. John Wiley & Sons, 2002.
- [5] F. E. Bustamante and Y. Qiao, "Friendships that Last: Peer Lifespan and its Role in P2P Protocols," in *Proc. Intl. Workshop on Web Content Caching and Distribution*, Sep. 2003.
- [6] M. Castro, M. Costa, and A. Rowstron, "Performance and Dependability of Structured Peer-to-Peer Overlays," in *Proc. DSN*, Jun. 2004.
- [7] B.-G. Chun, B. Zhao, and J. Kubiatowicz, "Impact of Neighbor Selection on Performance and Resilience of Structured P2P Networks," in *Proc. IPTPS*, Feb. 2005, pp. 264–274.
- [8] H. Frank, "Maximally Reliable Node Weighted Graphs," in *Proc. 3rd Ann. Conf. Information Sciences and Systems*, Mar. 1969, pp. 1–6.
- [9] 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.
- [10] M. F. Kaashoek and D. Karger, "Koorde: A Simple Degree-Optimal Distributed Hash Table," in *Proc. IPTPS*, Feb. 2003, pp. 98–107.

- [11] A. K. Kelmans, "Connectivity of Probabilistic Networks," *Auto. Remote Contr.*, vol. 29, pp. 444–460, 1967.
- [12] S. G. Krantz, *Handbook of Complex Variables*. Birkhäuser, 1999.
- [13] S. Krishnamurthy, S. El-Ansary, E. Aurell, and S. Haridi, "A Statistical Theory of Chord under Churn," in *Proc. IPTPS*, Feb. 2005, pp. 93–103.
- [14] S. S. Lam and H. Liu, "Failure Recovery for Structured P2P Networks: Protocol Design and Performance Evaluation," in *Proc. ACM SIGMETRICS*, Jun. 2004, pp. 199–210.
- [15] D. Leonard, V. Rai, and D. Loguinov, "On Lifetime-Based Node Failure and Stochastic Resilience of Decentralized Peer-to-Peer Networks," in *Proc. ACM SIGMETRICS*, Jun. 2005, pp. 26–37.
- [16] D. Leonard, Z. Yao, X. Wang, and D. Loguinov, "On Static and Dynamic Partitioning Behavior of Large-Scale Networks," in *Proc. IEEE ICNP*, Nov. 2005, pp. 345–357.
- [17] J. Li, J. Stribling, T. M. Gil, R. Morris, and M. F. Kaashoek, "Comparing the Performance of Distributed Hash Tables under Churn," in *Proc. IPTPS*, Feb. 2004, pp. 87–99.
- [18] J. Li, J. Stribling, R. Morris, M. F. Kaashoek, and T. M. Gil, "A Performance vs. Cost Framework for Evaluating DHT Design Tradeoffs under Churn," in *Proc. IEEE INFOCOM*, Mar. 2005, pp. 225–236.
- [19] 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.
- [20] 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.
- [21] G. Pandurangan, P. Raghavan, and E. Upfal, "Building Low-Diameter Peer-to-Peer Networks," *IEEE J. Sel. Areas Commun.*, vol. 21, no. 6, pp. 995–1002, Aug. 2003.
- [22] L. Plissonneau, J.-L. Costeux, and P. Brown, "Analysis of Peer-to-Peer Traffic on ADSL," in *Proc. PAM*, Mar. 2005, pp. 69–82.
- [23] S. Rhea, D. Geels, T. Roscoe, and J. Kubiatowicz, "Handling Churn in a DHT," in *Proc. USENIX Ann. Tech. Conf.*, Jun. 2004, pp. 127–140.
- [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] 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.
- [26] D. Stutzbach and R. Rejaie, "Understanding Churn in Peer-to-Peer Networks," in *Proc. ACM IMC*, Oct. 2006, pp. 189–202.
- [27] K. Sutner, A. Satyanarayana, and C. Suffel, "The Complexity of the Residual Node Connectedness Reliability Problem," *SIAM J. Comput.*, vol. 20, pp. 149–155, 1991.
- [28] X. Wang, Z. Yao, and D. Loguinov, "Residual-Based Measurement of Peer and Link Lifetimes in Gnutella Networks," in *Proc. IEEE INFOCOM*, May 2007, pp. 391–399.
- [29] 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.