

In-Degree Dynamics of Large-Scale P2P Systems

Zhongmei Yao
Department of Computer
Science
University of Dayton
Dayton, OH 45469
zyao@udayton.edu

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

Dmitri Loguinov^{*}
Department of Computer
Science and Engineering
Texas A&M University
College Station, TX 77843
dmitri@cs.tamu.edu

ABSTRACT

This paper builds a complete modeling framework for understanding user churn and in-degree dynamics in unstructured P2P systems in which each user can be viewed as a stationary alternating renewal process. While the classical Poisson result on the superposition of n stationary renewal processes for $n \rightarrow \infty$ requires that each point process become sparser as n increases, it is often difficult to rigorously show this condition in practice. In this paper, we first prove that despite user heterogeneity and non-Poisson arrival dynamics, a superposition of edge-arrival processes to a live user under uniform selection converges to a Poisson process when system size becomes sufficiently large. Using this finding, we then obtain closed-form results on the transient behavior of in-degree, paving novel ways for a variety of additional analysis of decentralized P2P systems.

Categories and Subject Descriptors

C.4 [Performance of Systems]: Modeling techniques

General Terms

Algorithms, Performance, Theory

Keywords

Peer-to-Peer, User Churn, Superposition, In-Degree

1. INTRODUCTION

Behavior of P2P networks under churn (i.e., in the face of node join/departure) is one of the most fundamental issues of P2P research [1], [3], [5], [6], [7], [8], [9], [10], [12], [15],[16]. Prior work typically uses a homogeneous Poisson process to model user arrivals, assumes that user lifetimes are exponentially distributed, and addresses only certain limited aspects of P2P performance under churn [6], [10], [16]. Among the recent modeling studies, [20] models each user as an alternating renewal process that is ON when the user is logged in and OFF otherwise. This model allows the online/offline duration distributions to be unique to each user, thus capturing their inherent heterogeneity. While this generic churn model can be adopted in many cases, [20] does not expose any restrictions on the selection of ON/OFF distributions or rigorously establish the Palm-Khintchine result.

To address these problems, we build a novel modeling framework for understanding user churn and in-degree dynamics in P2P systems. Each user i is modeled by a stationary alternating renewal process $Z_i(t)$ representing the user's

online/offline states. Unlike [20], the ON/OFF durations of $Z_i(t)$ are based on the *type* of user i , which is chosen randomly from some finite set of pre-defined peer types (e.g., desktop or mobile). As in many P2P networks and [7], we assume that each joining node selects k initial out-neighbors and then continuously replaces them as they fail with existing users in the system. We build a complete closed-form model characterizing the evolution of edge-arrival and in-degree under the assumption of uniform neighbor selection.

While the Palm-Khintchine Theorem [4] states that the superposition of n stationary renewal processes converges to a Poisson process as $n \rightarrow \infty$, it requires that each point process associated with the renewal process becomes sparser as $n \rightarrow \infty$ and the various processes be independent. However, it is often difficult to prove that these condition hold in real systems, which frequently impedes application of well-known Poisson theory in practice. In this paper, we first prove that *despite user heterogeneity, non-Poisson arrival dynamics, and two types of edge-creation processes (i.e., initial and replacement), the aggregate edge-arrival process to each live user under uniform selection converges to a Poisson process when system size becomes sufficiently large.* This novel result, an important consequence of our modified churn model, offers relatively simple analysis of the edge-arrival process, backed by a solid theoretical justification. We then derive the expected transient in-degree as a function of the lifetime distribution $F(x)$, *including cases with non-exponential peer lifetimes*, and show that users who stay online longer quickly accumulate non-trivial in-degree and become much more resilient to isolation over time.

2. GENERAL EDGE-CREATION MODEL

We start by formalizing the user join/departure model and the out-link model of unstructured P2P networks. Due to limited space, all proofs omitted from this paper can be found in the technical report [19].

2.1 Overview

As observed in real traces [2], [17], each peer can be viewed as alternating between online/offline states. Define $Z_i := \{Z_i(t)\}$ to be an ON/OFF right-continuous process on time interval $[0, \infty)$, indicating online/offline states of user i :

$$Z_i(t) := \begin{cases} 1 & \text{user } i \text{ is ON at time } t \\ 0 & \text{otherwise (OFF)} \end{cases}, t \geq 0, 1 \leq i \leq n, \quad (1)$$

where n is the number of participating peers. Following the illustration at the bottom of Figure 1, define random vari-

^{*}Supported by NSF grant CNS-0720571.

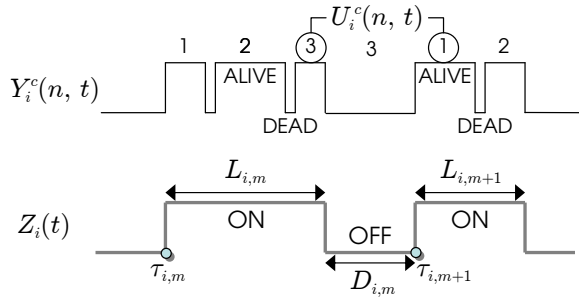


Figure 1: Process $\{Z_i(t)\}$ depicts ON/OFF behavior of user i , for $i = 1, \dots, n$. Process $\{Y_i^c(n, t)\}$ indicates DEAD/ALIVE behavior of the c -th out-link of user i , for $c = 1, \dots, k$. Process $\{U_i^c(n, t)\}$ counts the number of DEAD \rightarrow ALIVE transitions within the current ON/OFF cycle of i .

ables $L_{i,m} > 0, D_{i,m} > 0$ to be durations of user i 's ON (life) and OFF (death) periods, respectively, and $\{\tau_{i,m}\}_{m=1}^{\infty}$ to be arrival times of user i , where $\tau_{i,m+1} = \tau_{i,m} + L_{i,m} + D_{i,m}$, for $m \geq 1$. Further define $\{M_i(t)\}$ to be the point process that counts the number of arrivals of user i in interval $[0, t]$:

$$M_i(t) := \sum_{m=1}^{\infty} 1_{\tau_{i,m} \in [0, t]}. \quad (2)$$

To remain connected and participate in the system, each user i creates k out-links to other peers upon joining the network and then continuously repairs them as they fail. Denote by $Y_i^c := \{Y_i^c(n, t)\}$ an alternating process representing the states of i 's out-link c :

$$Y_i^c(n, t) := \begin{cases} 1 & \text{out-link } c \text{ of user } i \text{ is ALIVE at } t \\ 0 & \text{otherwise (DEAD)} \end{cases}, \quad (3)$$

for $c = 1, \dots, k$ and $t \geq 0$, where ALIVE means that the neighbor adjacent to this link is currently alive and DEAD means that the neighbor has departed from the system and a replacement is being sought. If node i is offline at t , all of its links are considered DEAD.

As depicted at the top of Figure 1, user i 's c -th link becomes ALIVE at arrival times $\{\tau_{i,m}\}_{m \geq 1}$ and then alternates between DEAD and ALIVE states during i 's ON periods. Note that ALIVE durations of Y_i^c are *residual lifetimes* of selected neighbors and DEAD durations are *search delays* for finding replacement neighbors, with the exception of the very last ALIVE cycle in each ON period, which is terminated by i 's departure rather than neighbor failure. Whenever Y_i^c transitions from DEAD to ALIVE, user i creates an edge (i.e., performs one selection). Define *initial edges* to be those added when users arrive in the system (i.e., whenever Z_i transitions from OFF to ON) and *replacement edges* to be those added in response to neighbor failures. Note that processes $\{Y_i^c\}$ are dependent in general and rather complex since multiple users may concurrently connect to the same neighbor and each out-link may point to a peer v again after v re-appears in the system.

Figure 1 also shows right-continuous process $\{U_i^c(n, t)\}$, which is the number of transitions DEAD \rightarrow ALIVE of Y_i^c within the current ON/OFF cycle of i up to time t . We set $U_i^c(n, \tau_{i,m}) := 1$ for all arrival times $\tau_{i,m}$, use notation t^- to

represent the instant just prior to t , and denote by

$$U_i^c(n, \tau_{i,m+1}^-) = \sup_{\tau_{i,m} \leq t < \tau_{i,m+1}} U_i^c(n, t) \quad (4)$$

the number of selections for link c in the m -th ON cycle of user i .

Finally, the number of selections that i performs for link c across all ON durations of i in $[0, t]$ is given by

$$W_i^c(n, t) := \sum_{m=1}^{M_i(t)} U_i^c(n, \tau_{i,m}^-) - U_i^c(n, 0) + U_i^c(n, t), \quad (5)$$

where $M_i(t)$ is the number of arrivals of i in $[0, t]$. Then, the edge-creation process of user i is the superposition of k processes W_i^c :

$$W_i(n, t) := \sum_{c=1}^k W_i^c(n, t). \quad (6)$$

Observe that $\sum_{i=1}^n W_i(n, t)$ is the number of out-degree edges generated by n users in $[0, t]$, which is the same as the number of in-degree edges received by living users in $[0, t]$.

2.2 Assumptions

We next make several assumptions about the system, without which the main result of this paper may not hold.

ASSUMPTION 1. *The number of out-links k a user creates upon joining the system is a constant for all n .*

This assumption often holds in unstructured P2P networks where individual users are unaware of system size (e.g., Gnutella) and some structured P2P networks with constant node degree (e.g., de Bruijn [11]).

The next assumption classifies users into a number of types and explains how users generate their online/offline durations according to their types.

ASSUMPTION 2. 1) *There exists some set \mathcal{F} of distinct pairs of non-lattice CDFs defining non-negative random variables:*

$$\mathcal{F} := \left\{ \left(F^{(1)}(x), G^{(1)}(x) \right), \dots, \left(F^{(\mathcal{T})}(x), G^{(\mathcal{T})}(x) \right) \right\},$$

where $\mathcal{T} \geq 1$ is a fixed number of user types and CDFs $F^{(j)}(x) > 0$ and $G^{(j)}(x) > 0$ for all $x > 0$, all $j = 1, \dots, \mathcal{T}$. Further, each mean $l^{(j)} := \int_0^{\infty} (1 - F^{(j)}(x)) dx$ and $d^{(j)} := \int_0^{\infty} (1 - G^{(j)}(x)) dx$ satisfies $0 < l^{(j)}, d^{(j)} < \infty$ for all types $j = 1, \dots, \mathcal{T}$;

2) *Each Z_i is associated with a pair of CDFs $(F_i(x), G_i(x))$ that is independently drawn from set \mathcal{F} , where type j is selected with probability $p_j \geq 0$ and $\sum_{j=1}^{\mathcal{T}} p_j = 1$;*

3) *Defining \mathcal{S} to be set of selections made by each user and conditioning on \mathcal{S} , set $\{Z_i(t)\}_{i=1}^n$ consists of mutually independent, stationary alternating renewal processes, where ON durations $\{L_{i,m}\}_{m=1}^{\infty}$ are independently identically distributed (i.i.d.) with CDF $F_i(x)$ and OFF durations $\{D_{i,m}\}_{m=1}^{\infty}$ are i.i.d. with distribution $G_i(x)$.*

Part 1) in Assumption 2 uses \mathcal{T} as the ‘‘diversity’’ factor of user behavior (e.g., $\mathcal{T} = 1$ reduces the system to a network of homogeneous users) and mandates that all average

online/offline durations are both positive and finite. Part 2) allows for bias in the selection process and lets certain user types be more popular than others. Part 3) ensures that once users have chosen their types (i.e., ON/OFF duration CDFs), $\{Z_i(t)\}_{i=1}^n$ evolve as independent stationary renewal processes.

For simplicity of notation, conditioning on user type, define $l_i := E[L_{i,m}|i\text{'s type}]$ and $d_i := E[D_{i,m}|i\text{'s type}]$ to be the mean ON/OFF duration of each user i , respectively. Denote by

$$\lambda_i := (l_i + d_i)^{-1} \quad (7)$$

the arrival rate of the user into the system. To ensure stationarity, the first arrival time $\tau_{i,1}$ is defined as

$$\tau_{i,1} := \begin{cases} L_i^e + D_i & \text{with probability (w.p.) } a_i \\ D_i^e & \text{w.p. } 1 - a_i \end{cases}, \quad (8)$$

where L_i^e has the equilibrium distribution of $F_i(x)$ (i.e., $P(L_i^e < x) = (l_i)^{-1} \int_0^x (1 - F_i(u)) du, x \geq 0$), D_i^e has the equilibrium distribution of $G_i(x)$, and a_i is user i 's availability:

$$a_i := l_i(l_i + d_i)^{-1}. \quad (9)$$

Due to stationarity, the expected number of arrivals of i in interval $[0, t]$ is $E[M_i(t)] = \lambda_i t$ for any $t \geq 0$, where λ_i is given in (7).

In the next subsection, we focus our attention on properties of aggregate user lifetimes, the system population at time t , and residual lifetimes of selected neighbors, upon which we compute the rate of edge-creation from each user.

2.3 Properties

For each instance of user i being present in the system during interval $[0, t]$, place its ON duration $L_{i,m}$ into set $S_i(t)$ and define $S(t) = \cup_{i=1}^n S_i(t)$. Then let $F(n, t, x)$ be the CDF of values collected in set $S(t)$ (i.e., the probability that the obtained lifetimes are less than or equal to x). Given n participating users, define $F(n, x) := \lim_{t \rightarrow \infty} F(n, t, x)$ to be the *aggregate lifetime distribution* of the system and $l(n)$ to be its mean in the equilibrium. By [20, Theorem 1], with Assumption 2 and any finite $n \geq 1$, the aggregate lifetime CDF $F(n, x)$ and its mean $l(n)$ are respectively given by

$$F(n, x) = \sum_{i=1}^n b_i F_i(x), \quad l(n) = \sum_{i=1}^n b_i l_i, \quad (10)$$

where $b_i := \lambda_i / \sum_{j=1}^n \lambda_j$ and λ_i is defined in (7). The asymptotic results on $F(n, x)$ and $l(n)$ are given below.

LEMMA 1. *Under Assumption 2, the following sequences converge almost surely (a.s.) as $n \rightarrow \infty$:*

$$F(n, x) \xrightarrow{\text{a.s.}} F(x) := \frac{\sum_{j=1}^{\mathcal{T}} p_j \lambda^{(j)} F^{(j)}(x)}{\lambda}, \quad (11)$$

$$l(n) \xrightarrow{\text{a.s.}} l := \frac{a}{\lambda}, \quad (12)$$

where $\lambda^{(j)} := 1/(l^{(j)} + d^{(j)})$ is the arrival rate of type j , $\lambda := \sum_{j=1}^{\mathcal{T}} p_j \lambda^{(j)}$ is the average arrival rate across all user types, $a := \sum_{j=1}^{\mathcal{T}} p_j a^{(j)}$, and $a^{(j)} := l^{(j)}/(l^{(j)} + d^{(j)})$ of type j . Further, $F(x)$ is a proper CDF function and $0 < l < \infty$.




Figure 2: Residual $R_i(t)$ is the duration of user i from time t until i departs, conditioned on $Z_i(t) = 1$.

Now, suppose node v picks a random currently-alive user i as a potential neighbor. Denote by $R_i(t)$ the remainder of the current ON cycle of user i , as illustrated in Figure 2, and by $H_i(x)$ the CDF of $R_i(x)$:

$$H_i(x) := P(R_i(t) \leq x | Z_i(t) = 1), \quad (13)$$

which is invariant in t due to stationarity of Z_i . It is known from [18] that $H_i(x) = (l_i)^{-1} \int_0^x (1 - F_i(u)) du, x \geq 0$.

Next, define $R(n, t)$ to be the residual lifetime of the user *uniformly randomly selected* from among peers that are alive at time t . Let $N(n, t) := \sum_{i=1}^n Z_i(t)$ be the system population at time t and denote by $H(n, x)$ the distribution of $R(n, t)$ conditioned on that at least one user is alive:

$$H(n, x) := P(R(n, t) \leq x | N(n, t) \geq 1). \quad (14)$$

It is not hard to see from (14) that $H(n, x)$ depends on the distribution of $N(n, t)$, the CDF of $R_i(t)$ of each user, and neighbor selection strategies. In the following, we examine the properties of $N(n, t)$, upon which we derive $H(n, x)$ when user v uniformly randomly selects neighbors.

LEMMA 2. *Given Assumption 2 and $N(n, t) \geq 1, \mu_n/N(n, t)$ converges to 1 in r -th mean for all $r \geq 1$:*

$$\lim_{n \rightarrow \infty} E \left[\left| \frac{\mu_n}{N(n, t)} - 1 \right|^r | N(n, t) \geq 1 \right] = 0, \quad (15)$$

where $\mu_n := E[N(n, t)] = na$ is the mean population and a is given immediately following (12).

Applying the above two lemmas leads to the main result of this subsection.

THEOREM 1. *Given Assumption 2 and uniform selection, $H(n, x)$ defined in (14) converges a.s. to the following as $n \rightarrow \infty$:*

$$H(n, x) \xrightarrow{\text{a.s.}} H(x) := \frac{1}{l} \int_0^x (1 - F(u)) du, \quad (16)$$

where $F(x)$ and l are given in (11)-(12).

Theorem 1 shows that the residual CDF $H(n, x)$ can be reduced to a simple function of the aggregate lifetime CDF $F(x)$. Recalling that ALIVE durations of Y_i^c are residual lifetimes of selected users and applying (16), we next deal with the edge-arrival process to each user i .

3. EDGE ARRIVAL PROCESS

For convenience, denote by $\delta_{i,z}$ the z -th time at which user i makes a selection, across all links of i , and by $I_{i,z}^j$ the indicator that user i selects peer j for its z -th connection. Note that the z -th selection time $\delta_{i,z}$ of user i is determined by the ON/OFF process Z_i and residual lifetimes of i 's all

previous selections $z' < z$. This indicates that the history observed by i up to time $\delta_{i,z}$ clearly helps i predict whether the previously selected peers are alive at $\delta_{i,z}$. The model for the number of users available for selection at time $\delta_{i,z}$ is rather intricate due to the dependence on i 's past selections.

3.1 Preliminaries

In the following, we first deal with uniform integrability of set $\{(W_i(n,t))^r\}_{n \geq 1}$ and then proving that residuals of selected neighbors in $[0, t]$ are asymptotically independent as $n \rightarrow \infty$. These two important results make the edge-creation process tractable for systems with large size and pave the way for computing the CDF of residual $R(n, \delta_{i,z})$ of neighbor selected at time $\delta_{i,z}$, as presented in Lemma 3. Note that uniform integrability of $\{(W_i(n,t))^r\}_{n \geq 1}$ shows that $E[(W_i(n,t))^r]$ converge as $n \rightarrow \infty$, for all $r > 0$. We then apply the distribution of residuals $R(n, \delta_{i,z})$ and the regenerative property of process $\{U_i^c(n,t)\}$ (see (4)) to obtain the limit for $E[W_i(n,t)]$.

LEMMA 3. *With Assumptions 1-2 and uniform selection, we have for each user i and all $t \geq 0$:*

- 1) Collections $\{(W_i(n,t))^r\}_{n \geq 1}$ are uniformly integrable in n , for any $r > 0$;
- 2) Residuals $\{R(n, \delta_{i,z})\}_{z \geq 1}$ of selected neighbors at random times $\delta_{i,z}$ are asymptotically independent and converge in distribution as $n \rightarrow \infty$ to i.i.d. r.v.'s with CDF $H(x)$ in (16);
- 3) Define $\{U(s)\}_{s \geq 0}$ to be a pure renewal process whose cycle lengths follow the CDF $H(x)$ in (16) and whose mean is $E[U(x)] = 1 + \sum_{r=1}^{\infty} H^{*r}(x)$, where $H^{*r}(\cdot)$ is the r -fold convolution of $H(\cdot)$. Given user i 's type, the mean number of edges that i creates in $[0, t]$ is

$$\begin{aligned} \lim_{n \rightarrow \infty} E[W_i(n,t) | i \text{'s type}] &= k\lambda_i t E[U(L_i) | i \text{'s type}] \\ &= k\lambda_i t \int_0^{\infty} E(U(x)) dF_i(x), \end{aligned} \quad (17)$$

where $F_i(x)$ is i 's lifetime CDF and λ_i is its arrival rate. Unconditioned on user types,

$$E[W_i(n,t)] = (k + \theta)\lambda t, \quad (18)$$

where $\theta := k \sum_{r=1}^{\infty} \int_0^{\infty} H^{*r}(x) dF(x)$, λ is given in (12), and $F(x)$ is the aggregate lifetime CDF in (11).

We next make several notes on Lemma 3. Intuitively, as system size n increases, the probability that user i selects any other peer more than once in the interval $[0, t]$ becomes smaller. In other words, for sufficiently large n , it is more likely that the set of neighbors that i connects to in $[0, t]$ contains *distinct* peers and thus the CDF of neighbor residuals selected at $\delta_{i,z}$ approaches that of residual $R(n,t)$ of a randomly selected user shown in (16), which is the rationale behind part 2 of Lemma 3.

The result in (18) demonstrates that as $n \rightarrow \infty$, user i brings k initial edges and an average of $\theta < \infty$ replacement edges per ON/OFF cycle into the system and that the edge creation rate from each user approaches the constant $(k + \theta)\lambda$. As we will see in the next subsection, $(k + \theta)\lambda$ is equal to the edge-arrival rate from the system to a given existing user.

3.2 Poisson Result

Given a set of n participating users, our approach is to set aside a *given* user v (so the values of $Z_v(\cdot)$ are given) and examine edge arrival to this user, from $n - 1$ other peers under uniform selection.

Define $\xi_{n,i}(t) := \sum_{z=1}^{W_i(n,t)} I_{i,z}^v$ to be the number of edges delivered by user i to node v in $[0, t]$, for $i \neq v$, where $I_{i,z}^v$ is the indicator that i selects v at time $\delta_{i,z}$. Then, the edge arrival process from the system to user v is the superposition of processes $\xi_{n,i}$:

$$\xi_n(t) := \sum_{i=1, i \neq v}^n \xi_{n,i}(t) = \sum_{i=1, i \neq v}^n \sum_{z=1}^{W_i(n,t)} I_{i,z}^v. \quad (19)$$

The properties of process ξ_n are presented in Theorem 2, which follows from Lemma 3 and the result on the convergence of random measures in [14, Proposition 3.22].

THEOREM 2. *Under Assumptions 1-2 and uniform selection, conditioned on Z_v , the point process ξ_n defined in (19) converges in distribution as $n \rightarrow \infty$ to a non-homogeneous Poisson process ξ with local rate $\gamma Z_v(t)$, where $Z_v(t)$ is a deterministic function of t ,*

$$\gamma := (k + \theta)/l, \quad (20)$$

θ is in (18), and l is the mean lifetime given in (12).

Theorem 2 shows that when user v is alive (i.e., $Z_v(\cdot) = 1$), the instantaneous rate of edge arrival to v is the constant γ ; otherwise it is 0. That is, the edge arrival process to v is a Poisson process whose rate varies according to process Z_v .

The above important result states that despite multiple user-types and non-Poisson user-arrival dynamics, the edge-arrival process to each user v is Poisson. This allows us to obtain relatively simple models for in-degree that we study later in this paper, and other interesting metrics (e.g., traffic load on each user).

3.3 Simulations

We next show simulations explaining this result and its accuracy in systems with *finite* age and size. We generate a network of n users whose arrival/departure follows the introduced churn model. The system evolves for at least 50 virtual hours before being examined. We start by generating $\mathcal{T} = 1,000$ pairs of means l_i and d_i , which are drawn randomly from two Pareto distributions with $\alpha = 3$ as described next. For mean ON durations, we use $\beta = 1$ and obtain $E[l_i] = 1/2$ hour; for mean OFF durations, we use $\beta = 2$ and get $E[d_i] = 1$ hour. We study three cases throughout the paper: 1) heavy-tailed system \mathcal{H} with $F^{(j)}(x) \sim \text{Pareto}(3, 2l^{(j)})$ and $G^{(j)}(x) \sim \text{Pareto}(3, 2d^{(j)})$; 2) very heavy-tailed system \mathcal{VH} with $F^{(j)}(x) \sim \text{Pareto}(1.5, l^{(j)}/2)$ and $G^{(j)}(x) \sim \text{Pareto}(1.5, d^{(j)}/2)$; and 3) exponential system \mathcal{E} with $F^{(j)}(x) \sim \exp(1/l^{(j)})$ and $G^{(j)}(x) \sim \text{Pareto}(3, 2d^{(j)})$, where notation $\text{Pareto}(\alpha_i, \beta_i)$ represents

$$F_i(x) = 1 - (1 + x/\beta_i)^{-\alpha_i}, \quad \alpha_i > 1, \beta_i > 0, \quad (21)$$

The actual pairs $(F_i(x), G_i(x))$ are selected uniformly randomly from \mathcal{F} .

Figure 3 shows the distribution of edge inter-arrival delays to a single node obtained in simulations with two types of systems, given that the node is alive. Notice in the sub-figures that for finite n , the distribution of inter-arrival delay

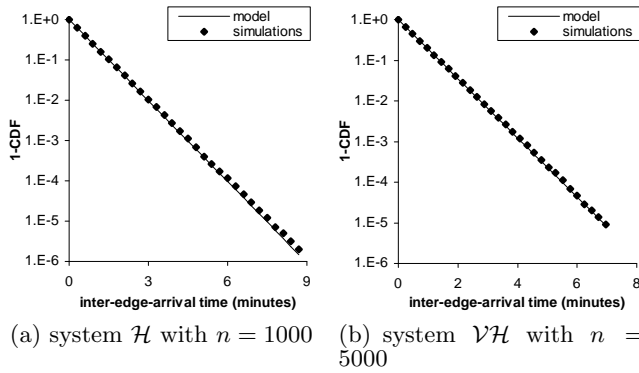


Figure 3: Distribution of edge inter-arrival delays approaches exponential with rate γ in (20) for $n = 1000$, $k = 10$, and $\theta = 10$ using 10^9 iterations.

is approximately exponential with the rate given by (20). Additionally, Figure 4 shows that the distribution of the number of edge arrivals to a node in an interval of size Δt approaches a Poisson distribution with the same rate γ in (20).

Finally, note that the Poisson result in Theorem 2 is not an *assumption* of the paper as in prior work [6], [10], [13], but rather a *consequence* of the churn model introduced earlier.

4. IN-DEGREE

In this section, we aim to understand how the in-degree of each *live* user changes with user age. We start by building a generic in-degree model for any lifetime distribution $F(x)$.

4.1 Expected In-Degree

In a stationary system, define $X_n(t)$ to be the in-degree of a random online user v at age $t \geq 0$. Applying Theorem 1 and Lemma 3, we show transient and limiting properties of $X_n(t)$ under Assumptions 1-2 and uniform selection of neighbors in the next theorem.

THEOREM 3. *Given that a user is alive in the system, its expected in-degree at fixed age $t \geq 0$ converges as $n \rightarrow \infty$ to a monotonically increasing function of age*

$$E[X_n(t)] \rightarrow k \int_0^\infty (E[U(x) - U(x-t)]) dH(x), \quad (22)$$

where $H(x)$ is in (16), $\{U(s)\}_{s \geq 0}$ is a pure renewal process defined in Lemma 3, $E(U(x)) = 1 + \sum_{r=1}^\infty H^{*r}(x)$ for $x \geq 0$, and $E(U(x)) = 0$ for $x < 0$.

To better interpret the mean in-degree of user v at fixed age t , we transform (22) into

$$\lim_{n \rightarrow \infty} E[X_n(t)] = kE[U(R) - U(R-t)], \quad (23)$$

where r.v. $R \sim H(x)$ denotes the age of a *random* user i that is ON at time t , for $i \neq v$, $R-t$ is the duration that i is alive but v is offline, and $U(x)$ represents the number of selections that i makes for link c in an interval with a given length x (recall $U(x) = 0$ for $x < 0$). It is then not hard to see that the left-hand side of (23) is the expected number of connections that peer i builds since both i and v are alive. By Theorem 3, this metric is equal to v 's mean in-degree (i.e., the expected number of existing users that select

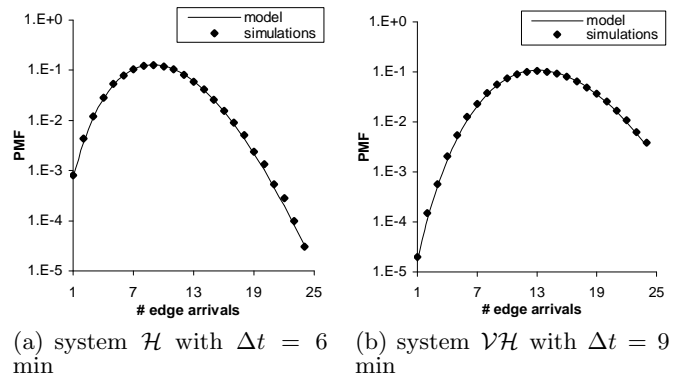


Figure 4: Distribution of the number of edge arrivals to a node in the interval $[t, t + \Delta t]$ in a system with $n = 1000$ users, $k = 10$, and $\theta = 10$. The lines show Poisson fits with γ in (20) at $t = 500$ hours and after 10^5 iterations.

v since v 's arrival), which is consistent with the symmetry of out-going/in-coming edges under uniform selection. Our last note is that as $t \rightarrow \infty$, (23) tends to $kE[U(R)]$, which provides a simple upper-bound at which the in-degree of each user saturates.

We next show that (23) can be expressed in simple closed-form for exponential lifetimes and further apply Theorem 2 to obtain the distribution of in-degree $X_n(t)$ at age t .

THEOREM 4. *For exponential lifetimes L and $n \rightarrow \infty$, the mean in-degree at failure $\theta = k$ and*


$$E[X_n(t)] \rightarrow 2k(1 - e^{-t/E[L]}). \quad (24)$$

Moreover, the distribution of $X_n(t)$ converges to a Poisson distribution with mean $2k(1 - e^{-t/E[L]})$ as $n \rightarrow \infty$.

In (24), the mean in-degree of a node increases monotonically from $X_n(0) = 0$ when it arrives into the system to $E[X_n(\infty)] = 2k$ when its age tends to infinity. For the exponential case we directly use (24), while for the Pareto lifetime case with $\alpha > 2$ we numerically solve (23). Simulation results in Figure 5 demonstrate that the models are very accurate and indeed saturate at predicted values $2k$ and $kE[U(R)]$ as age $t \rightarrow \infty$. Furthermore, if a node survives for more than 1 hour in the system under the parameters of this simulation, it develops an average of 12 – 15 in-degree neighbors (depending on the distribution of L) and is unlikely to be isolated from the graph from that point on. It is also interesting to observe in the figure that the Pareto curve increases slower, but saturates at larger values, which suggests more resilience support for users with very large lifetimes.

5. CONCLUSION

This paper introduced a generic framework for modeling user join/departure, edge arrival, and in-degree. Using this foundation, it then developed closed-form results on the edge-arrival process to each user and their transient in-degree. Future work involves modeling of non-uniform neighbor selection, measurement of real P2P dynamics, and analysis of non-stationary user arrivals.



(a) exponential (24) with $\mu = 2$ (b) Pareto (22) with $\alpha = 3$

Figure 5: Comparison of the model for $E[X_n(t)]$ to simulation results for $n = 2000$, $E[L] = 0.5$ hours, and $k = 8$ after 10^6 iterations.

6. REFERENCES

- [1] M. Castro, M. Costa, and A. Rowstron, "Performance and Dependability of Structured Peer-to-Peer Overlays," in *Proc. DSN*, Jun. 2004.
- [2] P. B. Godfrey, S. Shenker, and I. Stoica, "Minimizing Churn in Distributed Systems," in *Proc. ACM SIGCOMM*, Sep. 2006.
- [3] 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.
- [4] D. Heyman and M. Sobel, *Stochastic Models in Operations Research, Volume 1*. McGraw-Hill, 1982.
- [5] M. F. Kaashoek and D. Karger, "Koorde: A Simple Degree-Optimal Distributed Hash Table," in *Proc. IPTPS*, Feb. 2003, pp. 98–107.
- [6] 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.
- [7] 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.
- [8] 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.
- [9] 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.
- [10] 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.
- [11] D. Loguinov, J. Casas, and X. Wang, "Graph-Theoretic Analysis of Structured Peer-to-Peer Systems: Routing Distances and Fault Resilience," *IEEE/ACM Trans. Networking*, vol. 13, no. 5, pp. 1107–1120, Oct. 2005.
- [12] P. Maymoukov and D. Mazieres, "Kademlia: A Peer-to-Peer Information System Based on the XOR Metric," in *Proc. IPTPS*, Mar. 2002, pp. 53–65.
- [13] 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.
- [14] S. I. Resnick, *Extreme Values, Regular Variation, and Point Processes*. Boston, NY: Springer-Verlag, 1987.
- [15] 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.
- [16] 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.
- [17] 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.
- [18] R. W. Wolff, *Stochastic Modeling and the Theory of Queues*. Englewood Cliffs, NJ: Prentice Hall, 1989.
- [19] Z. Yao, D. B. Cline, and D. Loguinov, "In-Degree Dynamics of Large-Scale P2P Systems," University of Dayton, Tech. Rep., Apr. 2010. [Online]. Available: <http://academic.udayton.edu/zhongmeiyao/papers/indegree-report.pdf>.
- [20] 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.