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

Zhongmei Yao

Joint work with Xiaoming Wang, Derek Leonard, and Dmitri Loguinov

Internet Research Lab
Department of Computer Science
Texas A&M University, College Station, TX 77843

May 10, 2007

Agenda

- Motivation and background
  - Terminology, assumptions, and related work
- Generic node isolation model
- Max-age selection
- Age-proportional random-walk selection
- Wrap-up

Terminology

- Resilience of unstructured P2P networks
  - Ability of a network to remain connected under node failure, which is fundamental to system performance
- User churn
  - Each user stays in the system for L random time units
- Out-degree
  - Joining users select k neighbors
- Neighbor replacement
  - Detection of failed neighbors and replacement with existing peers occur within S time units (can be fixed or random)

Background

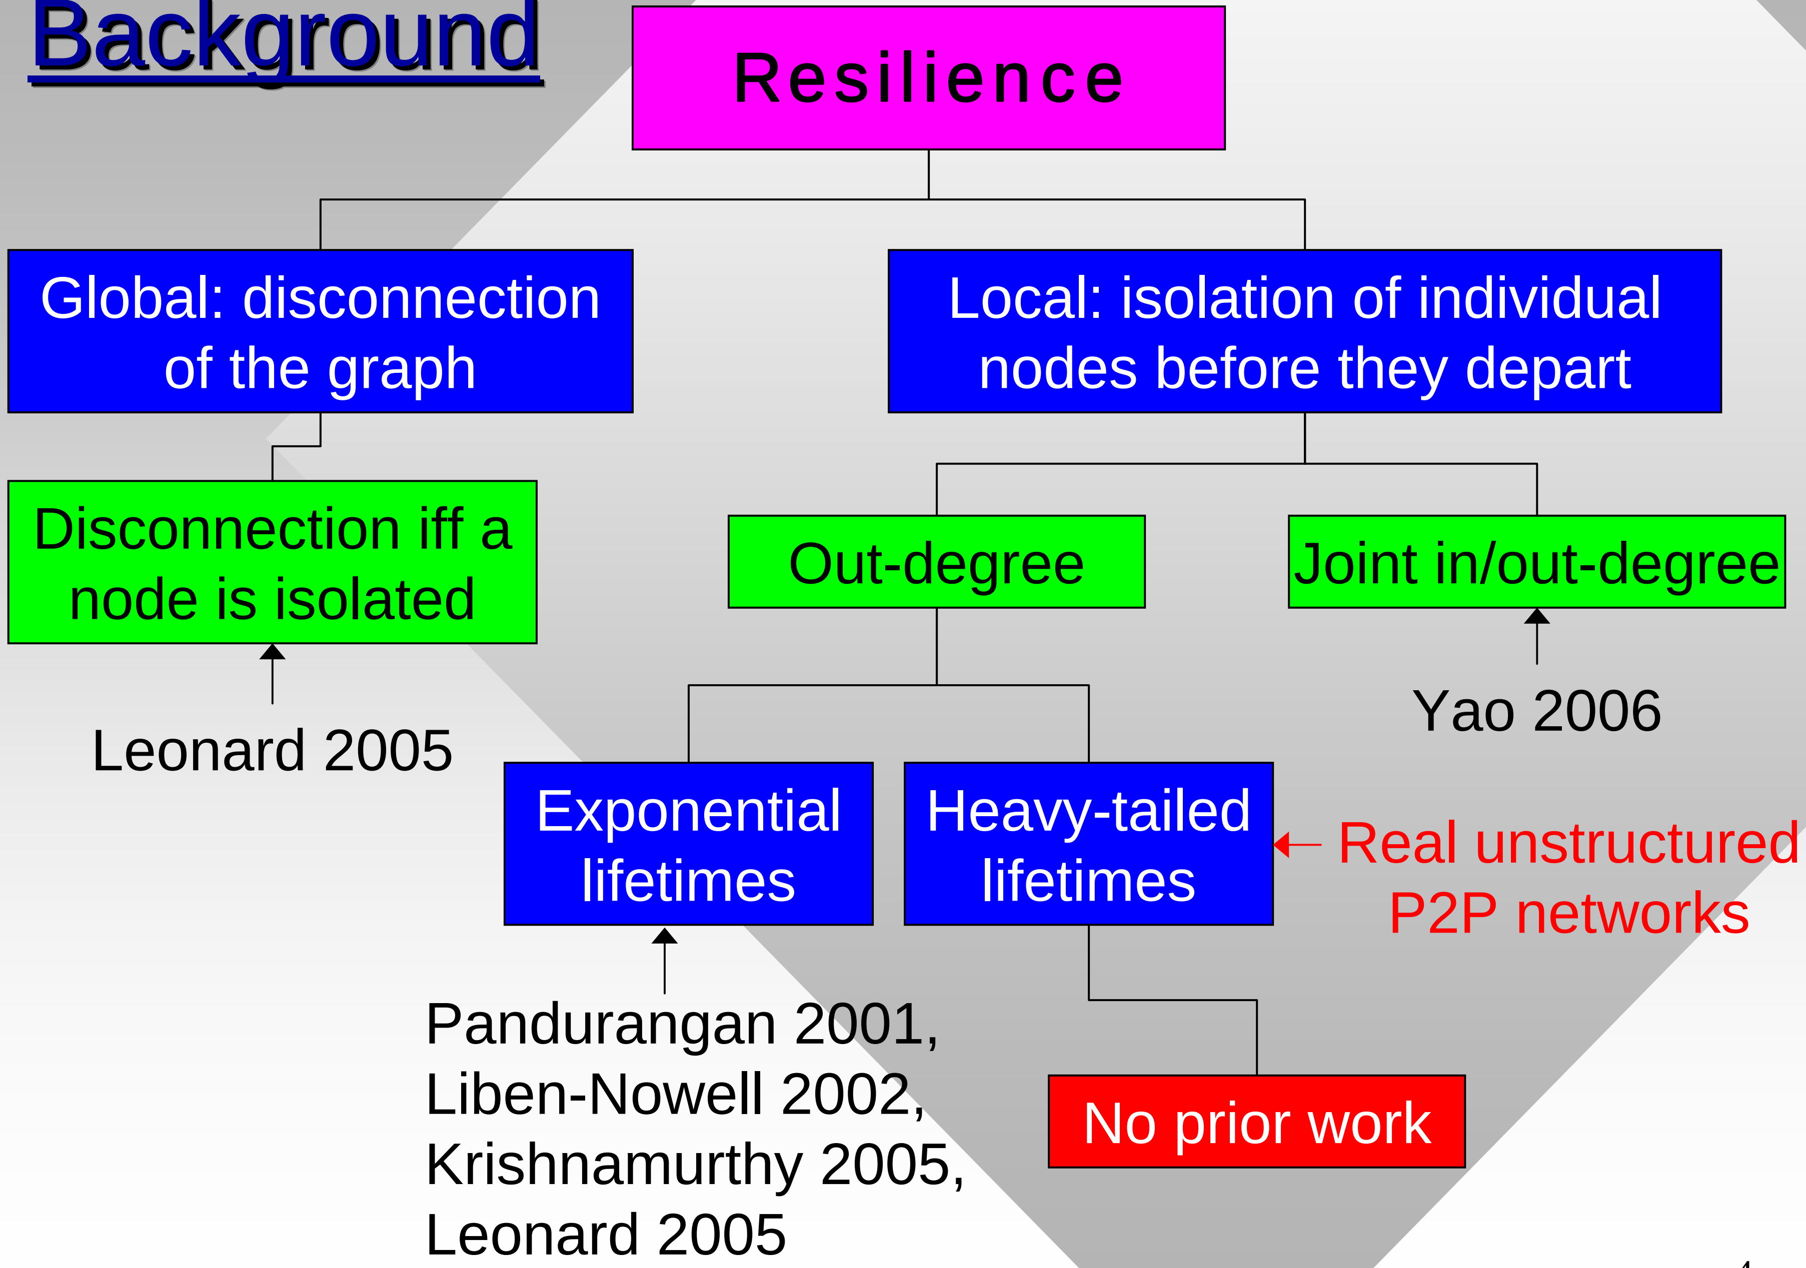

Resilience

Global: disconnection of the graph

Local: isolation of individual nodes before they depart

Disconnection iff a node is isolated

Leonard 2005

Out-degree

Joint in/out-degree

Yao 2006

Exponential lifetimes

Heavy-tailed lifetimes

Real unstructured P2P networks

Pandurangan 2001,
Liben-Nowell 2002,
Krishnamurthy 2005,
Leonard 2005

No prior work

Agenda

- Motivation and background
  - Terminology, assumptions, and related work
- Generic node isolation model
- Max-age selection
- Age-proportional random-walk selection
- Wrap-up

Model Basics

- Neighbor residual lifetimes R
  - The time duration from the instance a peer is selected by user v as a neighbor until the peer departs

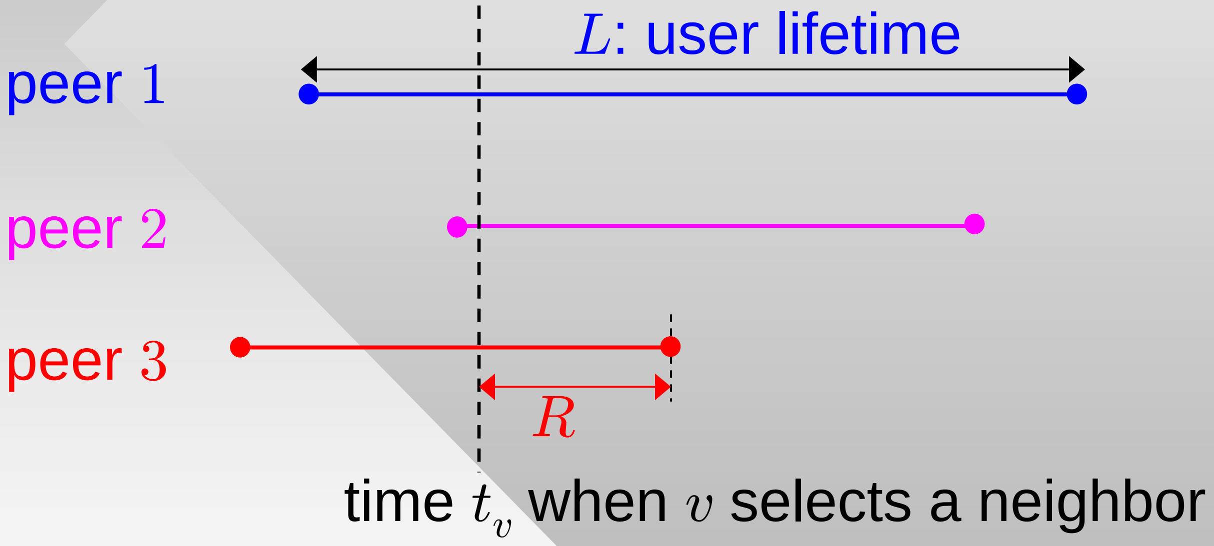

L: user lifetime

peer 1

peer 2

peer 3

R

time t_v when v selects a neighbor

- This metric depends on neighbor selection strategies
  - Some strategies may find users with large residual lifetimes with high probability while others may not

Model Basics 2

- Neighbor failure/replacement is an ON/OFF process

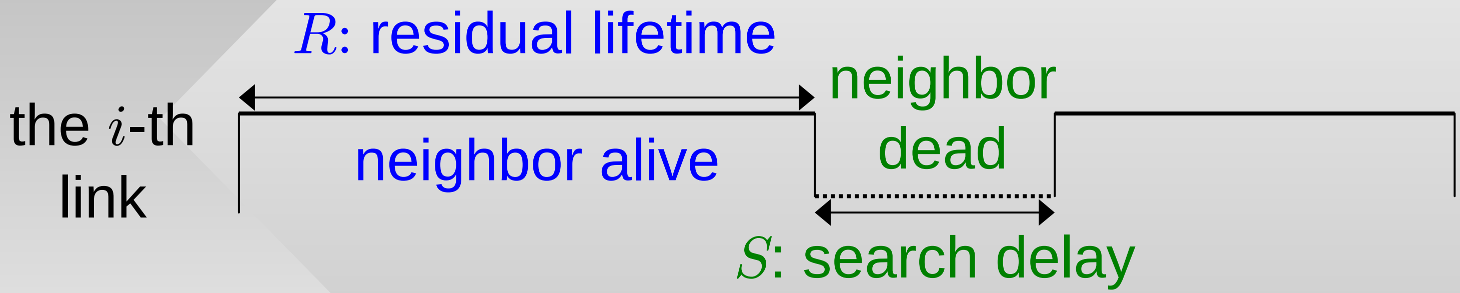

- Node out-degree evolution

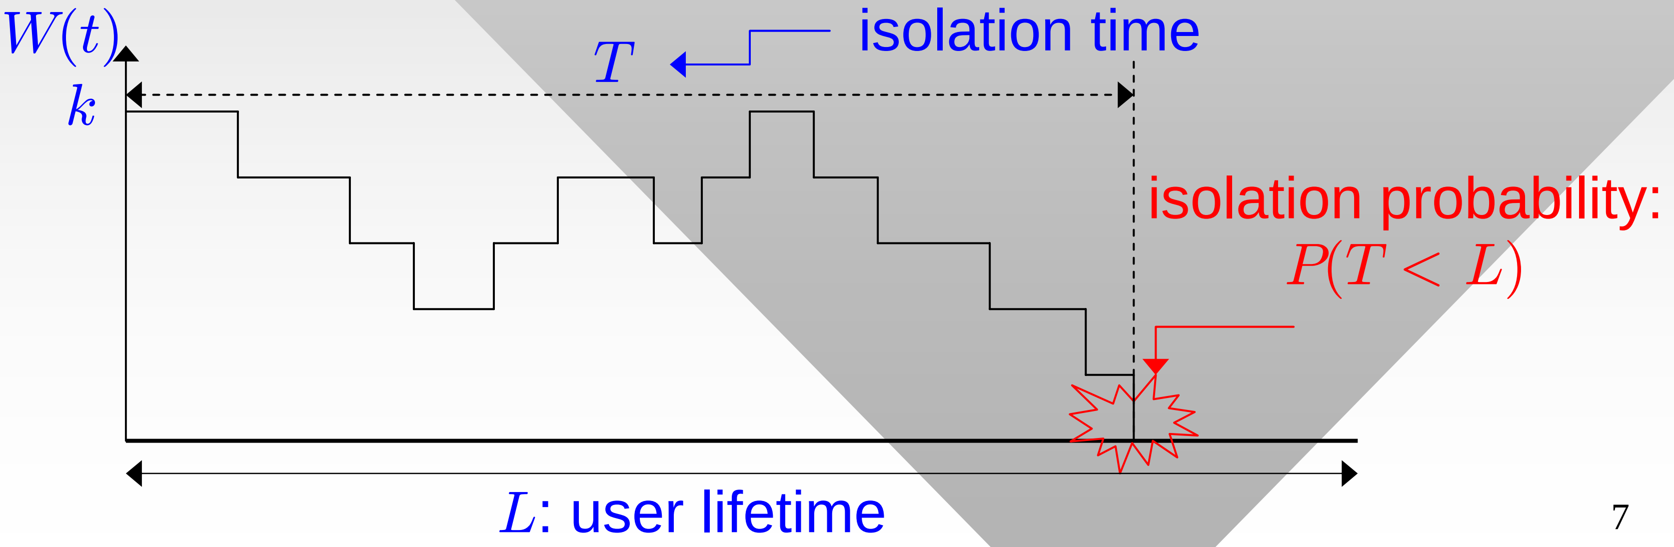

Out-Degree Process

- Determining the first-hitting time of $W(t)$ to zero (i.e., isolation time T) is difficult unless $W(t)$ is Markovian
  - — Idea: replace the distribution of ON/OFF durations with a hyper-exponential approximation (see paper for details)
  - — It is well-known (Feldmann 1998) that any completely monotone density function (e.g., Pareto, Weibull) can be approximated by a hyper-exponential PDF with arbitrary accuracy
- Theorem 1: For hyper-exponential neighbor residual lifetimes R and hyper-exponential search delays S, the out-degree process $\{W(t)\}$ is a continuous-time Markov process

Node Isolation Probability

- Theorem 2: Given that $\{W(t)\}$ is a Markov process, the PDF $f_T(t)$ of the isolation time T can be obtained using the transition rate matrix of process $\{W(t)\}$ shown in the paper
  - Then, it is straightforward to obtain:

$$\underset{\uparrow}{\boxed{\phi}} = P(T < L) = \int_0^\infty \underset{\uparrow}{\boxed{P(L > t)}} f_T(t)dt$$

node isolation probability

the CCDF of user lifetimes

Accuracy of Node Isolation Model

- Simulation results on isolation probability ϕ for $E[L] = 0.5$ hours and $k = 7$ under uniform selection

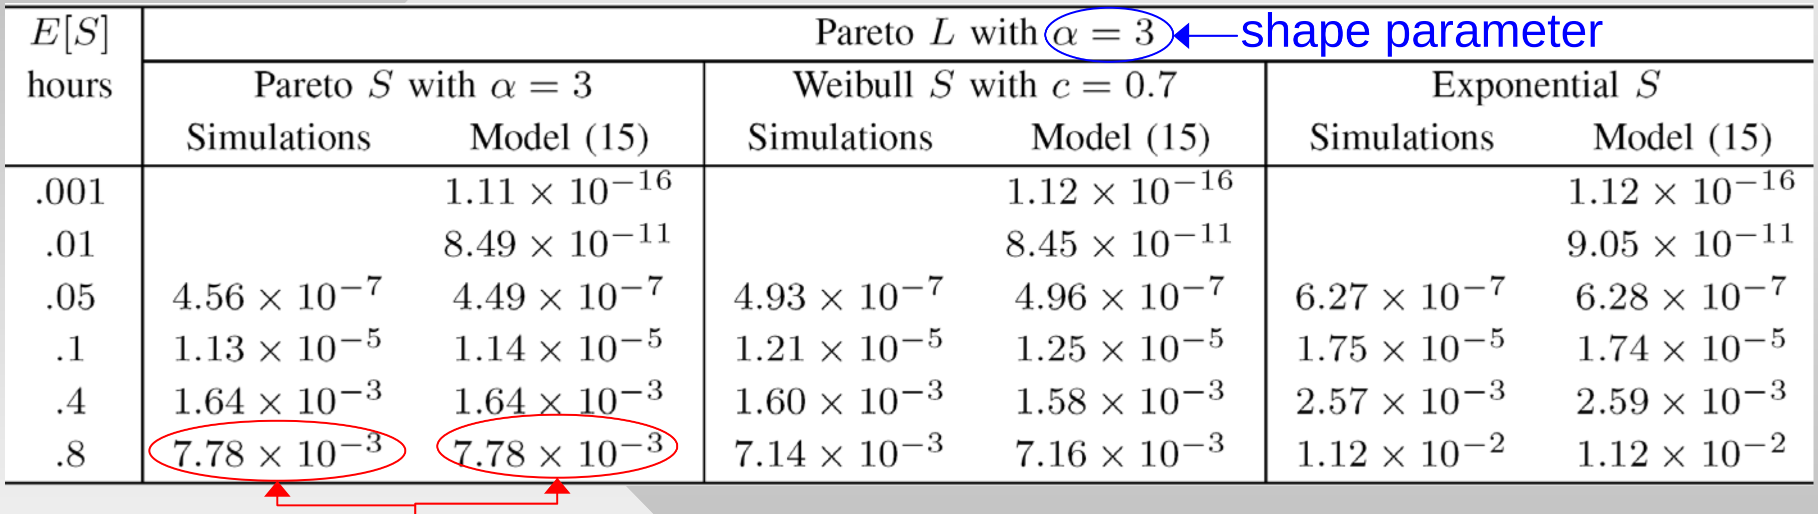

$E[S]$ hours	Pareto L with $\alpha = 3$					
	Pareto S with $\alpha = 3$		Weibull S with $c = 0.7$		Exponential S	
	Simulations	Model (15)	Simulations	Model (15)	Simulations	Model (15)
.001		1.11×10^{-16}		1.12×10^{-16}		1.12×10^{-16}
.01		8.49×10^{-11}		8.45×10^{-11}		9.05×10^{-11}
.05	4.56×10^{-7}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

shape parameter

consistent

- Our model can be used to compute ϕ in networks with various types of lifetimes and different neighbor selection strategies
  - As long as the distribution of neighbor residual lifetimes can be approximated by a hyper-exponential distribution

Rules for Selecting Neighbors

- Higher resilience (i.e., smaller isolation probability) is achieved by selecting neighbors with larger residual lifetimes
    - When it is impossible to obtain future knowledge of user remaining lifetimes R, user age A may be used as a robust predictor of R
- In systems with heavy-tailed lifetimes (e.g., Pareto, Weibull, and Cauchy), users with larger age demonstrate stochastically larger residual lifetimes.
- For light-tailed lifetimes (e.g., uniform distributions), it is the opposite

Agenda

- Motivation and background
  - Terminology, assumptions, and related work
- Generic node isolation model
- Max-age selection
- Age-proportional random-walk selection
- Wrap-up

Basics of Max-Age Selection

- Suppose that each user v publishes its joining time t_v to its neighbors, so that they knows v's current age: $t - t_v$, where t is the current time
- Each user uniformly selects m alive users at random from the system and chooses the one with the maximal current age as its neighbor
  - Uniform selection can be implemented using special random walks on the graph (Zhong 2005)
  - When $m = 1$, the max-age approach reduces to the simple uniform approach
- Denote by U_m the residual lifetime of the user whose age is maximal among m uniformly selected peers

Neighbor Residual Lifetimes

- Theorem 3: For any heavy-tailed lifetime distribution, larger m implies a stochastically larger neighbor residual lifetime U_m:

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

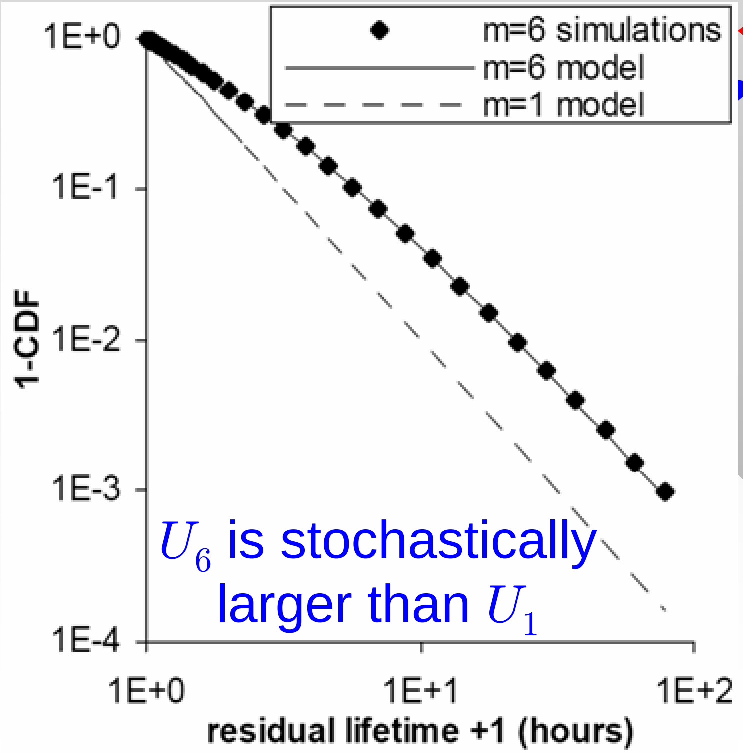

U_6 is stochastically larger than U_1

Simulation results on the tail distribution of U_m for $\alpha = 3$

See the paper for the formula of the distribution of U_m

Open question: how does m affect the obtained benefits?

Neighbor Residual Lifetimes 2

- Theorem 4: For Pareto lifetimes L with CDF $F(x) = 1 - (1 + x/\beta)^{-\alpha}$, the mean residual lifetime $E[U_m]$ is proportional to $m^{1/(\alpha - 1)}$ for $\alpha > 2$ and non-trivial m
  - If $\alpha = 3$, $E[U_m] \sim \sqrt{m}$
  - For $\alpha \to 2$, the increase in $E[U_m]$ is more aggressive:

$$E[U_m] \sim m$$

  - If $\alpha \leq 2$, the mean is infinite
- Max-age selection is much more effective in systems with more heavy-tailed lifetimes (e.g., smaller α)

Node Isolation under Max-Age Selection

- By approximating the distribution of U_m with a hyper-exponential distribution, we readily obtain isolation probability using our general node isolation model

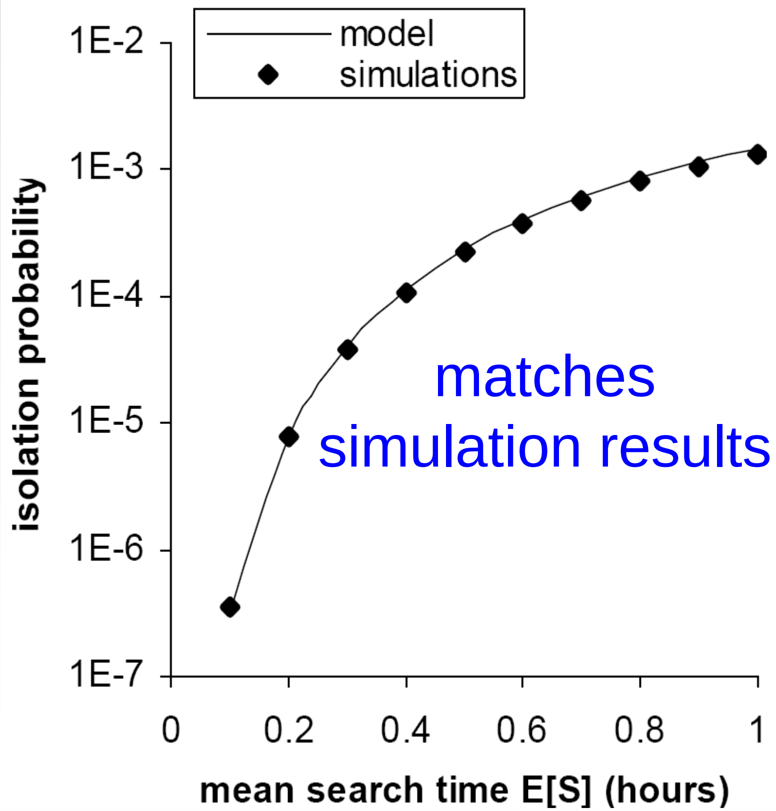

matches simulation results

$m=6$, $\alpha=3$, and $k=7$

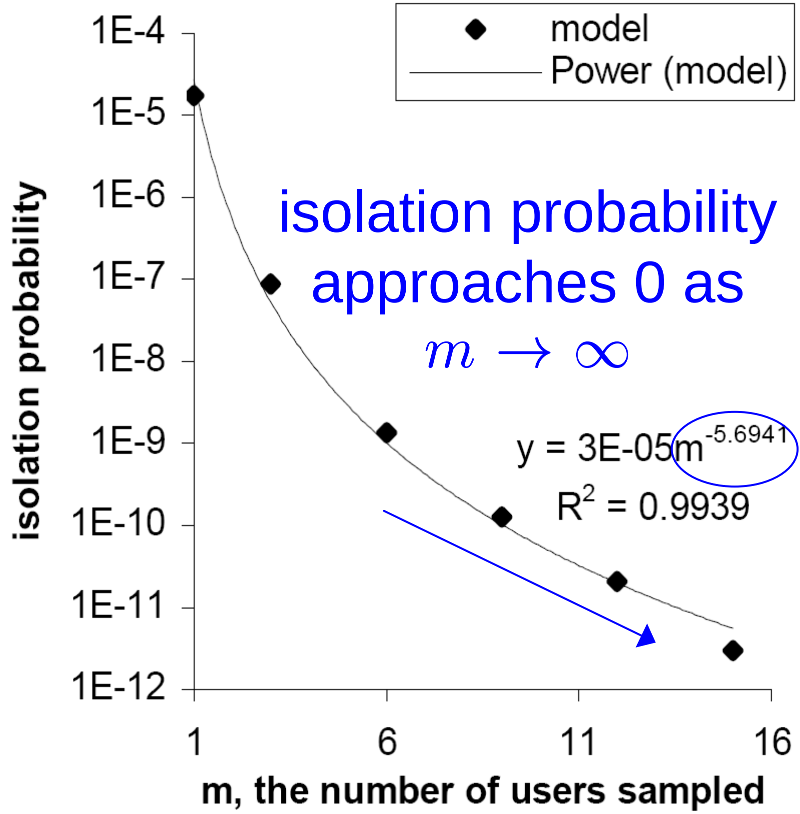

isolation probability approaches 0 as $m \to \infty$

$\alpha=2$, and $k=7$, $E[S]=6$ mins

Discussion

- The max-age selection strategy requires sampling m users per link
  - The overhead may not scale well for large m
- Much higher resilience can be achieved by more aggressively preferring users with large age
- We thus next propose a more efficient and effective neighbor selection strategy for heavy-tailed lifetimes

Agenda

- Motivation and background
  - Terminology, assumptions, and related work
- Generic node isolation model
- Max-age selection
- Age-proportional random-walk selection
- Wrap-up

Basics of Age-Proportional Selection

- We introduce a new age-biased neighbor selection method to ensure that the probability that user v is selected by another peer is proportional to its current age A_v:

$$P(v \text{ is selected}) = \frac{A_v}{\sum_{u \in \boxed{V}} A_u}$$

the set including all existing users

- This approach is based on random walks on directed and weighted graphs
  - It provides a distributed solution that requires only one sample per link

Random Walks

- Assume that each user makes its current age and in-degree known to its in-degree neighbors
  - The weight of each link is determined by the current age and in-degree of the node that the link points to

$$\text{weight} = \frac{\text{age of } v}{\text{in-degree of } v}$$

- Random walks are performed by alternating between walking along incoming and outgoing links
  - The probability that a link is chosen is proportional to its weight
  - The stationary distribution of the random-walk algorithm is:

$$P(v \text{ is selected}) = \frac{A_v}{\sum_{u \in V} A_u}$$

← achieves the desired result

Neighbor Residual Lifetimes

- Theorem 5: For random-walks where the above stationary distribution holds, the tail distribution of residual lifetimes R of selected neighbors is:

$$P(R > x) = \frac{1}{E[L]E[A]} \int_0^\infty y(1 - F(x+y))dy$$

the mean age

the user lifetime distribution

- For Pareto lifetimes $F(x)=1-(1+x/\beta)^{-\alpha}$, $\alpha > 2$, the above yields:

$$P(R > x) = \left(1 + \frac{x}{\beta}\right)^{-(\alpha-2)}$$

The shape is reduced by 2

- The mean $E[R]$ is $\beta/(\alpha-3)$ if $\alpha > 3$, and is infinite otherwise

Node Isolation Probability for $\alpha > 2$

- Simulation results on node isolation probability under age-proportional selection for Pareto L and $k = 7$

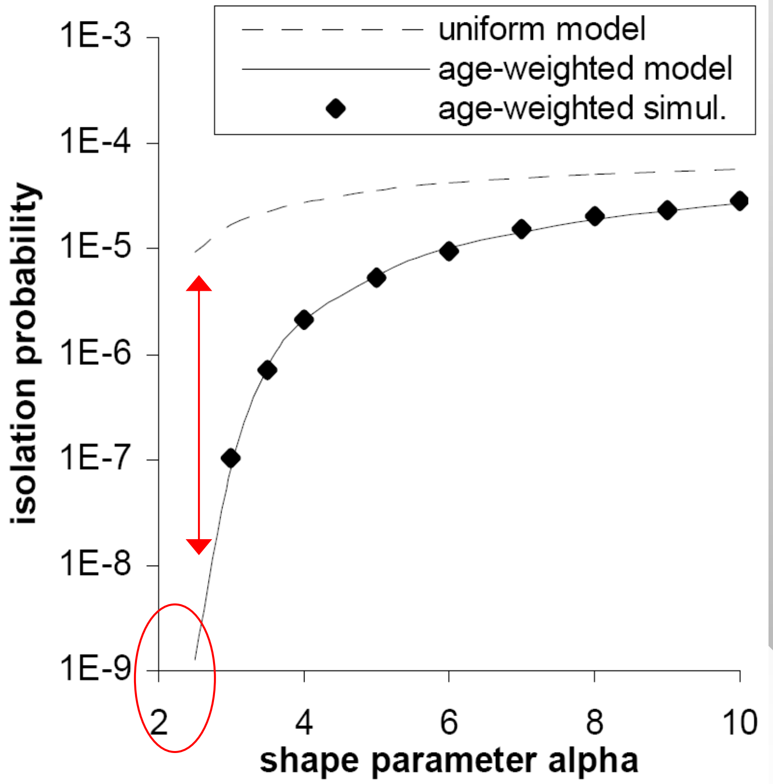

(a) $E[S] = 6$ minutes

  - Isolation probability is 10^4 times smaller than that under uniform selection for $\alpha = 2.5$
  - Isolation probability converges to 0 as $\alpha \to 2$
  - In contrast, this metric under max-age selection does not tend to 0 unless $m \to \infty$ or $\alpha \to 1$ (both impossible to achieve in practice)

Node Isolation Probability for $\alpha \le 2$

- Theorem 4: For age-proportional random walks, Pareto lifetimes with $1 < \alpha \le 2$, any number of neighbors $k \ge 1$, and any type of search delay (including $\mathbb{S} = \infty$), as system age $\mathcal{T}$ and size n converge to infinity, node isolation probability approaches:

$$\lim_{n\to\infty} \lim_{\mathcal{T}\to\infty} \phi = 0$$

- Gnutella has been shown to have α between 1.06 (Bustamante 2003) and 1.09 (Wang 2007)
  - These networks under age-proportional random walks approach an ideal system with zero node isolation probability as users join/depart the system

Node Isolation Probability for $\alpha \leq 2$

- Simulation results of node isolation probability without replacing neighbors (i.e., $S = \infty$) for Pareto lifetimes

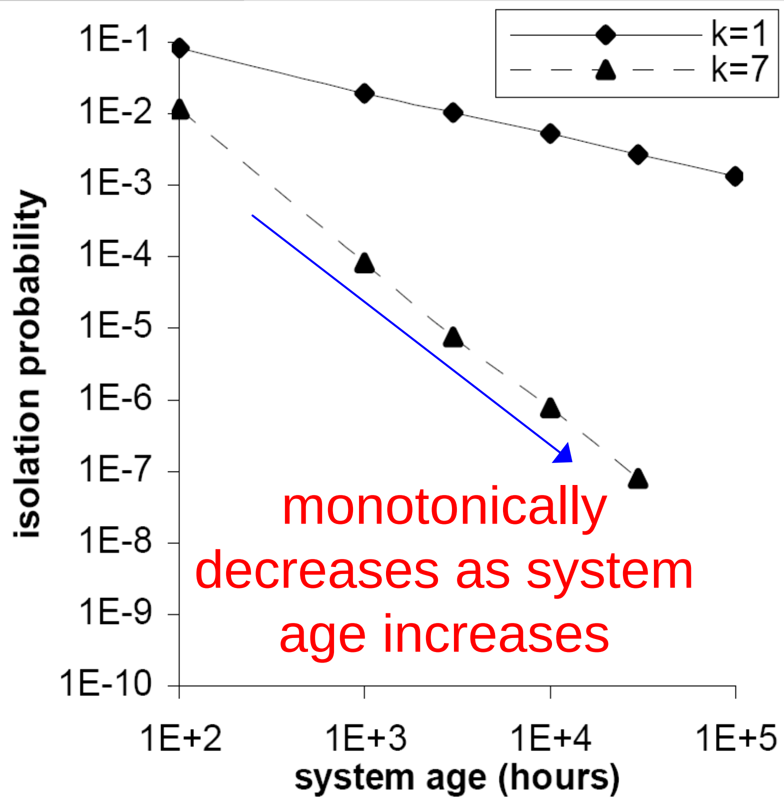

monotonically decreases as system age increases

(a) $\alpha = 1.5, S = \infty$

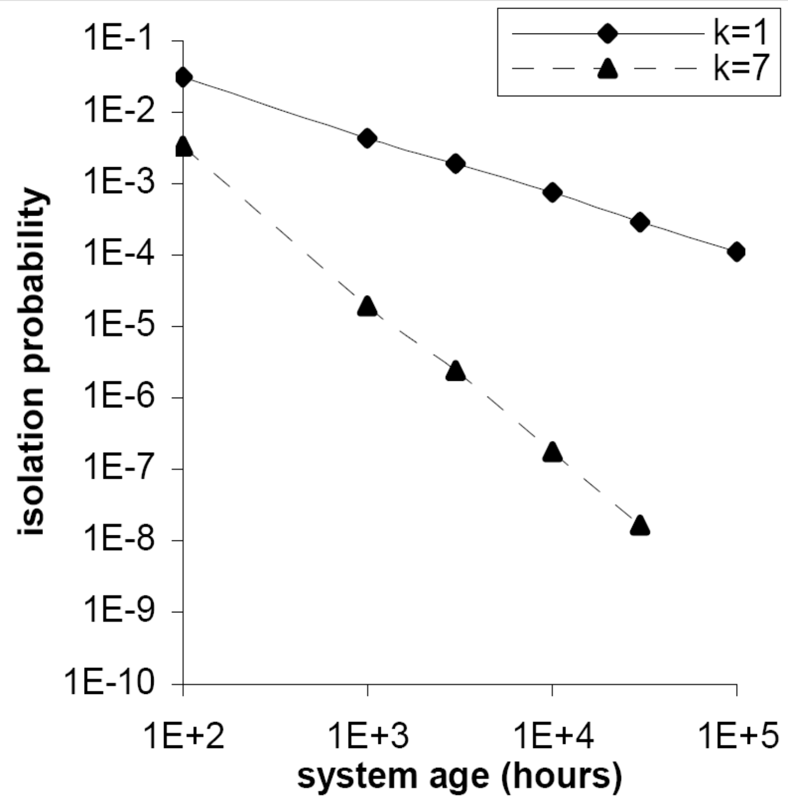

(b) $\alpha = 1.2, S = \infty$

Wrap-up

- We developed a general node isolation model for any completely monotone density function of neighbor residual lifetimes
  - We applied this model to study node isolation behavior under uniform, max-age, or age-proportional random-walk selection to demonstrate its versatility
- We proposed a new neighbor selection strategy, age-proportional random walks
  - Under proposed neighbor selection, P2P networks with heavy-tailed lifetimes with $\alpha \leq 2$ become progressively more resilient over time and approach a system with zero node isolation probability, as more users join the system