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

Zhongmei Yao[*], Daren B.H. Cline[**], and Dmitri Loguinov[**]

[*] The University of Dayton
[**] Texas A&M University

IEEE P2P, London, England
9-10-2014

Agenda

- Introduction
  - Motivation
- Understanding the tradeoff
  - Dynamic graph model
  - Performance measures
- Idle fraction
- Degree of selected neighbors
- Resilience and degree-overload tradeoff
- Conclusion

Introduction

- Non-memoryless node lifetime distributions $F_L(x)$ allow us to utilize the knowledge of user age A to predict residual lifetimes V
  - Reliable users are the ones with longer residual lifetimes V

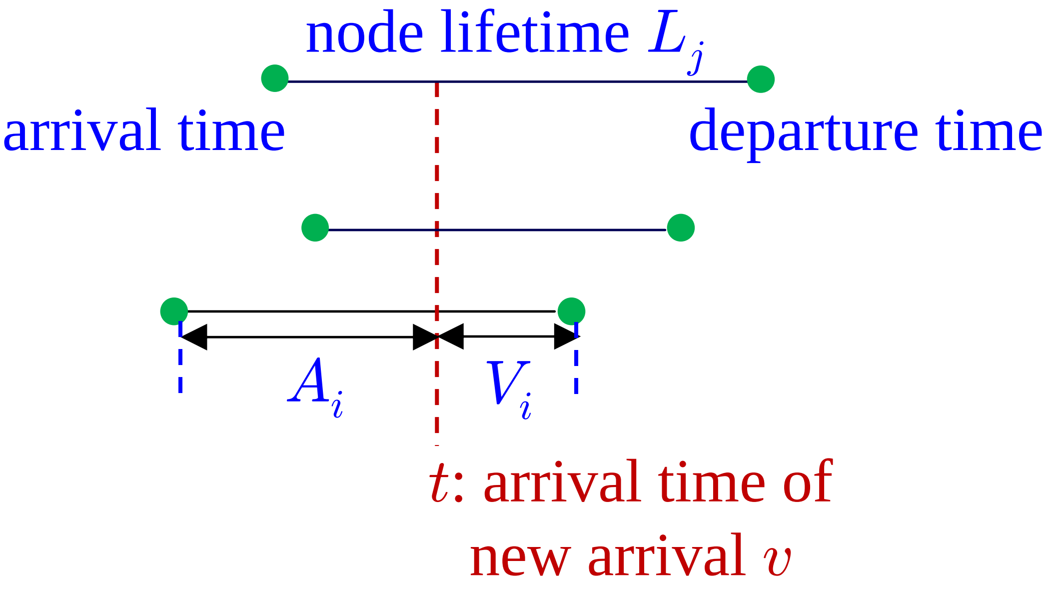

node lifetime L_j

arrival time

departure time

A_i

V_i

t: arrival time of new arrival v

  - Given heavy-tailed $F_L(x)$ (Pamies-Juarez 2010, Wang 2009), larger A implies stochastically larger V
    - New worse than used (NWU)
    - The opposite if $F_L(x)$ is light-tailed

Introduction 2

- Age-biased selection is implemented by using a general neighbor preference function $w(x)$
  - User v assigns non-negative weight $w(x)$ to users with age x
  - The probability $c_N(i)$ that v connects to a live peer i in a system of N users is proportional to $w(A_i)$

$$c_N(i) = P(v \to i \mid A_1, \ldots, A_{N-1}) \sim w(A_i)$$

age of a live user

N: number of users currently alive

- The residual lifetime distribution under $w(x)$:

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

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

Motivation

- Analysis of P2P systems is a well-studied area
  - Poisson arrivals and exponential lifetimes
  - On/off arrival/departure processes and non-exponential lifetimes
- Under general lifetime distributions, previous work studied passive and active systems
  - Active: k out-degree neighbors for routing. Failed ones are replaced
  - Passive: Broken connections are never re-connected. Inbound/outbound links are used for routing. Passive systems are surprisingly appealing:
    - Good resilience: users are well protected via out-links long enough for in-degree to take over
    - Simple operations: no edge rewiring, no keep-alive messages
- Age-biased selection allows us to send more links to old users, thus improving resilience
  - Is it possible that an $O(1)$ fraction of the system is forced to handle $\Theta(n)$ of load when we make $w(x)$ very aggressive?
  - Does there exist an optimal $w(x)$?

Agenda

- Introduction
  - Motivation
- Understanding the tradeoff
  - Dynamic graph model
  - Performance measures
- Idle fraction
- Degree of selected neighbors
- Resilience and degree overload tradeoff
- Conclusion

Dynamic Graph Model

- Passive systems
  - Out-link lifetimes $V_1, V_1 ,\ldots , V_k \sim F_V(x)$
  - In-link lifetimes $L_j \sim F_L(x)$

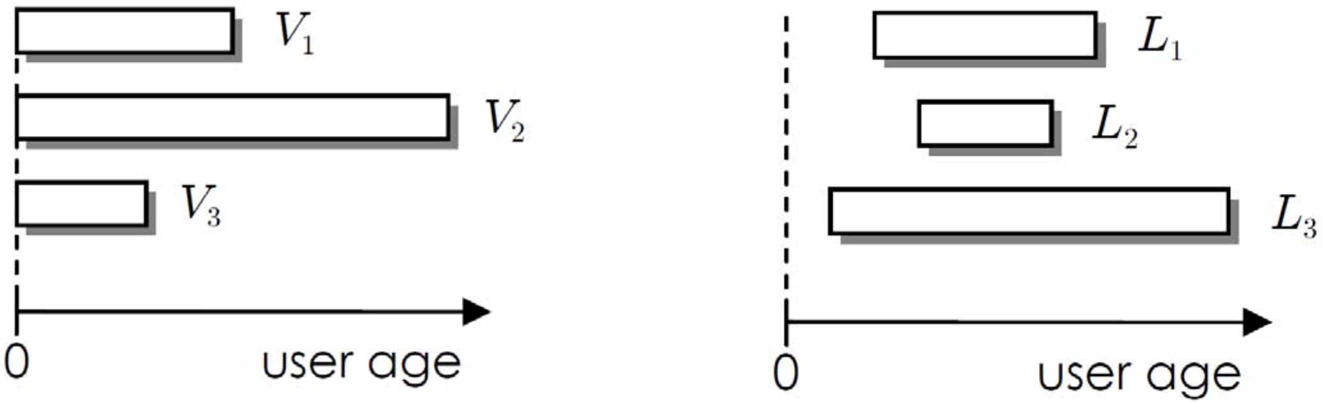

- Out-degree at age τ is a binomial random variable

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

- In-degree at age τ is a Poisson random variable

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

Performance Measures

- Treat degree $D(\tau)$ at age τ as a non-absorbing process
- Two metrics:
  - Measuring resilience: define idle fraction as the expected fraction of time a user's degree is zero within its lifetime

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

  - Indication of overload: let Y be the age of a user at the time it was selected. We focus on $E[D(Y)]$, the expected degree of selected users
- Weight functions
  - Uniform weight: $w(x) = 1$
  - Max-age weight: $w(x) = mF_A(x)^{m\text{-}1}$
  - Age-proportional: $w(x) = x$
  - Step function: $w(x) = 1$ if $x \geq x_0$; otherwise $w(x) = 0$
  - Truncated power function: $w(x) = \min((x/x_0)^\rho, 1)$

Performance Measures 2

- The Pareto-optimal curve
  - Each point is driven by some combination $(k, w(.))$
  - Point A is dominated by B since $x_A = x_B$, $y_A > y_B$
  - Define a point to be Pareto-optimal if it is dominated by no other point

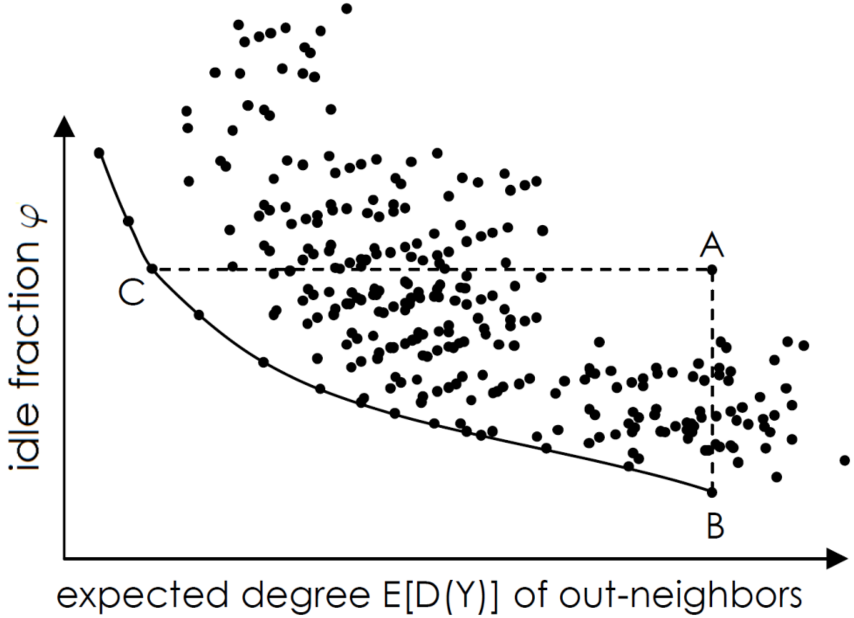

- The goal of our optimization problem
  - To obtain the best weight function that places points only along the Pareto-optimal curve
  - That is, achieve the smallest idle fraction for a given $E[D(Y)] = d$

Agenda

- Introduction
  - Motivation
- Understanding the tradeoff
  - Dynamic graph model
  - Performance measures
- Idle fraction
- Degree of selected neighbors
- Resilience and degree overload tradeoff
- Conclusion

Idle Fraction

- Theorem 1: The idle fraction can be reduced to

$$\varphi = P(D(A) = 0)$$

A: age of a random live user

- Under uniform weight, we obtain

$$\varphi = \int_0^1 \left(ye^{-y}\right)^k dy \le e^{-k}$$

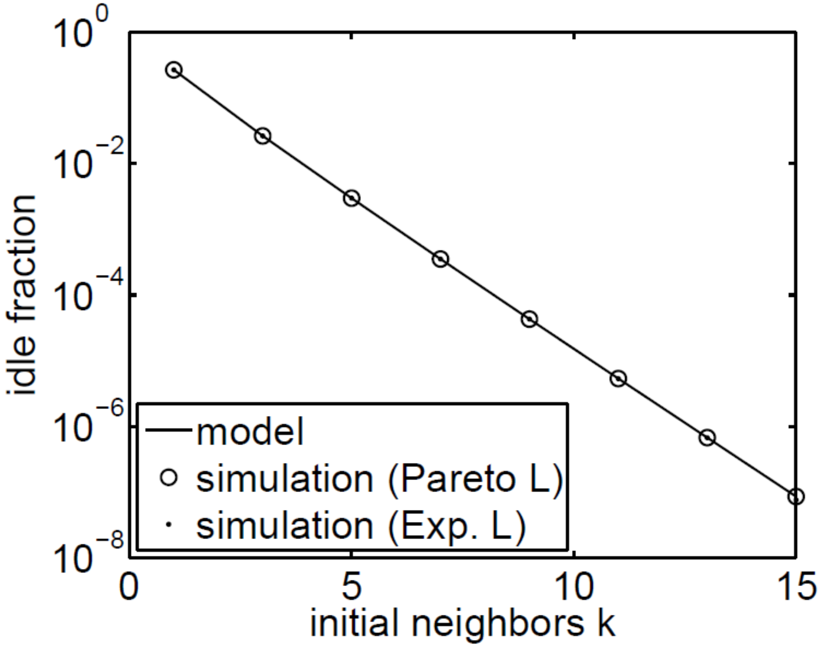

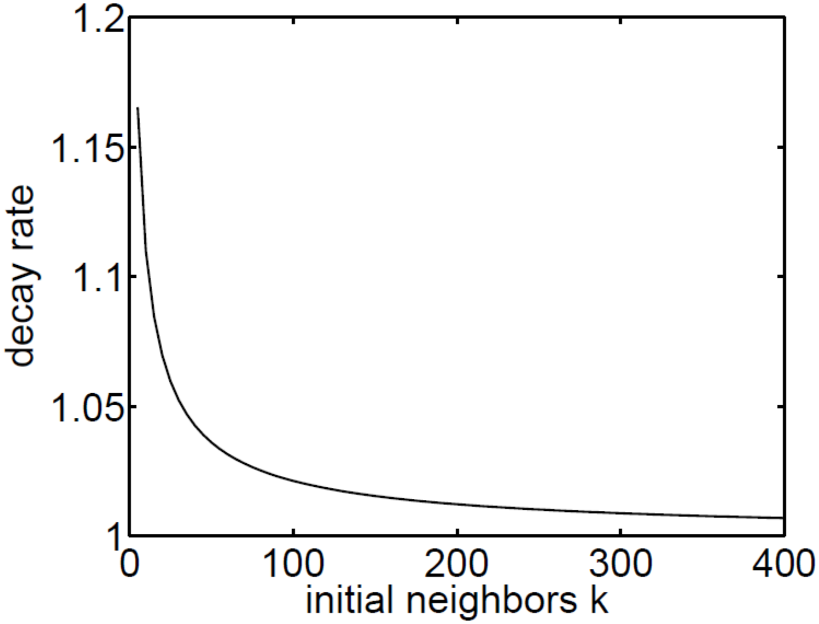

Idle Fraction 2

- Step functions
  - (a) Exponential lifetimes, (b) NWU lifetimes

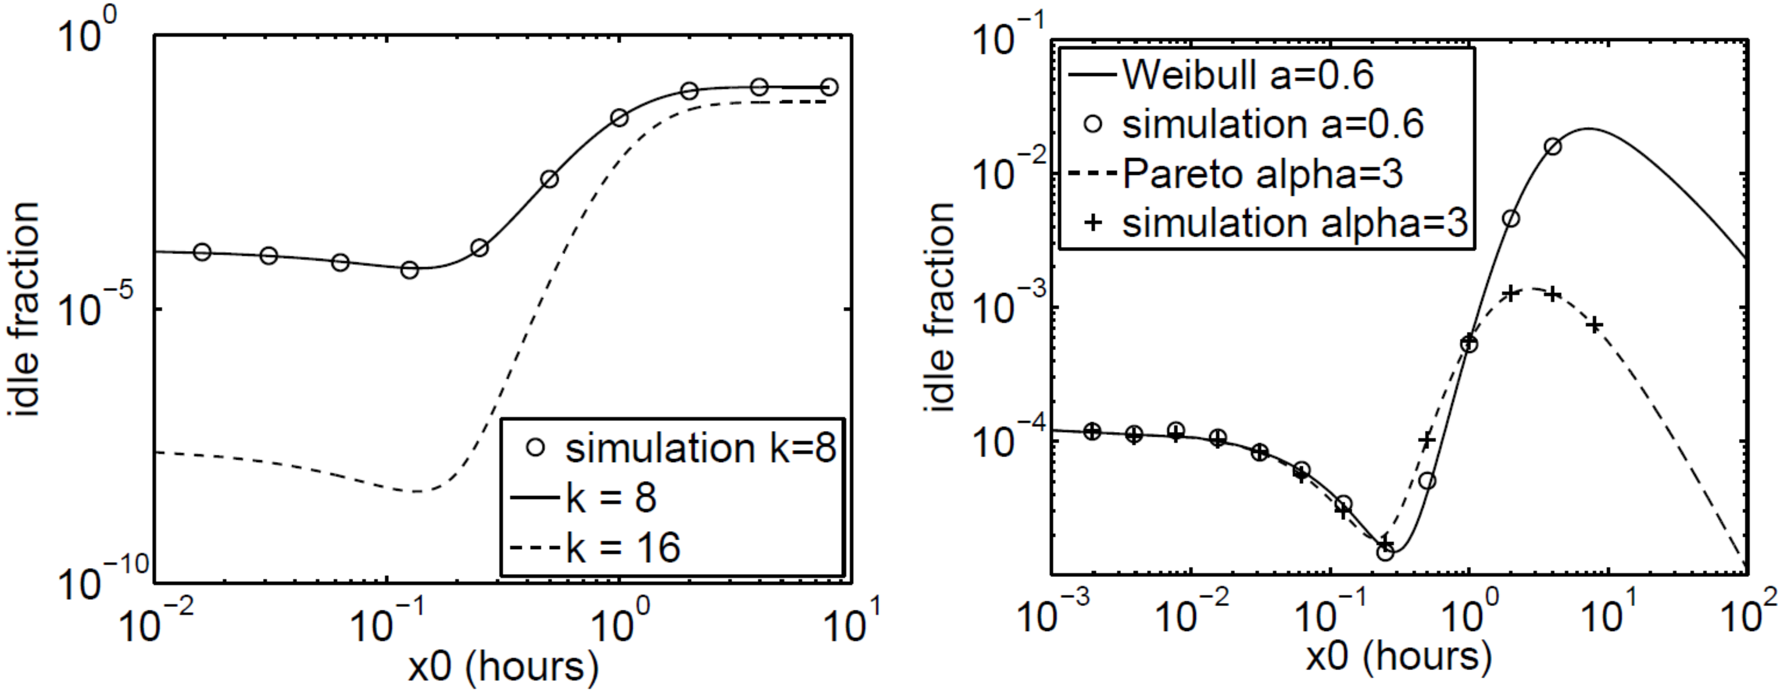

Idle Fraction 3

- Max-age weight, truncated power weight for NWU L

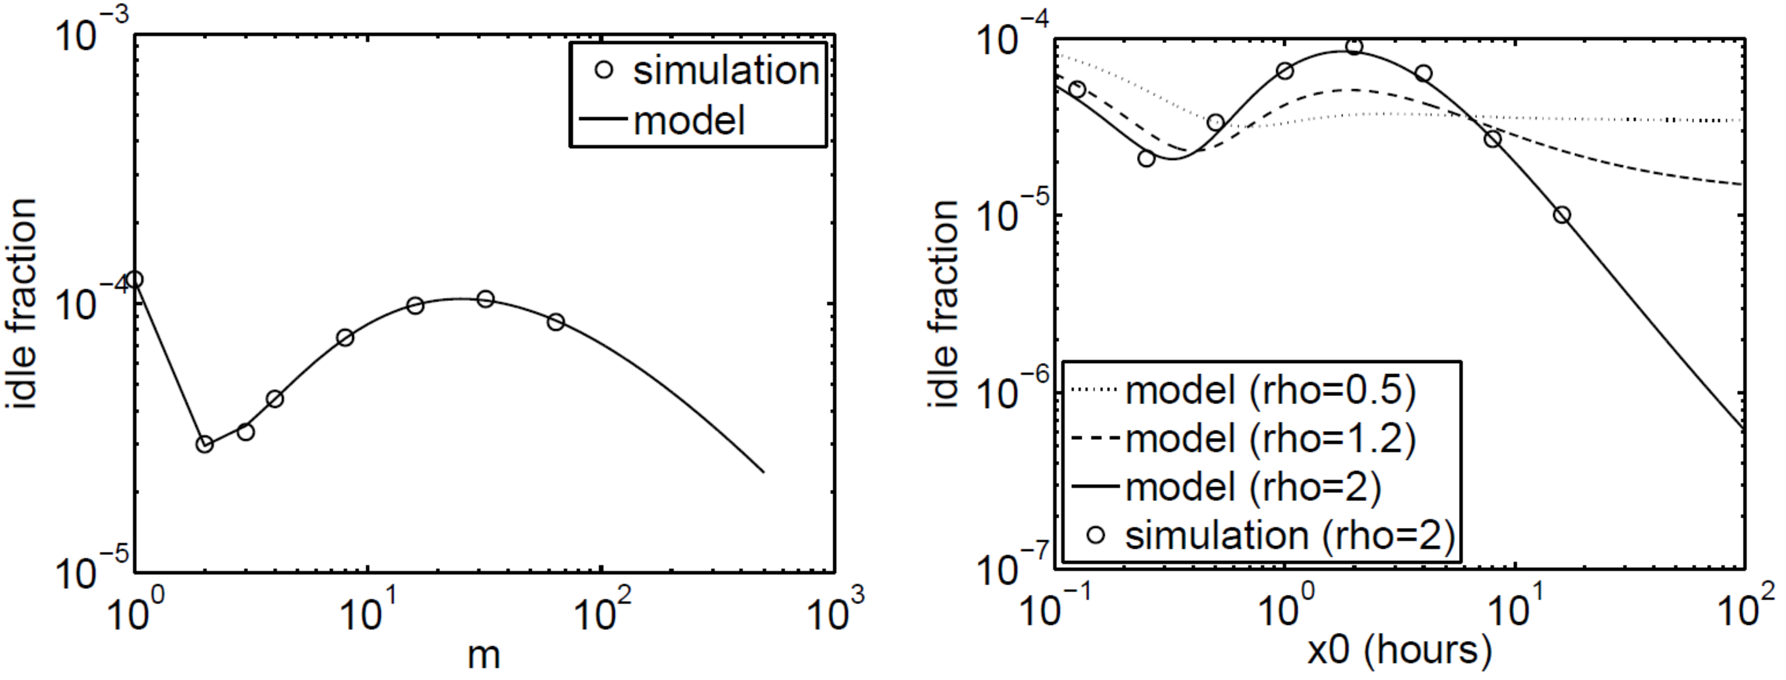

Asymptotic Decay Rate

- Theorem 6: For Pareto L , step weight, and $k > \alpha - 1$, we have

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

  - For large x_0, the idle fraction follows the tail of the age distribution, i.e., $\varphi = \Theta({x_0}^{1-\alpha})$
  - The power-law decay rate of idle fraction is slow compared to the exponential e^{-k}
  - Using extremely large x_0 is never beneficial; instead, the optimal technique is to set x_0 to ${x_0}^*$ and then increase k until the desired resilience is reached
  - The open issue is which weight function is the winner

Agenda

- Introduction
  - Motivation
- Understanding the tradeoff
  - Dynamic graph model
  - Performance measures
- Idle fraction
- Degree of selected neighbors
- Resilience and degree overload tradeoff
- Conclusion

Degree of Selected Users

- Theorem 7: The mean degree of a selected user is

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

A_1, A_2, A are iid with $F_A(x)$

  - Exponential L vs Pareto L:

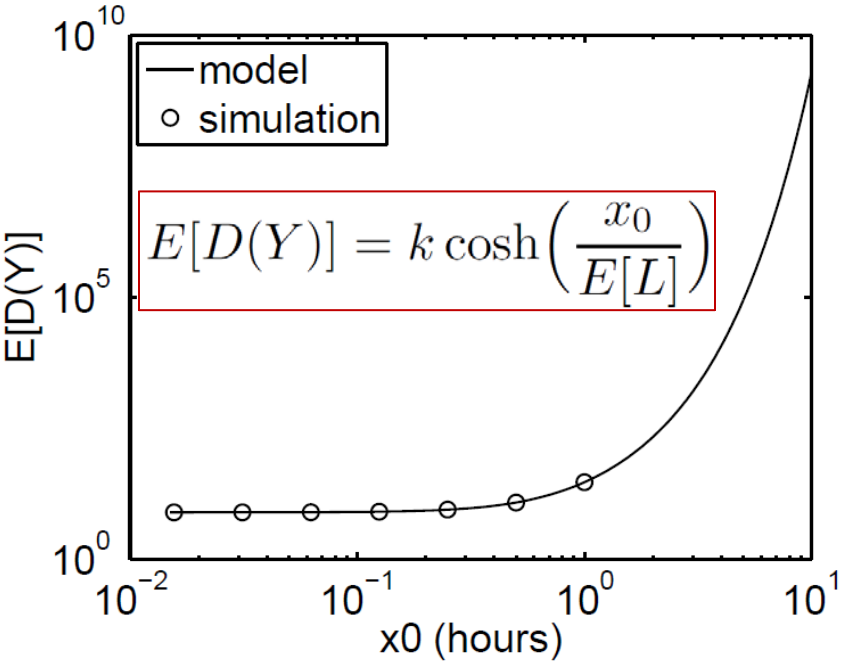

$$E[D(Y)] = k\cosh\left(\frac{x_0}{E[L]}\right)$$

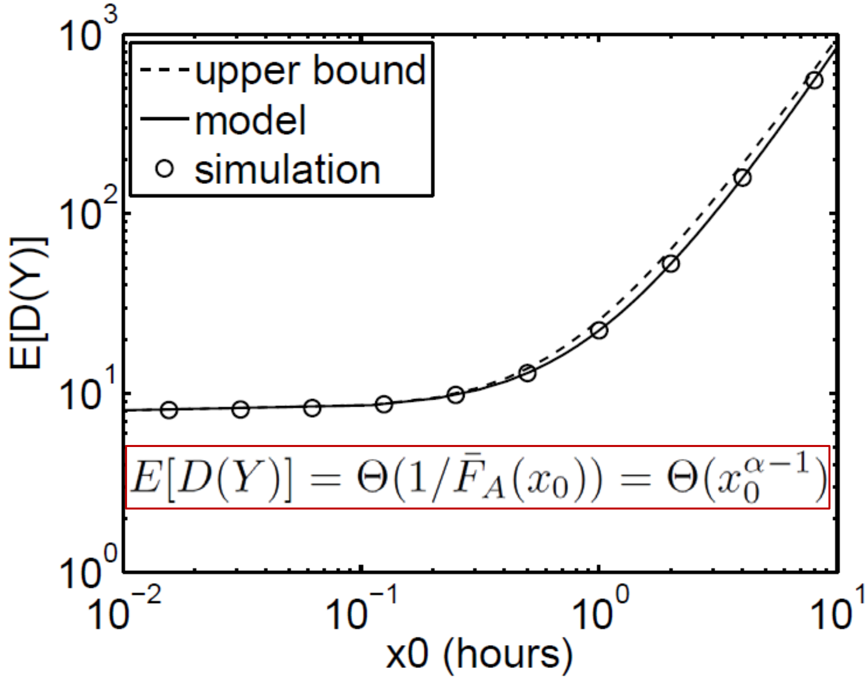

$$E[D(Y)] = \Theta(1/\bar{F}_A(x_0)) = \Theta(x_0^{\alpha-1})$$

Degree of Selected Users 2

- Max-age weight vs truncated power weight

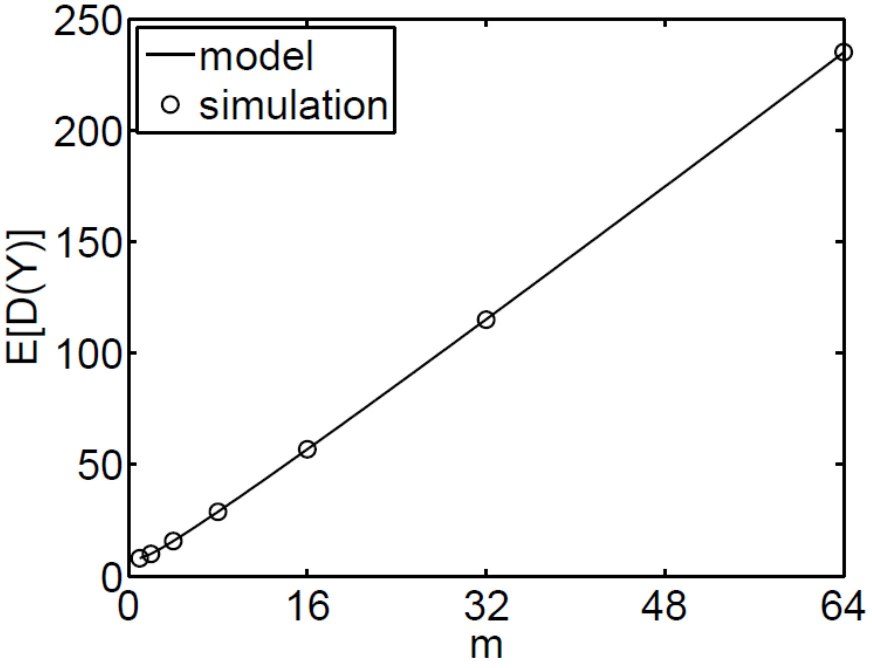

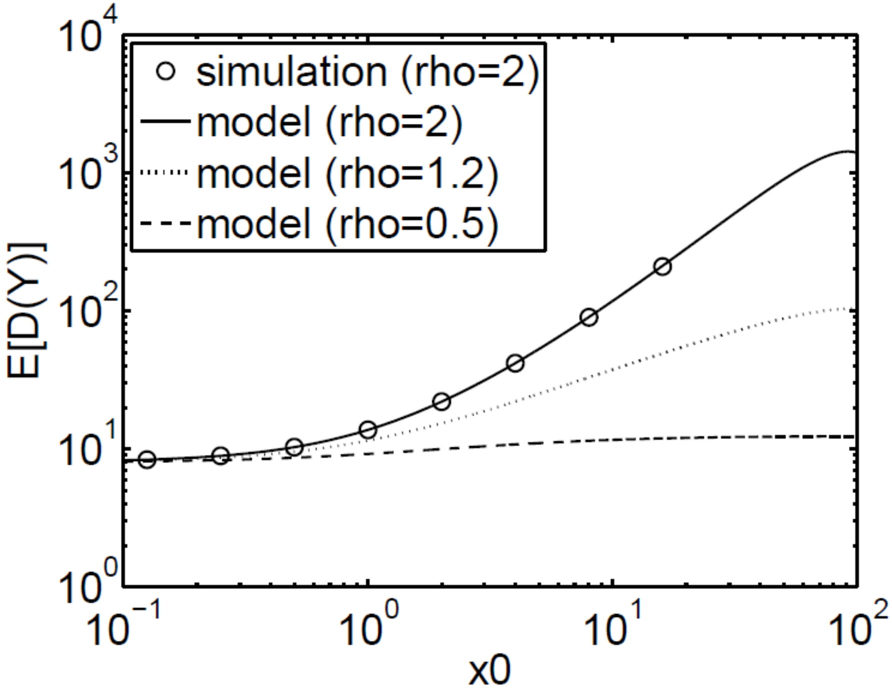

Agenda

- Introduction
  - Motivation
- Understanding the tradeoff
  - Dynamic graph model
  - Performance measures
- Idle fraction
- Degree of selected neighbors
- Resilience and degree overload tradeoff
- Conclusion

The Objective Function

- Suppose θ is some parameter of $w(.)$ that we aim to optimize
- Define the expected degree under parameter θ and $k = 1$:

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

- The objective function that we need to minimize is

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

Results

- Uniform, max-age, and step weight

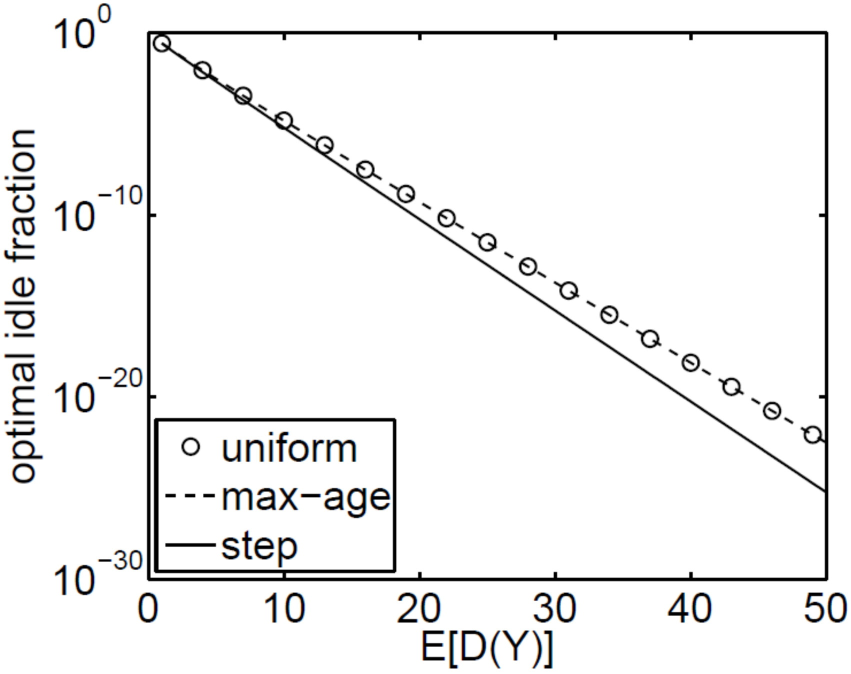

(a) $\alpha = 3$

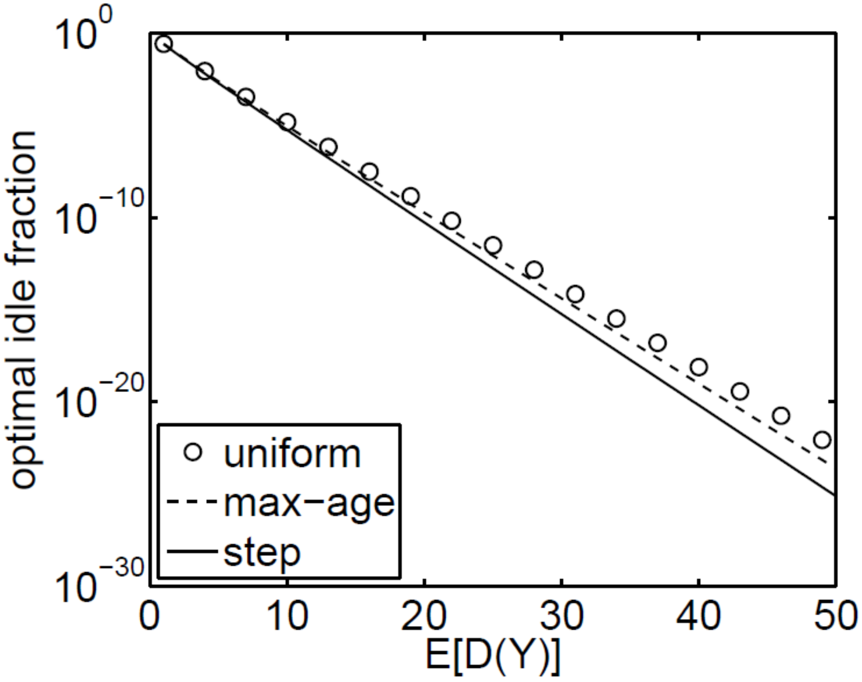

(b) $\alpha = 1.5$

Results 2

- Uniform, age-proportional, step, and truncated power

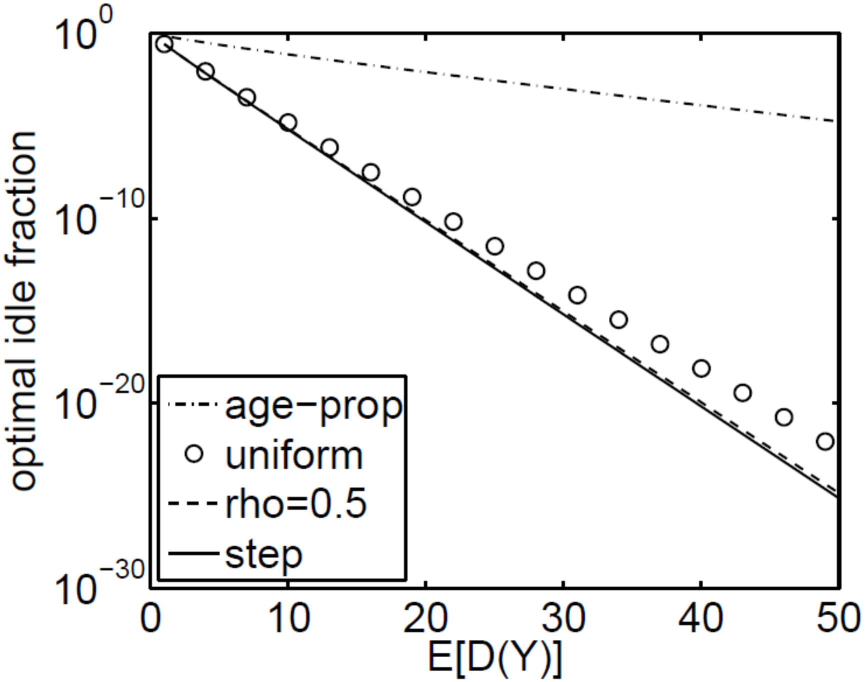

(c) $\alpha = 3.5$

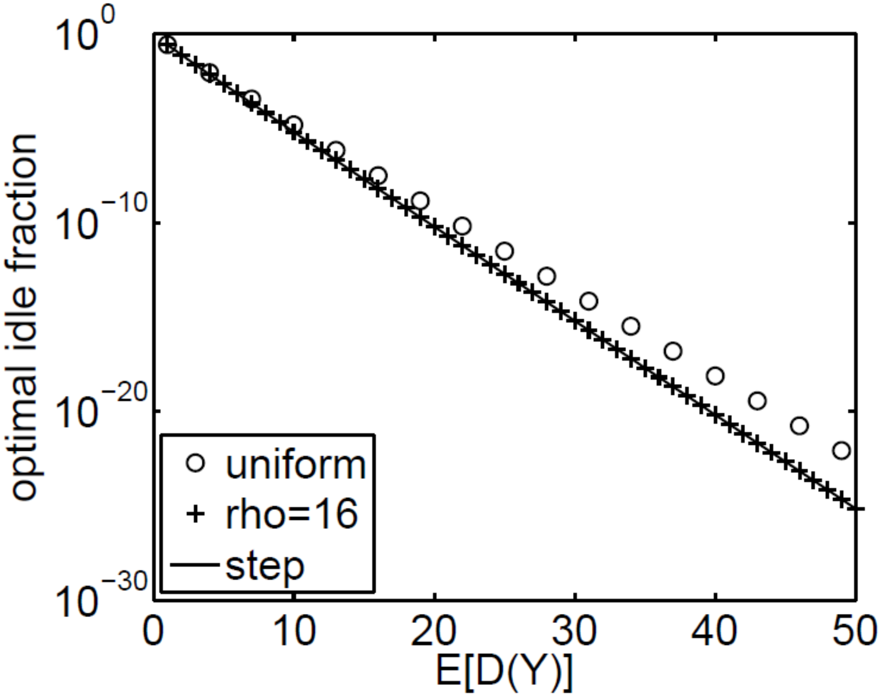

(d) $\alpha = 3.5$

Results 3

- Any non-decreasing function can be represented as a sum of step-functions
- Random mixture
  - Generate mixture of 20 random step-weights and examine result

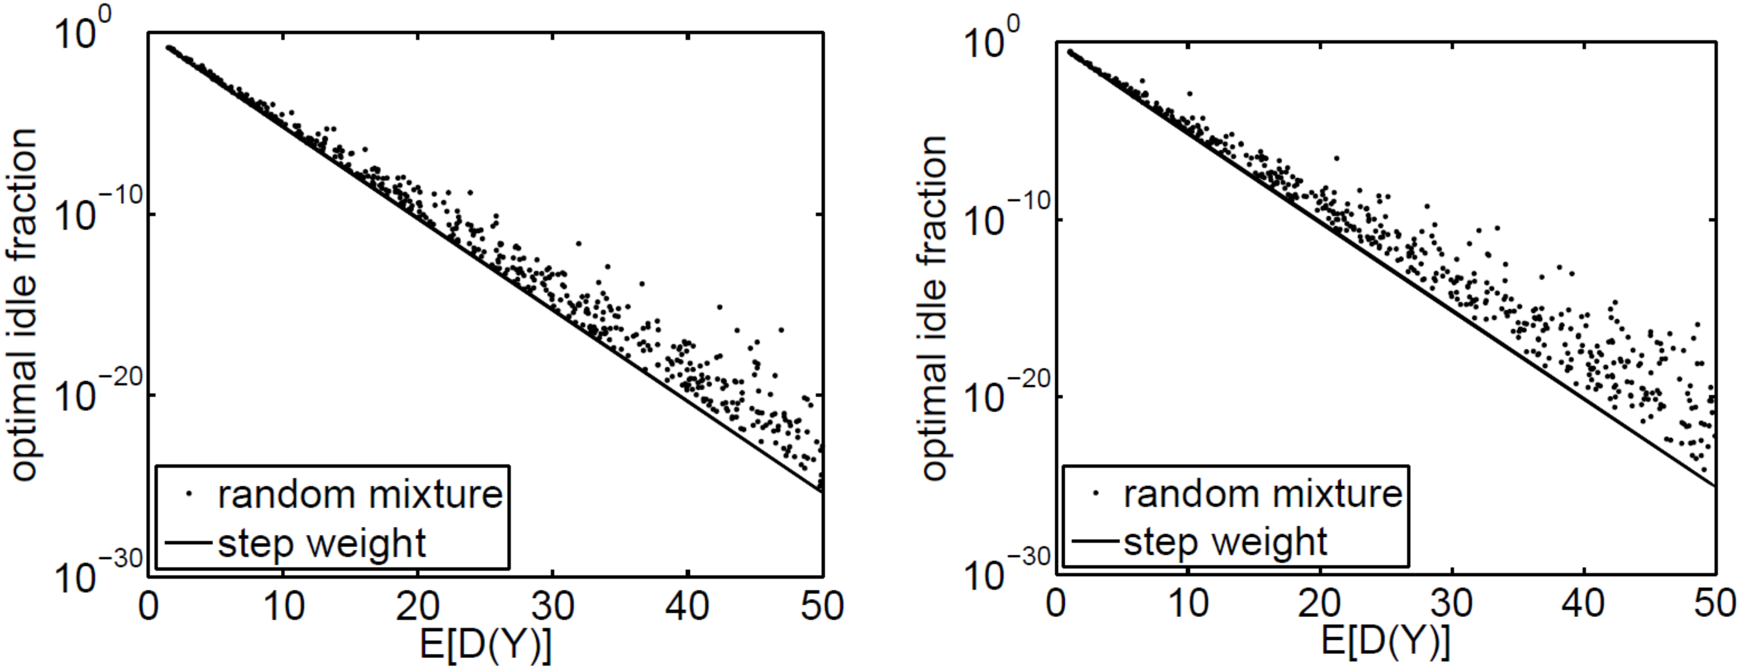

(a) $\alpha = 2.5$ (b) $\alpha = 4$

Agenda

- Introduction
  - Motivation
- Understanding the tradeoff
  - Dynamic graph model
  - Performance measures
- Idle fraction
- Degree of selected neighbors
- Resilience and degree overload tradeoff
- Conclusion

Conclusion

- We introduced models measuring resilience and load-balancing
  - Analyzed both uniform and non-uniform neighbor selection strategies in passive P2P systems
- We formulated a tradeoff problem
  - Given the constraint on degree of a randomly selected user, what was the best weight function that maximized resilience
- We showed that among the methods studied here, the step function was the clear winner