Residual-Based Measurement of Peer and Link Lifetimes in Gnutella Networks

Xiaoming Wang

Joint work with Zhongmei Yao and Dmitri Loguinov

Internet Research Lab
Computer Science Department
Texas A&M University

May 8, 2007

Agenda

- Introduction
- Related work
  - Create-Based Method (CBM)
- Analysis of CBM
- Proposed method
  - Residual-Based Estimator (RIDE)
- Comparison of overhead
- Experiments
- Conclusion

Introduction

- Current peer-to-peer networks
  - Have a large number of participating users
  - Are organized in a decentralized infrastructure
- Measurements are needed for validation purposes
  - Lifetime distribution — our focus
  - Inter-arrival delays
  - Availability
  - Topological information
  - Traffic flow rate

Introduction – Lifetime Sampling

- Measuring lifetime distribution
  - Collecting lifetime samples
- The lifetime of a user can be measured only if one observes both its birth and death
- To detect arrivals and departures, one must have periodic snapshots of the whole system

Introduction – Lifetime Measurement

- Take snapshots of the system every Δ time units
- Consider a user with lifetime L
  - The user is captured in the snapshots taken at time points $t_i, t_{i+1}, \ldots, t_j$
  - Then, the measured lifetime of this user is rounded up to $t_j - t_i + \Delta$
- Create-Based Method (CBM)

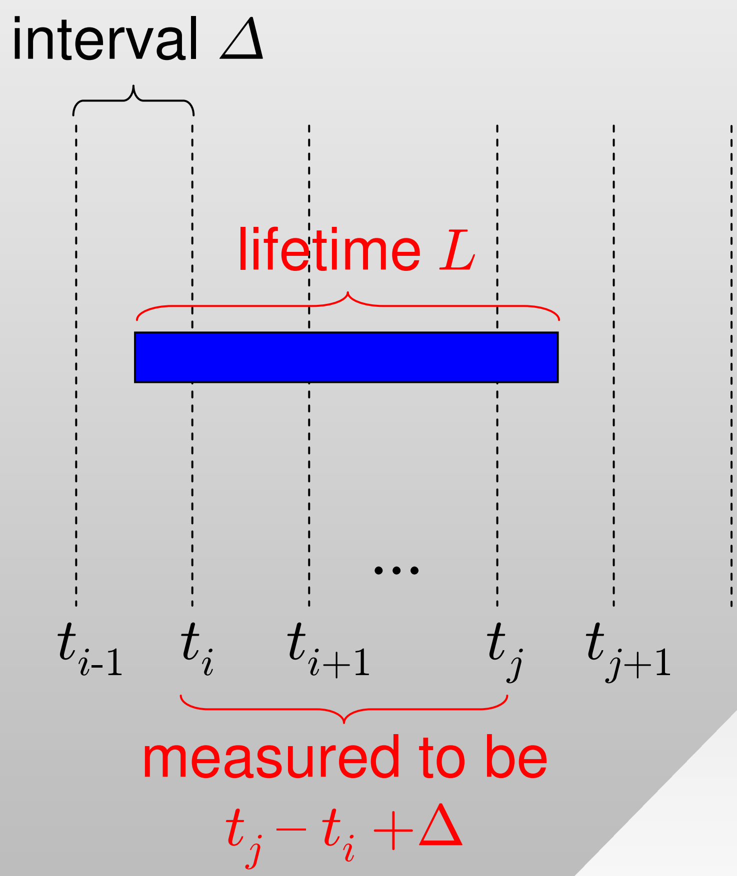

Agenda

- Introduction
- Related work
  - Create-Based Method (CBM)
- Analysis of CBM
- Proposed method
  - Residual-Based Estimator (RIDE)
- Comparison of overhead
- Experiments
- Conclusion

Related Work – CBM Sampling

1. Appear during the first half of the window

2. Disappear somewhere in the window

3. Stay in the system no longer than T time units

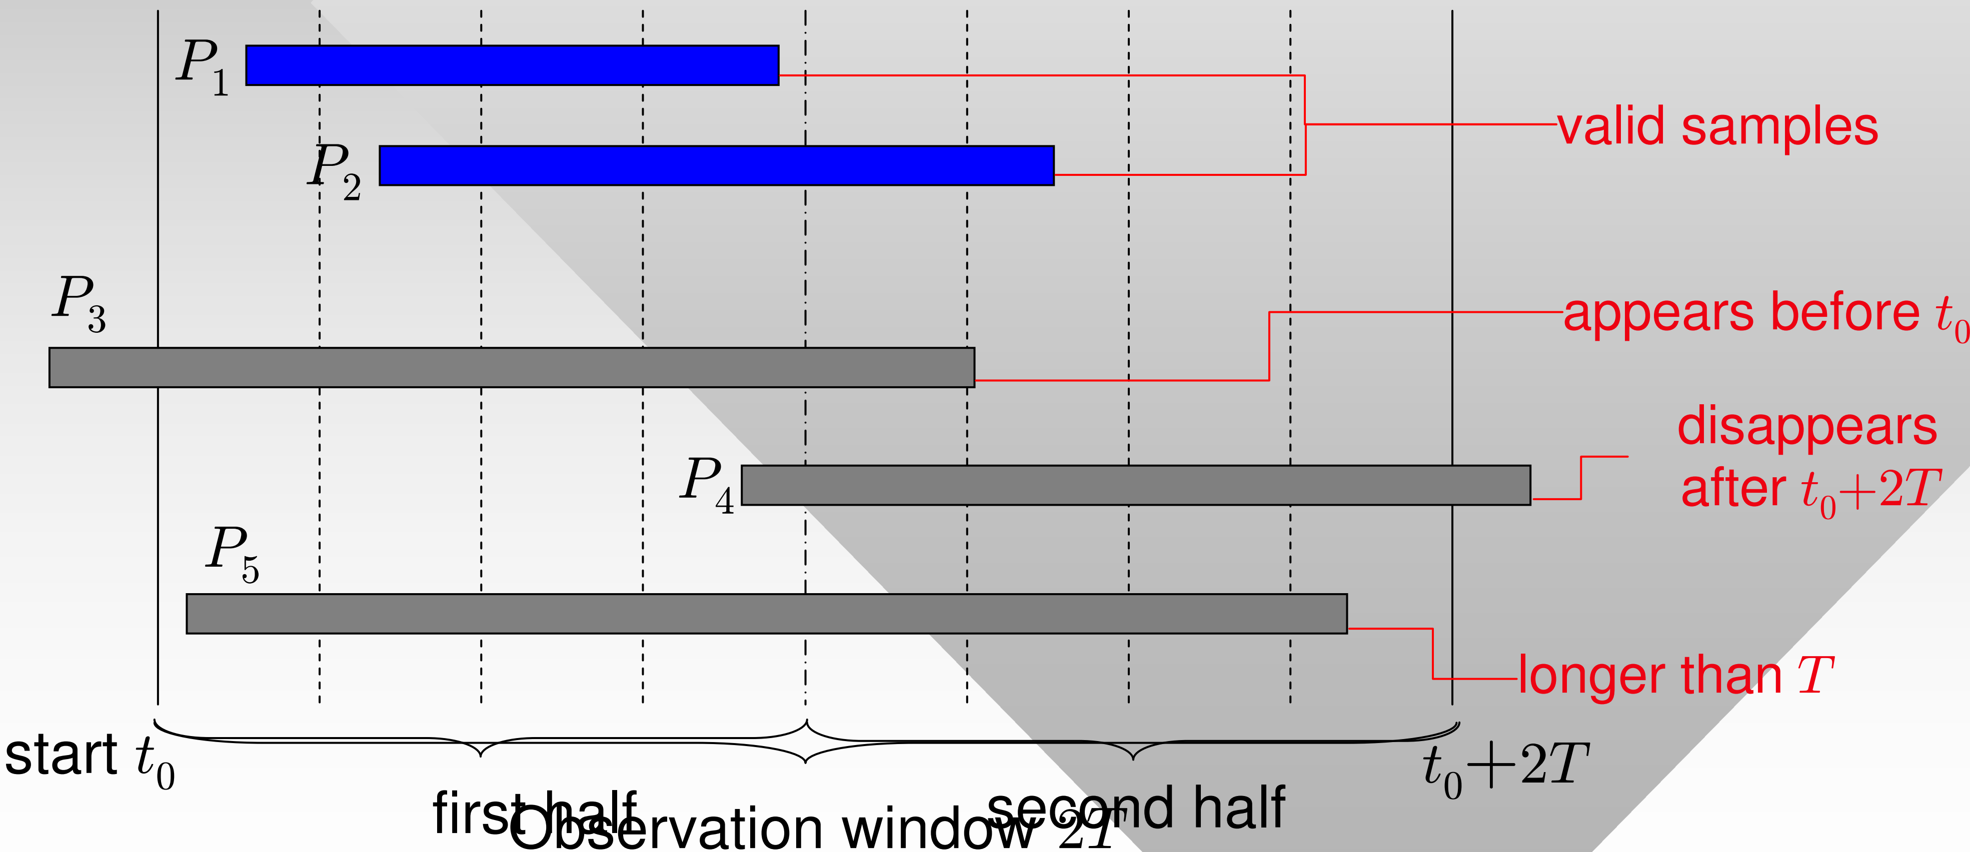

Related Work – CBM Estimators

criteria of unbiased estimator

CDF function $F(x)$

- Denote by $E(x)$ the estimated value for the probability a random lifetime L is no larger than x
  - We want $E(x) = P(L \leq x) = F(x)$ to hold for discrete points $x = j\Delta$, $j = 0, 1, 2, \ldots, T/\Delta$

- Two CBM estimators have been proposed
  - Roselli 2000

$$E_A(x_j) = \frac{N(x_j)}{N(T)}$$

$x_j = j\Delta$

valid samples smaller than x_j

total valid samples

  - Saroiu 2002, Bustamante 2003, Stutzbach 2006

$$E_B(x_j) = \frac{N(x_j)}{N}$$

total observed samples that appear in the first half window

Agenda

- Introduction
- Related work
  - Create-Based Method (CBM)
- Analysis of CBM
- Proposed method
  - Residual-Based Estimator (RIDE)
- Comparison of overhead
- Experiments
- Conclusion

CBM – Round-off Errors

- Round-off errors may cause bias in CBM

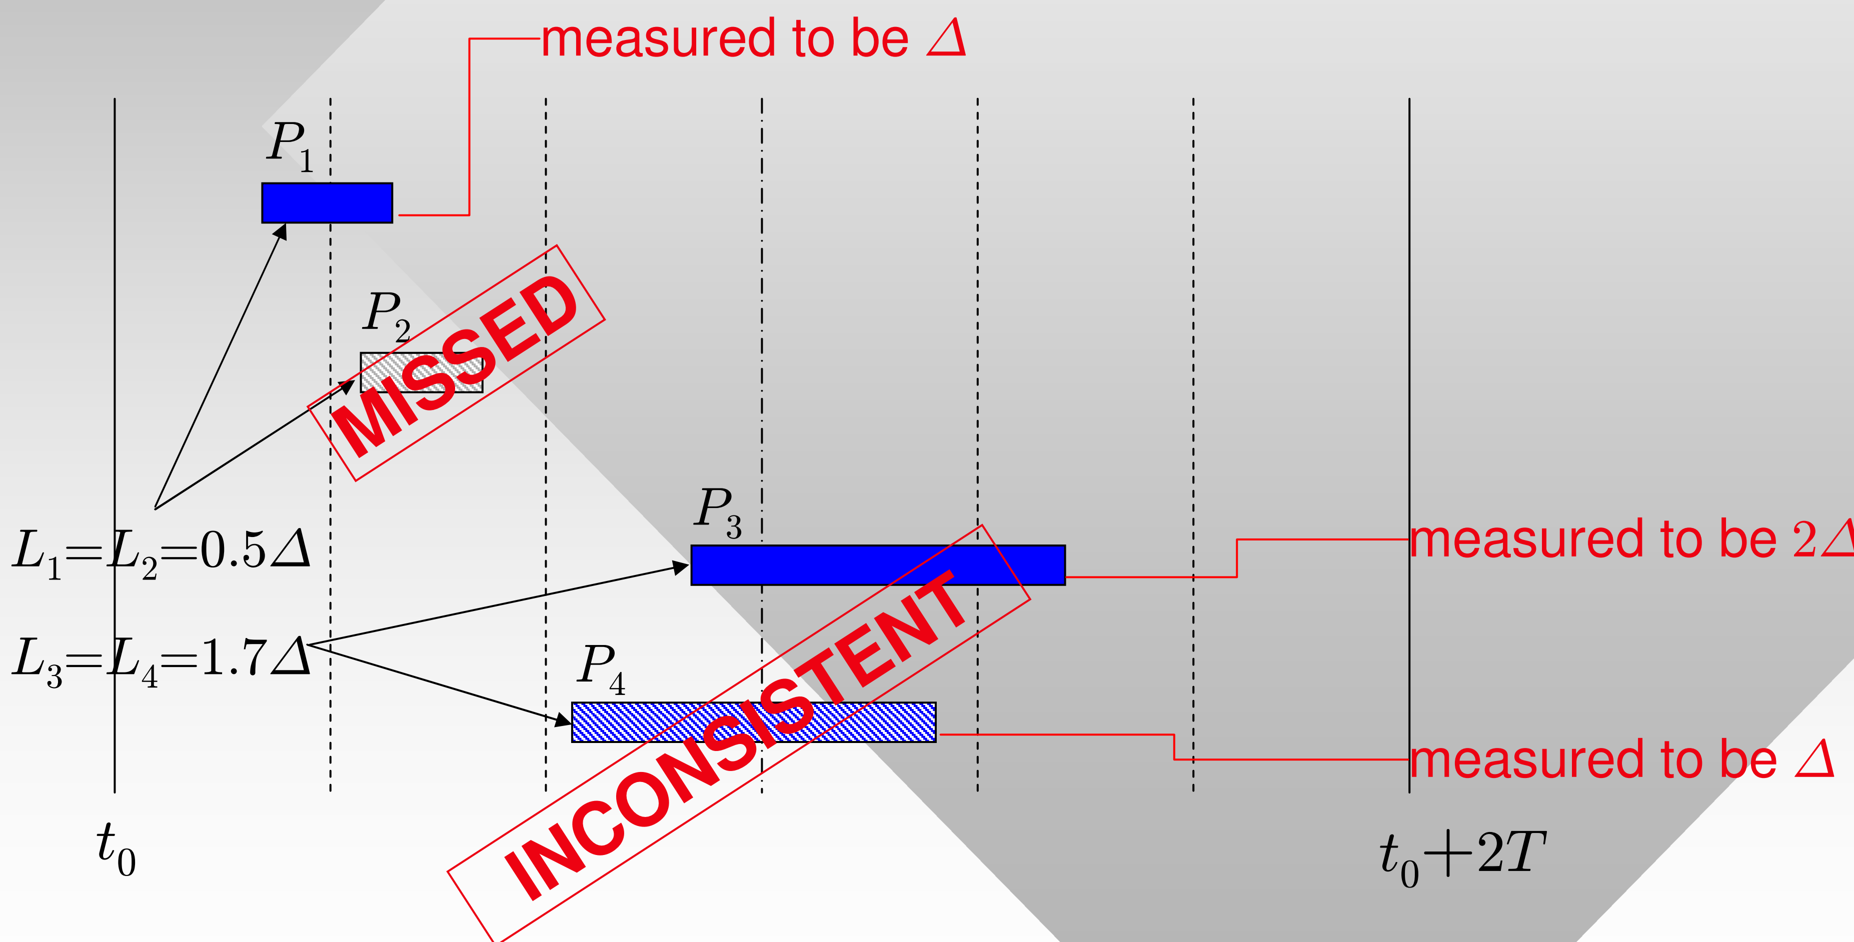

CBM – Model

target $P(L \leq x_j)$

probability of round-off errors

not affected by T more accurate

- Theorem 1: Both CBM estimators differ from the actual lifetime distribution as follows

$$E_A(x_j) = \frac{F(x_j)+\rho_j-\rho_0}{F(T)-\rho_0} \qquad E_B(x_j) = \frac{F(x_j)-\rho_0+\rho_j}{1-\rho_0}$$

probability $P(L \leq T)$

  – where

$$\rho_j = \frac{1}{\Delta}\int_0^{\Delta} F(x+x_j)dx - F(x_j)$$

- In fact, actual $F(x)$ can be recovered from E_B only if ρ_j are known
  – However, ρ_j can be neither measured in practice nor calculated without knowing $F(x)$
- Therefore, E_B as well as E_A are inherently biased

CBM – Simulations with E_A

$$F(x) = 1 - (1 + x/\beta)^{-\alpha}$$

- Actual lifetimes follow a Pareto distribution
  - α=1.1, β=0.05, $E[L]$=0.5 hours, T=24 hours

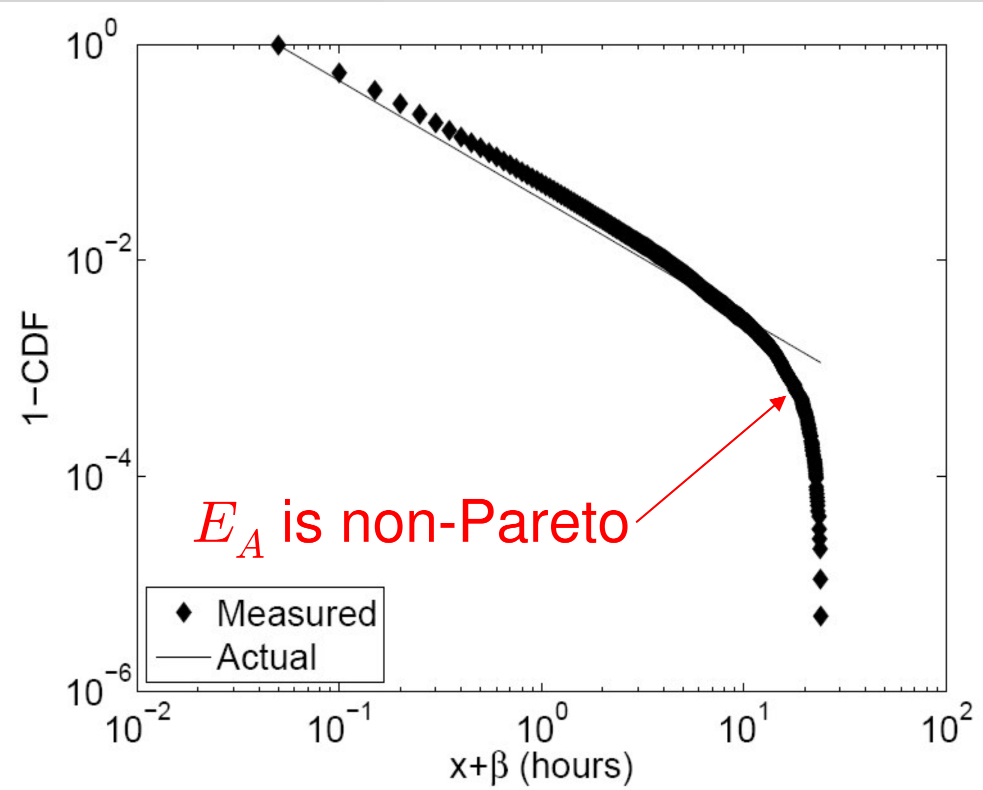

Δ=3 min

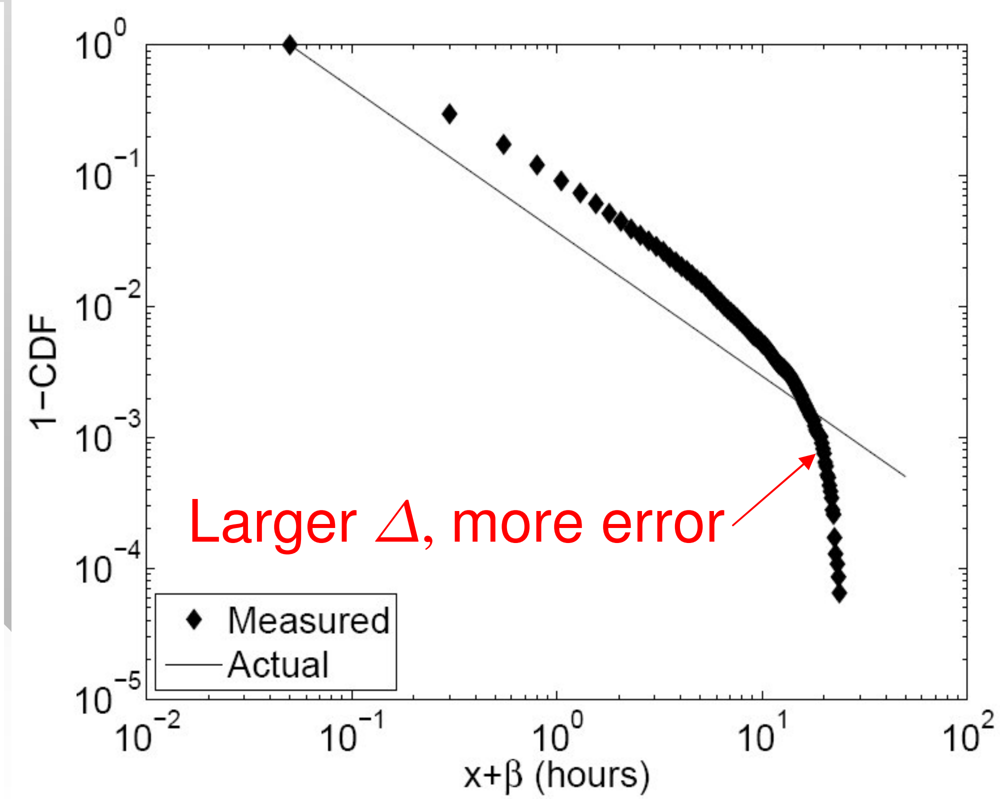

Δ=15 min

CBM – Simulations with E_B

- Estimator E_B preserves the Pareto shape α for small ρ_j
  - But makes the Pareto scale β inaccurate

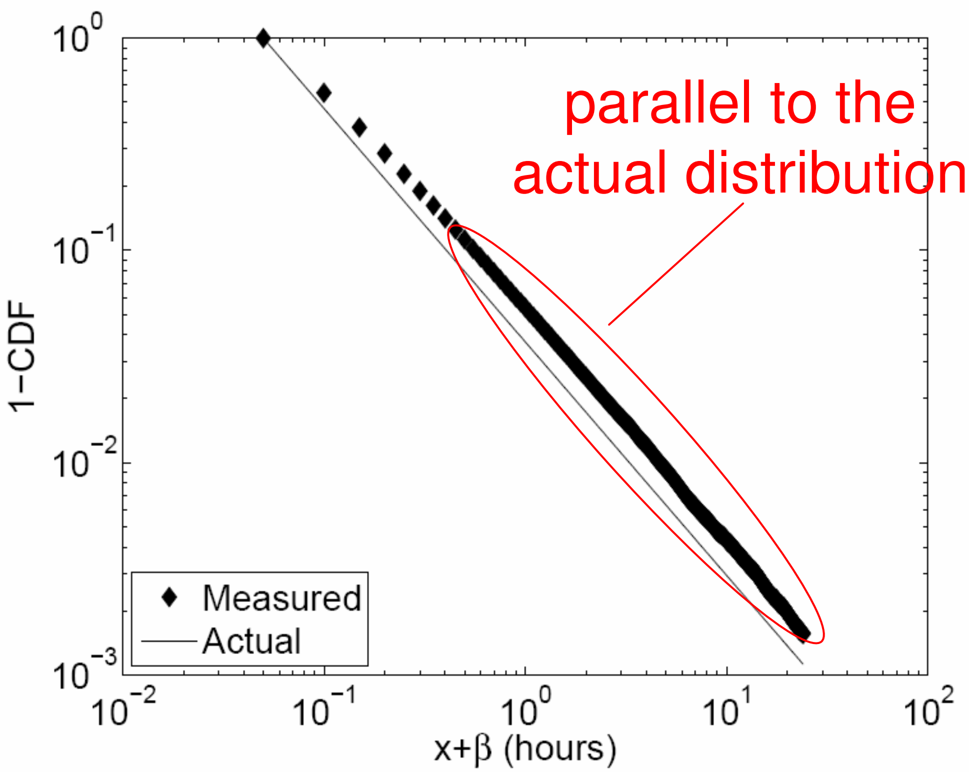

$\Delta=3$ min

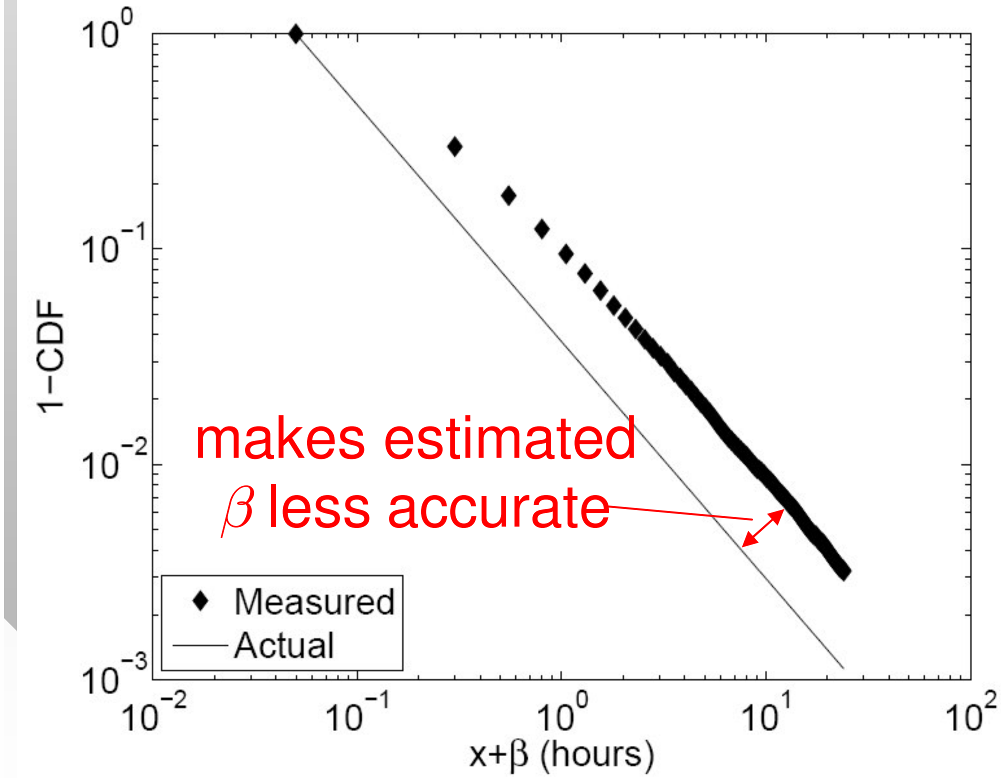

$\Delta=15$ min

Agenda

- Introduction
- Related work
  - Create-Based Method (CBM)
- Analysis of CBM
- Proposed method
  - Residual-Based Estimator (RIDE)
- Comparison of overhead
- Experiments
- Conclusion

Residual-Based Method – Motivation

- Larger Δ implies more bias in CBM
  - It is desirable to crawl the system as frequently as possible
- Smaller Δ means more traffic overhead
  - Moreover, Δ cannot be smaller than the time needed for crawling the entire system
- There is an inherent tradeoff between accuracy and overhead in CBM
- We propose Residual-Based Estimator (RIDE)

RIDE – Sampling

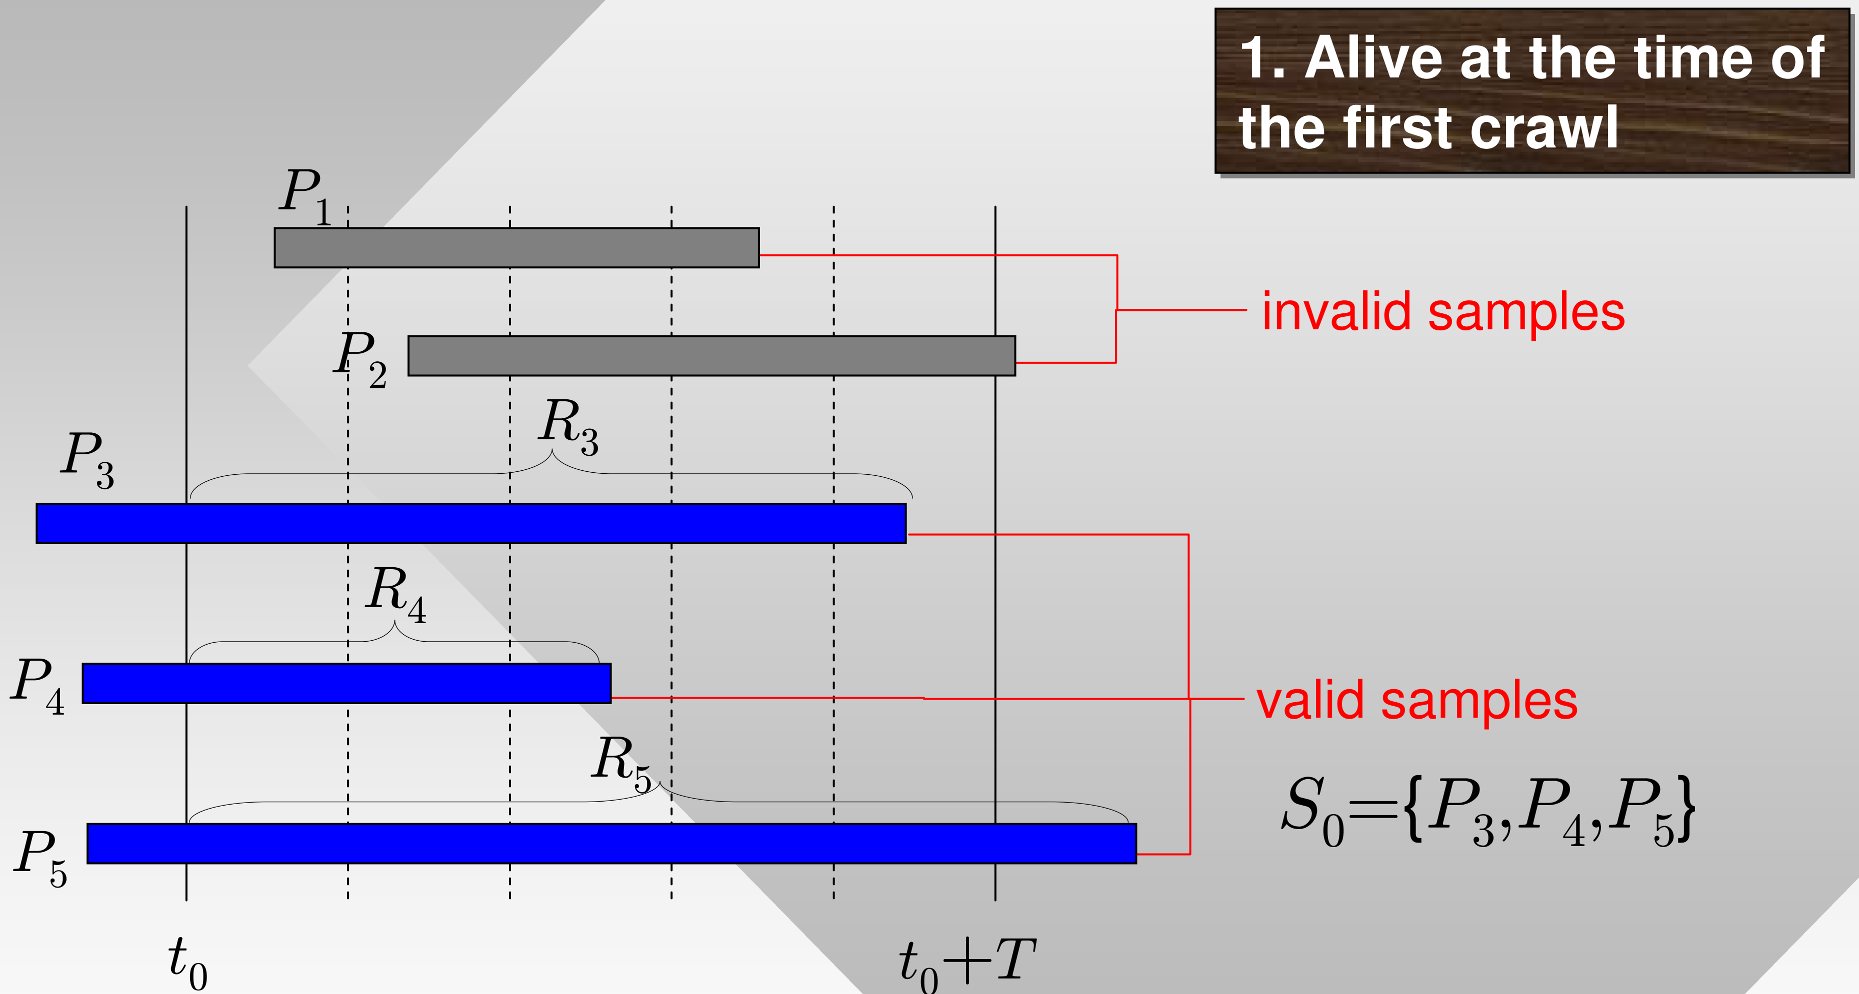

- RIDE samples the residual lifetime R of users in S_0
  - From time t_0 until the user dies

RIDE – Subsampling

- RIDE acquires all valid samples in the initial set S_0 during the very first crawl
  - This allows us to randomly subsample the users in set S_0
- Suppose we track the residuals of only ϵ percent of the entire initial set S_0
  - Significantly reduce traffic requirements
- Note that subsampling is not possible in CBM
  - It requires full system crawls to discover new users

RIDE – Estimator

- RIDE has all valid samples starting from t_0
  - It will never miss any sample nor have any round-off errors
- Theorem 2: The following equation defines an unbiased estimator

$$E_R(x_j) = 1 - \frac{h(x_j)}{h(0)}$$

estimated probability $P(L \le x_j)$

PDF of residuals lifetimes

RIDE – Simulations without Subsampling

- Lifetimes are Pareto with α=1.1, β=0.05
  - $E[L]$=0.5 hours, T=24 hours, Δ=15 min, $|S_0|$=1M

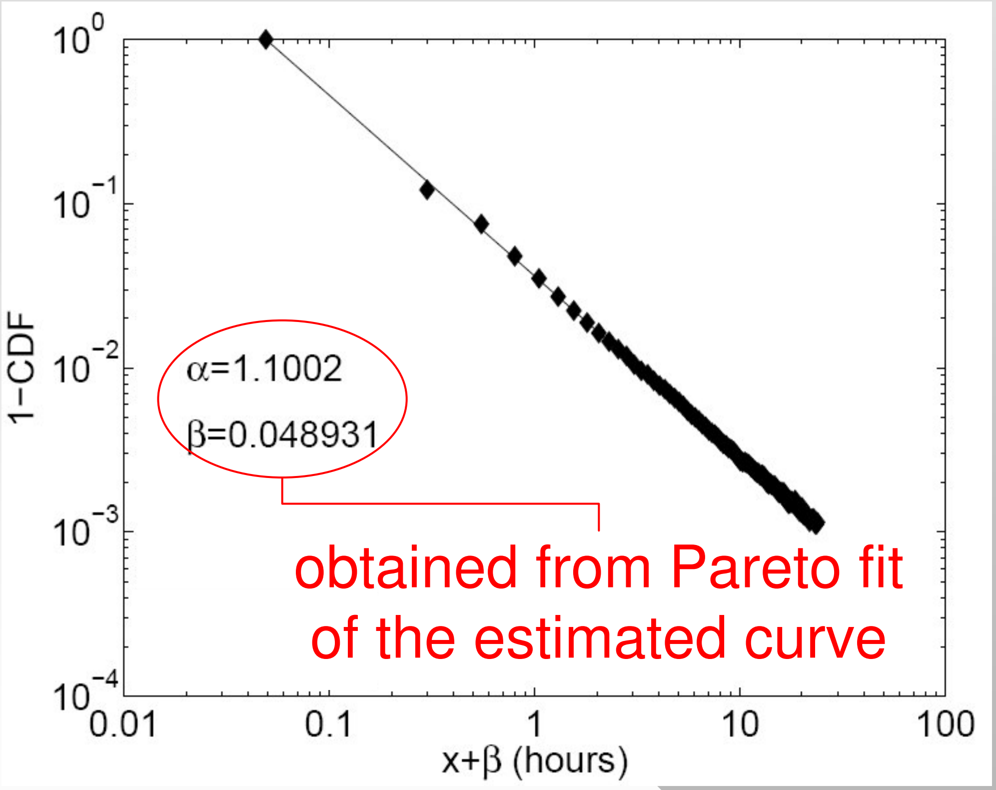

ϵ=1

RIDE – Simulations with Subsampling

- Applying inverse average

before inverse average

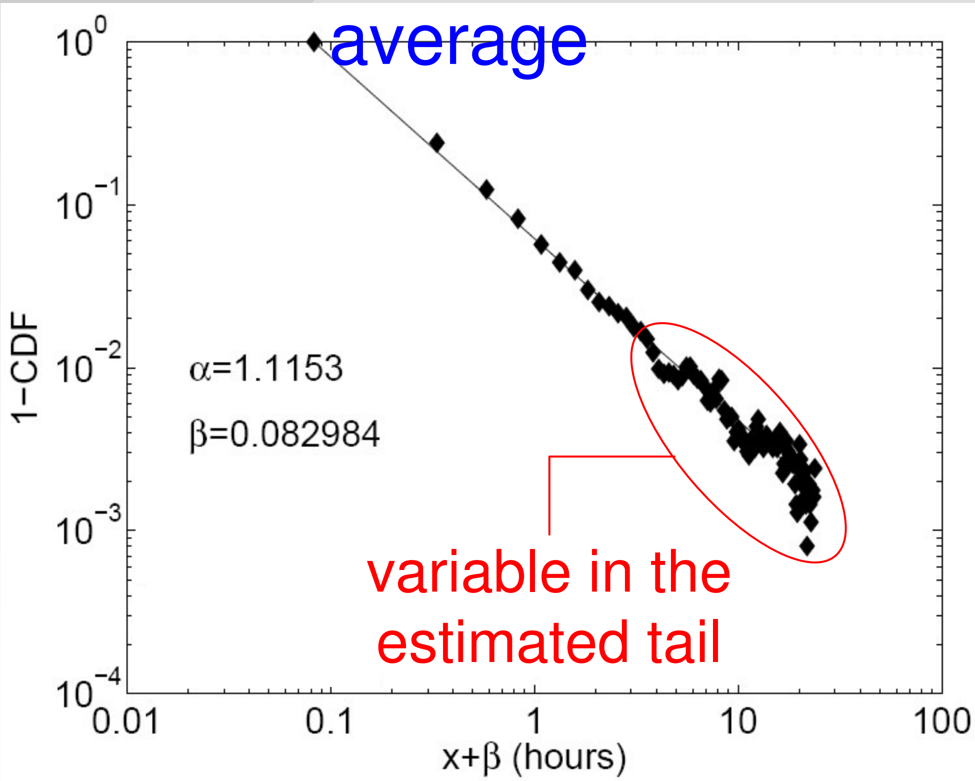

$\epsilon=0.01$

after inverse average

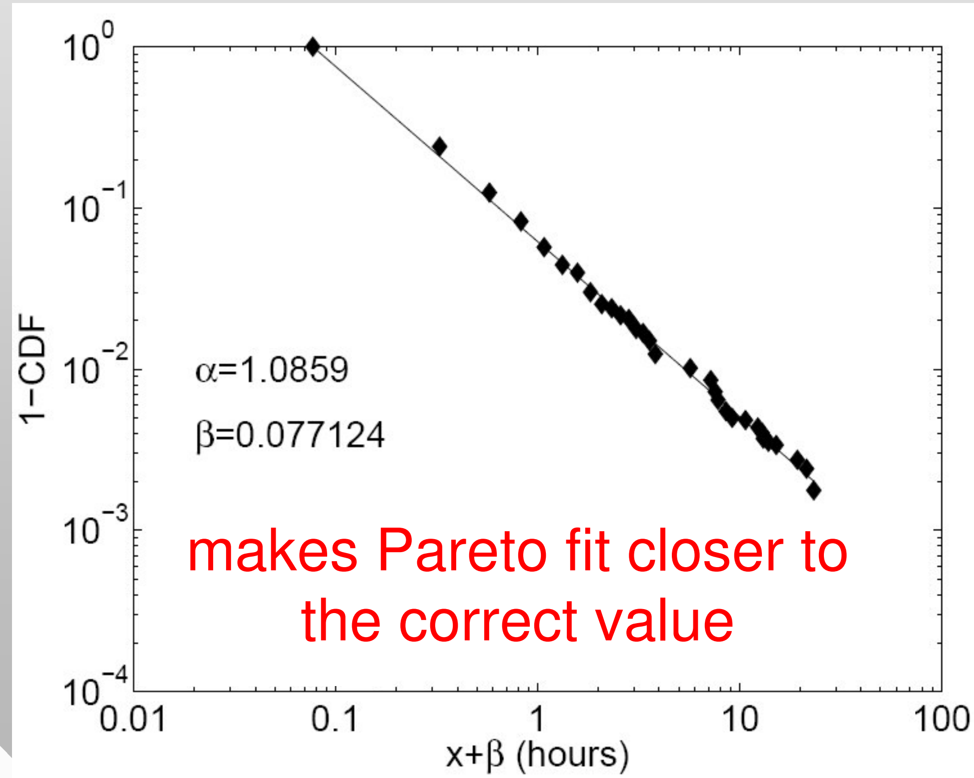

$\epsilon=0.01$

RIDE – Simulations with Subsampling

- Applying inverse average

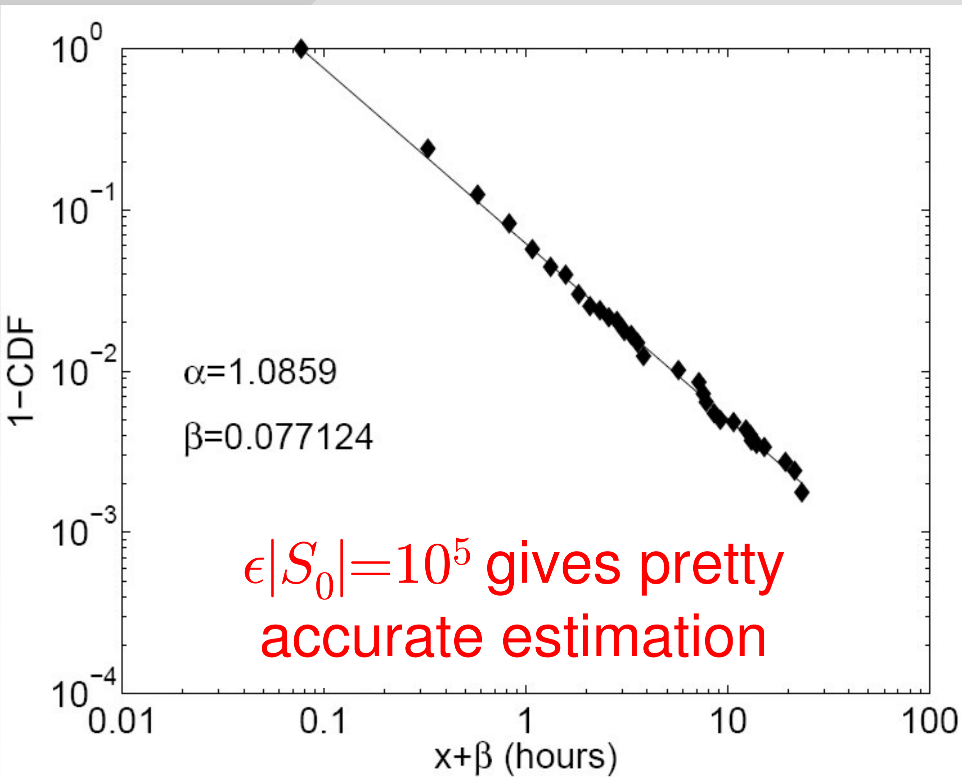

$\epsilon=0.01$

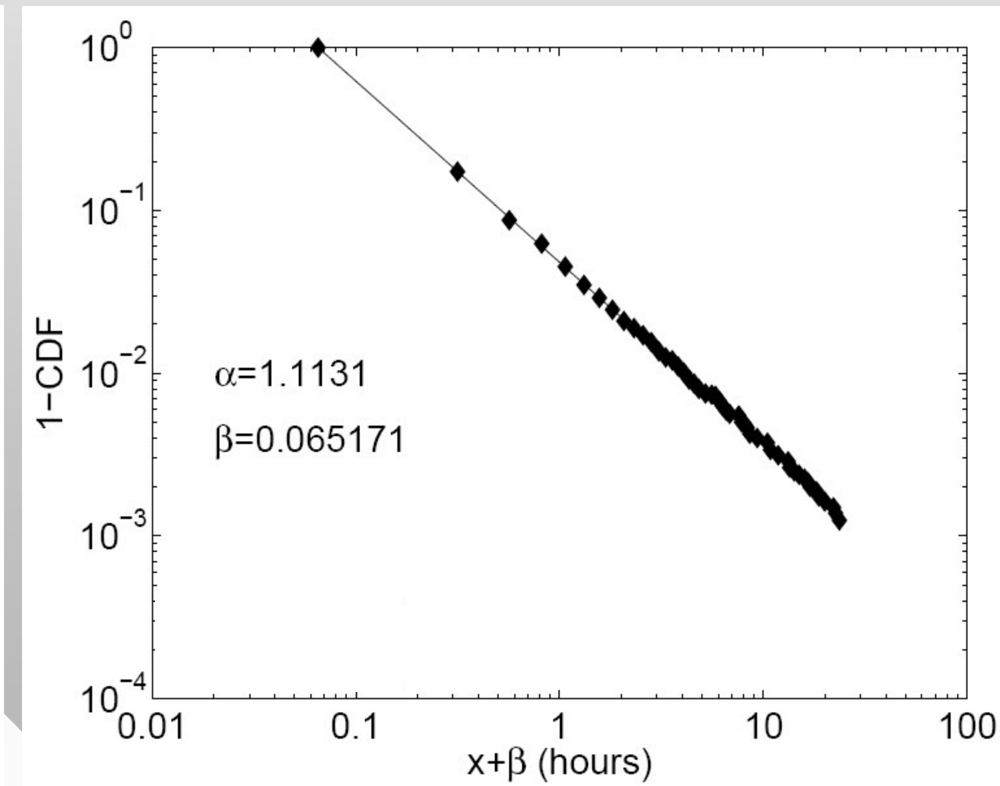

$\epsilon=0.1$

Agenda

- Introduction
- Related work
  - Create-Based Method (CBM)
- Analysis of CBM
- Proposed method
  - Residual-Based Estimator (RIDE)
- Comparison of overhead
- Experiments
- Conclusion

Overhead Comparison

- Theorem 3: Total bandwidth overhead of (Δ, T)-sampling using CBM and RIDE is given by:

$$b_{CBM} = \frac{Cn}{\Delta}\left(T + \int_0^T [H(T) - H(x)]dx\right)$$

$$b_{RIDE} = \frac{C|S_0|}{\Delta}\left(\Delta + \epsilon\int_0^T [1 - H(x)]dx\right)$$

- Define $q(\epsilon)$ the ratio of overhead of CBM and RIDE

$$q(\epsilon) = \frac{b_{CBM}}{b_{RIDE}}$$

Overhead Comparison (cont.)

- Pareto lifetime with $E[L]=1$ hours, $\Delta=3$ min

α	T	$q(0.1)$	$q(0.01)$
1.1	24 hrs	16	125
	48 hrs	17	151
	72 hrs	18	164

α	T	$q(0.1)$	$q(0.01)$
2	24 hrs	71	319
	48 hrs	116	573
	72 hrs	157	811

smaller ϵ, more savings

larger T, more savings

RIDE saves overhead by a fact of more than 800

- In fact, we can choose proper ϵ based on the size of the initial set S_0
  - $\epsilon|S_0|$ is fixed at some pre-determined value, e.g., 10^5

Agenda

- Introduction
- Related work
  - Create-Based Method (CBM)
- Analysis of CBM
- Proposed method
  - Residual-Based Estimator (RIDE)
- Comparison of overhead
- Experiments
- Conclusion

Experiments – Gnutella

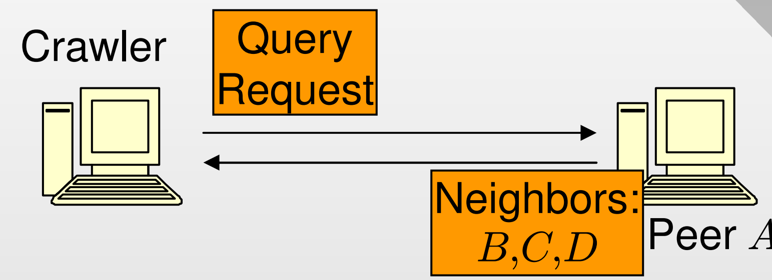

- Gnutella is fully distributed
  - However, it allows to query the neighbor list of any node
- Using Breadth-First-Search (BFS), we can take the snapshots of the users in the system
  - Peers A, B, C, D

peer snapshot

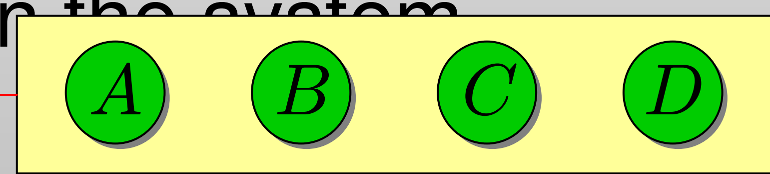

- Moreover, we can infer link information from query replies
  - Links (A,B), (A,C), (A,D)

link snapshot

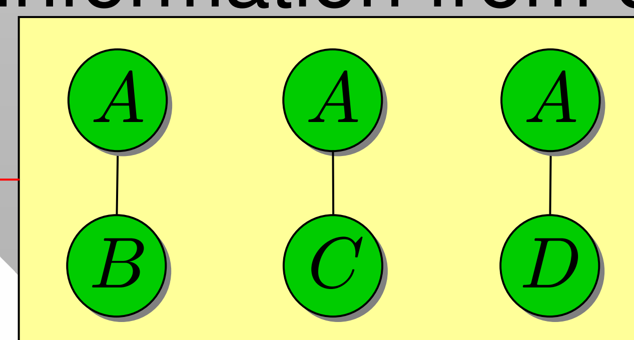

Experiments – Gnutella Crawler

- We implemented a Gnutella crawler – GnuSpider
  - BFS search among ultra-peers
  - Up to $60{,}000$ simultaneous connections
  - $216{,}000$ contacted ultra-peers per min
- Entire system crawled in 3 min on July 22, 2006
  - 6.4 million users (1.2 million ultra and 5.2 million leaves)
- First crawl obtains $468{,}000$ responsive ultra-peers
  - Subsampling $\epsilon|S_0| = 100{,}000$
  - $\Delta = 3$ min, $T = 72$ hours

Experiments – Lifetime Distributions

Lifetimes are Pareto with very heavy-tailed $\alpha \approx 1.09$ and 1.8 for ultrapeers and links

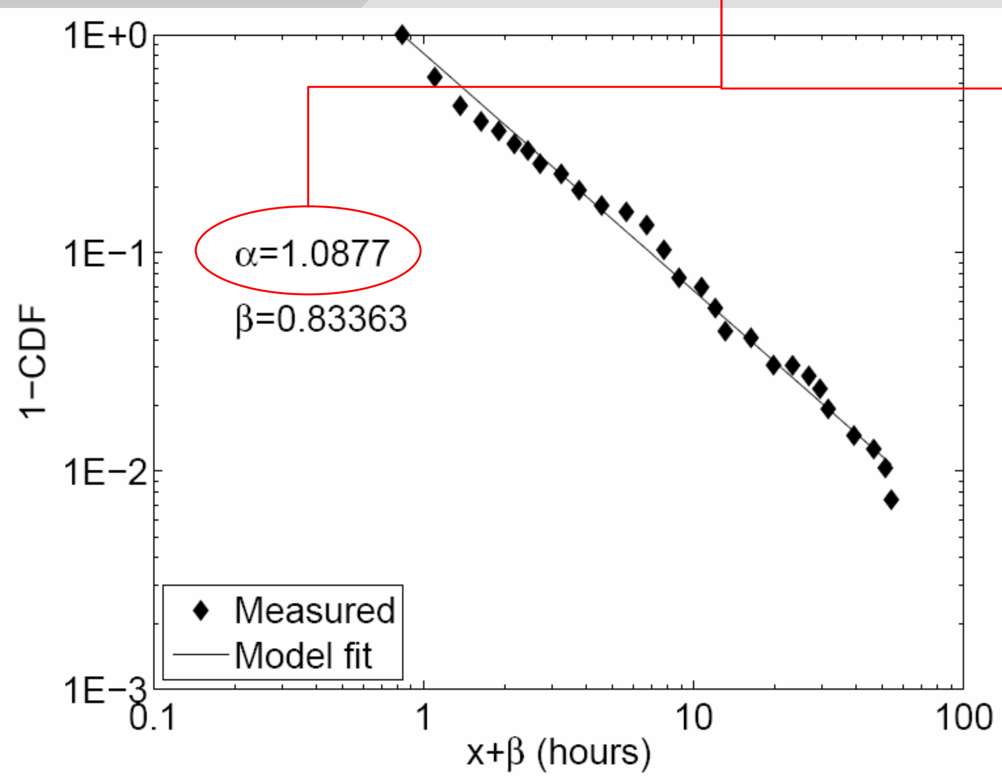

responsive ultrapeer

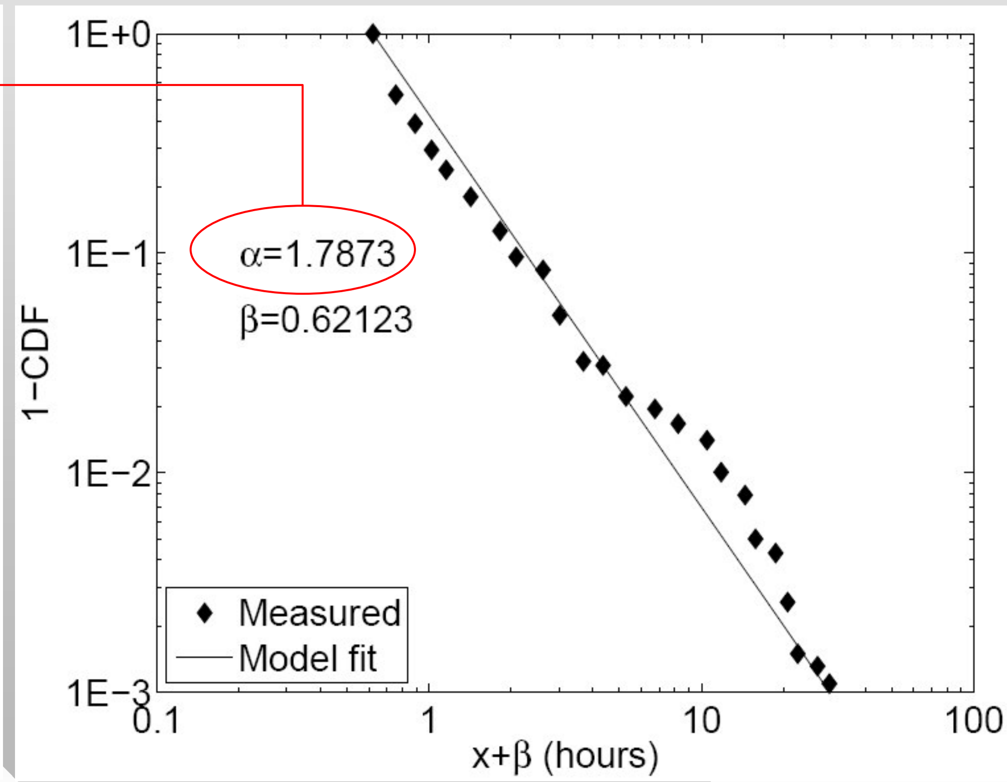

associated links

Conclusion

- CBM is generally biased for $\Delta > 0$
  - It may not scale to large networks
- RIDE can reduce traffic overhead by several orders of magnitude
  - Generally more accurate and scalable than CBM
- Ultrapeer lifetimes are Pareto with $\alpha \approx 1.09$
  - $\alpha \approx 1.06$ (Bustamante 2003)
- Link lifetimes exhibit much lighter tails with $\alpha \approx 1.8$
  - More volatile than ultrapeers