Estimation of DNS Source and Cache Dynamics under Interval-Censored Age Sampling

Di Xiao, Xiaoyong Li, Daren B.H. Cline, Dmitri Loguinov

Internet Research Lab

Department of Computer Science and Engineering

Texas A&M University

April 18, 2018

Agenda

- Introduction
- Notation
- Models
- Passive Measurement (Chameleon)
- Active Measurement (Shark)
- Experiments

Introduction

- Caches are important parts of many distributed systems in today's Internet
  - Search engines
  - Wireless mobile networks
  - P2P structures
  - CDNs
- Two common eviction policies
  - Capacity-based: unpopular elements get removed
  - TTL-based: elements are evicted only upon expiration
- TTL-based policy is more suitable when object staleness is of primary concern

Introduction

- DNS has long used TTL-based policy
  - However, staleness of records was never considered
- Instead, hit rate h was the sole metric of performance for many years
  - With wide adoption of dynamic DNS and CDNs, many authoritative domains frequently change IP addresses
  - Simply maximizing h, essentially setting TTL to infinity, is not a meaning pursuit
- The tradeoff between h and freshness f may be of interest in certain applications
  - Unfortunately, this interplay has not been studied

Introduction

- Besides modeling f, it is important to develop methods for estimating this value
  - Applications may aim to minimize staleness for a given hit rate, or maximize h for a given f
  - Researchers and network admins might be interested in measuring performance of existing systems to which they have no direct access
- Approaches to sampling f
  - Passive: observer is internal to the cache, but external to the source
  - Active: observer is external to both

Agenda

- Introduction
- Notation
- Models
- Passive Measurement (Chameleon)
- Active Measurement (Shark)
- Experiments

Notation

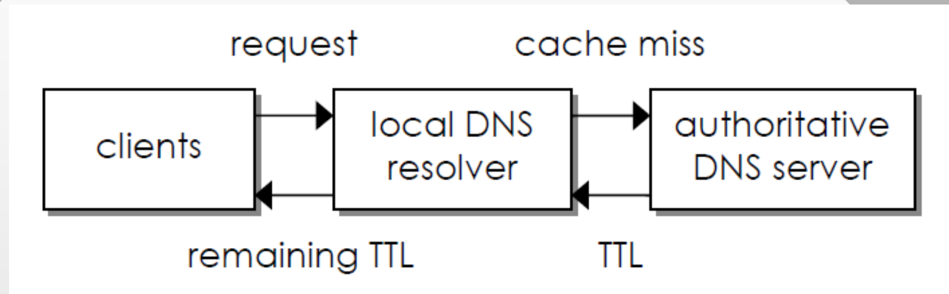

- Source experiences random updates via process N_U
- Clients query resolver via process N_Q
- TTLs are iid and follow distribution $F_T(x)$
- Interplay between N_Q and $F_T(x)$ defines download process N_D

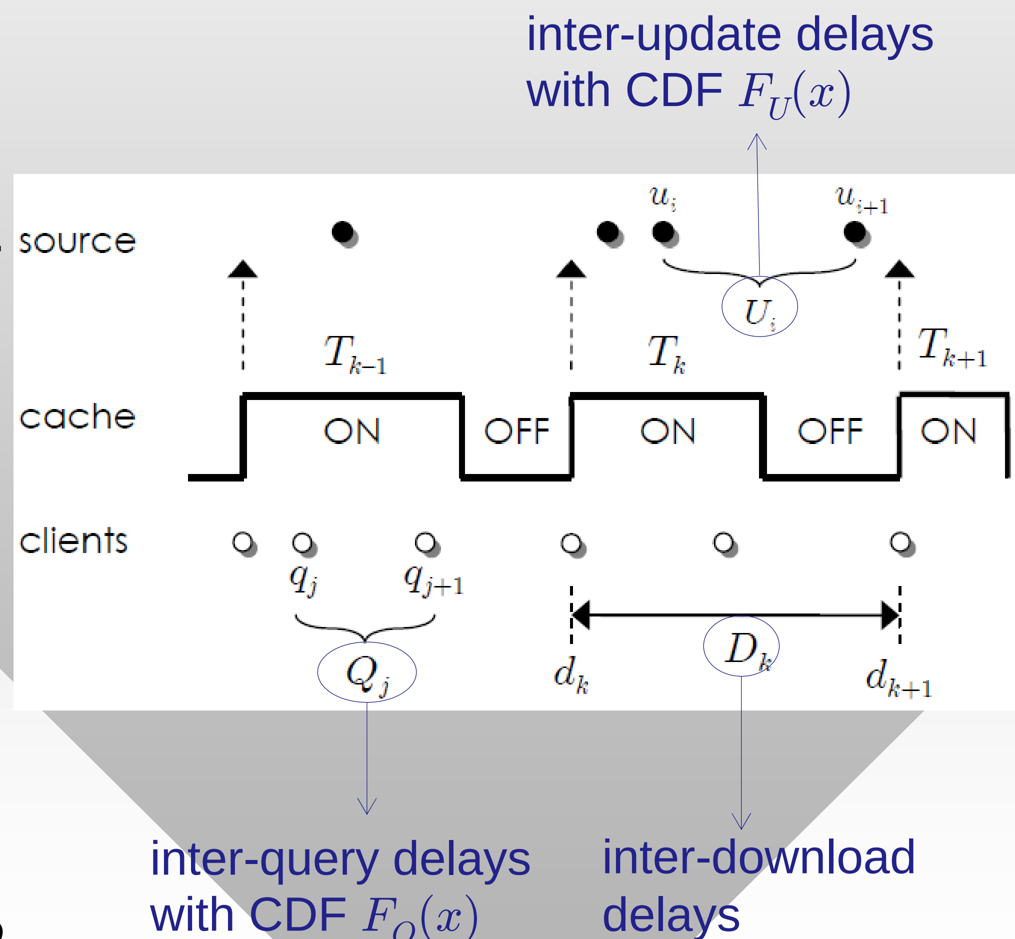

Agenda

- Introduction
- Notation
- Models
- Passive Measurement (Chameleon)
- Active Measurement (Shark)
- Experiments

Models

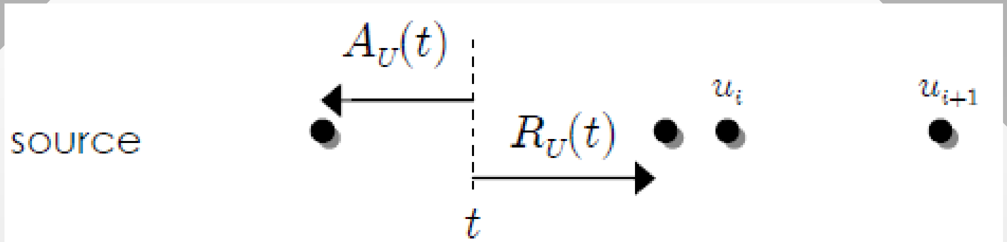

- Define update age $A_U(t)$ and update residual $R_U(t)$
  - For $t \to \infty$, they have the residual distribution of $F_U(x)$
  - We call this function $G_U(x)$
  - Neither $A_U(t)$ nor $R_U(t)$ is directly measurable!
- Assuming $T \sim F_T(x)$ is a random TTL, hit rate was already shown in previous work as

$$h = \frac{E[N_Q(T)]}{1+E[N_Q(T)]}$$

- Theorem: Assuming $R_U \sim G_U(x)$ is a random update residual, cache replies are fresh with probability

$$f = \frac{1+E[N_Q(\min(R_U,T))]}{1+E[N_Q(T)]}$$

Models

- Theorem: Increasing $E[T] \to \infty$ produces $h \to 1$ and $f \to 0$. Making Q stochastically smaller yields $h \to 1$ and $f \to E[\min(R_U, T)]/E[T]$ from above
- Tradeoff between h and f
  - Simulation with Pareto U, T, Q with $E[U] = 20$ sec

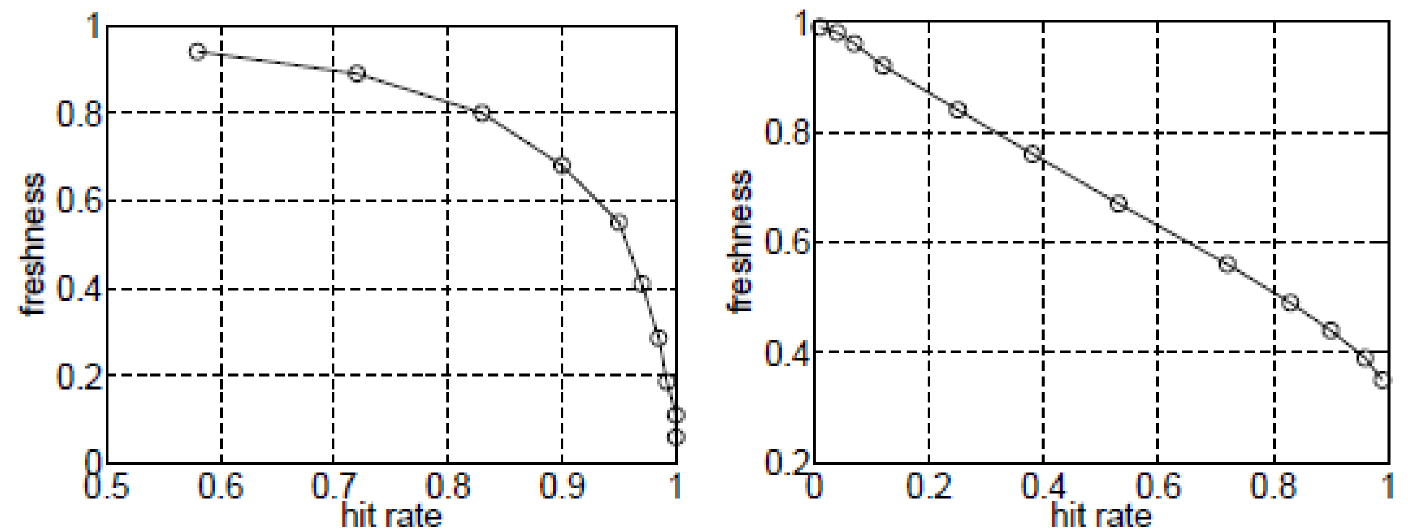

(a) changing $E[T]$ ($\lambda = 1$) (b) changing λ ($E[T] = 40$)

Models

- It is common to assume queries arrive from a large number of users and N_Q is Poisson, where

$$h = \frac{\lambda E[T]}{1+\lambda E[T]} \qquad f = \frac{1+\lambda E[\min(R_U,T)]}{1+\lambda E[T]}$$

- Let $G_T(x)$ be the residual distribution of $F_T(x)$ and $R_T \sim G_T(x)$ is a random residual TTL
- Theorem: Under Poisson queries, freshness is given by $f = 1 - h(1-p)$, where $p := P(R_T < R_U)$
- Note that f requires h and p, the latter of which is significantly more difficult to estimate as it needs the distribution of both R_U and R_T
  - Since $G_T(x)$ is relatively easy to obtain, the main focus of the paper is on sampling $G_U(x)$ and $F_U(x)$

Roadmap

- Passive measurement is the cache estimating f with access to all information except $G_U(x)$
  - E.g., an adaptive cache that trades spare bandwidth to maintain f at some desired threshold
- Active measurement is an observer estimating f only knowing residual TTLs $R_T(t)$
  - E.g., network administrators / researchers study replication efficiency and diagnose potential problems with bandwidth consumption or staleness
- Either way, it is useful to also recover $F_U(x)$ to characterize the source

Agenda

- Introduction
- Notation
- Models
- Passive Measurement (Chameleon)
- Active Measurement (Shark)
- Experiments

Passive Measurement

- The major challenge is to remotely estimate $G_U(x)$ assuming random inter-download delay
- Best previous method is M_6 from [Li 2015]
- Drawbacks of M_6
  - Slow convergence speed
  - Quadratic computation time
  - Non-concave estimator
- Our method Chameleon improves in all aspects

Passive Measurement

- It produces upper/lower bounds on $A_U(d_k)$
  - Based on detecting modification of records by comparing them at download points d_k and d_{k-1}
  - Generates one bound at each download point d_k
- Example: Suppose $d_{\gamma(k)}$ is the last point before d_k that detected a modification to the record

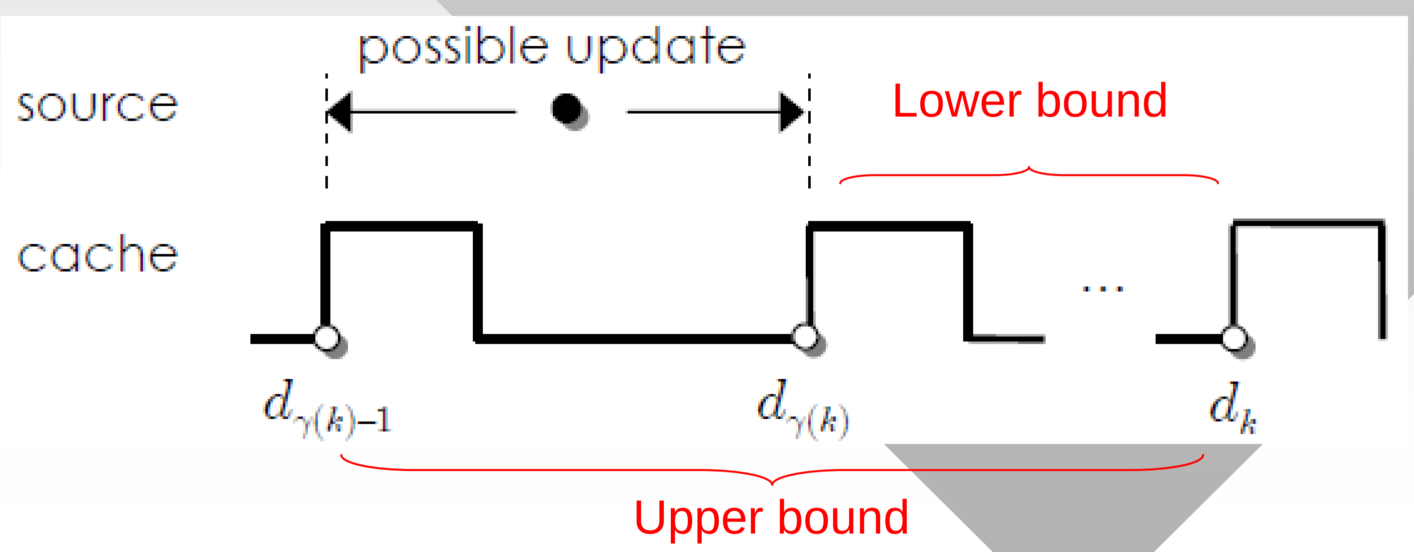

EM (Expectation Maximization) Estimator

- Problem maps to *interval-censored* observation
  - Naïve method inspired by [Turnbull 1976], which is asymptotically accurate as $n \to \infty$
  - However, complexity is $O(n^2)$ per iteration
- We propose an $O(n)$ time/space algorithm
- Further notice that $G_U(x)$ is a concave function
  - Which prior methods in this field cannot guarantee
- We thus propose a new concave EM estimator
  - Requires little additional CPU cost
  - Allows recovery of $F_U(x)$ from $G_U(x)$

Agenda

- Introduction
- Notation
- Models
- Passive Measurement (Chameleon)
- Active Measurement (Shark)
- Experiments

Active Measurement

- Without access to the cache, the observer needs to send probes to local resolver and remotely estimate all metrics
  - No previous work has studied this problem
  - Probes are iterative queries that don't pollute the cache
- Sampling points
  - A process N_S at points $\{s_k\}$ with iid sample delays S_k that follow distribution $F_S(x)$
- Two scenarios for TTL
  - Constant T: problem reduces to passive measurement
  - Random T: our focus next

Active Measurement

- Main challenge is we cannot see download points d_k, which makes passive techniques inapplicable
- We thus propose a new algorithm Shark
  - For each sample point s_k, we derive probabilistic bounds on d_k, which allows interval-censoring of update age $A_U(d_k)$
- We then modify our concave EM to operate with probabilistic intervals
  - This usually requires a huge number of samples to be processed; however, our EM is fast enough to handle billions of points on commodity machines

Agenda

- Introduction
- Notation
- Models
- Passive Measurement (Chameleon)
- Active Measurement (Shark)
- Experiments

Experiment Setup

- Passive scenario
  - An authoritative server A
  - A local DNS server that resolves at A
  - A background traffic generator from process N_Q

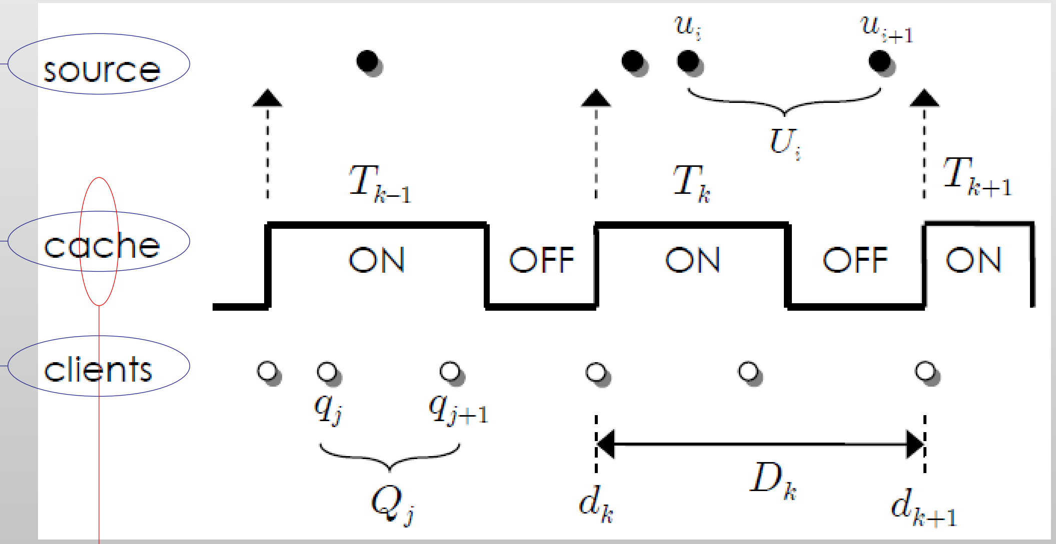

- Active scenario
  - We randomly selected a recursive server R from our port-53 UDP scan of the Internet, making sure it was stable and compliant with source-provided TTLs
  - Added an observer that sends probes to R

Experiments

- Performance of our fast EM (Algorithm 1)
  - Scales no worse than linearly in n
  - Handles 1 billion observations in 30 sec
- Estimation of simple parameters
  - Both Chameleon and Shark are asymptotically accurate

TABLE I
RUNTIME IN PASSIVE ESTIMATION OF $G_U(x)$

n	M_6	Naive EM $\epsilon = 10^{-4}$	Algorithm 1 $\epsilon = 10^{-4}$
10^4	0.3 sec	9.4 sec	0.06 sec
10^5	58 sec	2.9 min	0.13 sec
10^6	2.2 hours	43 min	0.19 sec
10^7	–	5.5 hours	0.70 sec
10^8	–	–	5.79 sec
10^9	–	–	27.9 sec

TABLE II
RELATIVE ESTIMATION ERROR OF SIMPLE PARAMETERS

n	Shark			Chameleon		
	h	$E[T]$	λ	h	$E[T]$	λ
10^2	2.94%	16.8%	51.6%	1.58%	4.18%	7.93%
10^3	0.88%	5.81%	11.5%	0.42%	1.58%	2.79%
10^4	0.29%	2.40%	3.73%	0.16%	0.41%	0.80%

Experiments

- Passive recovery of $G_U(x)$ with 10K samples
  - Pareto $F_U(x)$ with $\alpha = 3$ and $E[U] = 20$ sec
  - M_6 is poorly suited for such small sample size

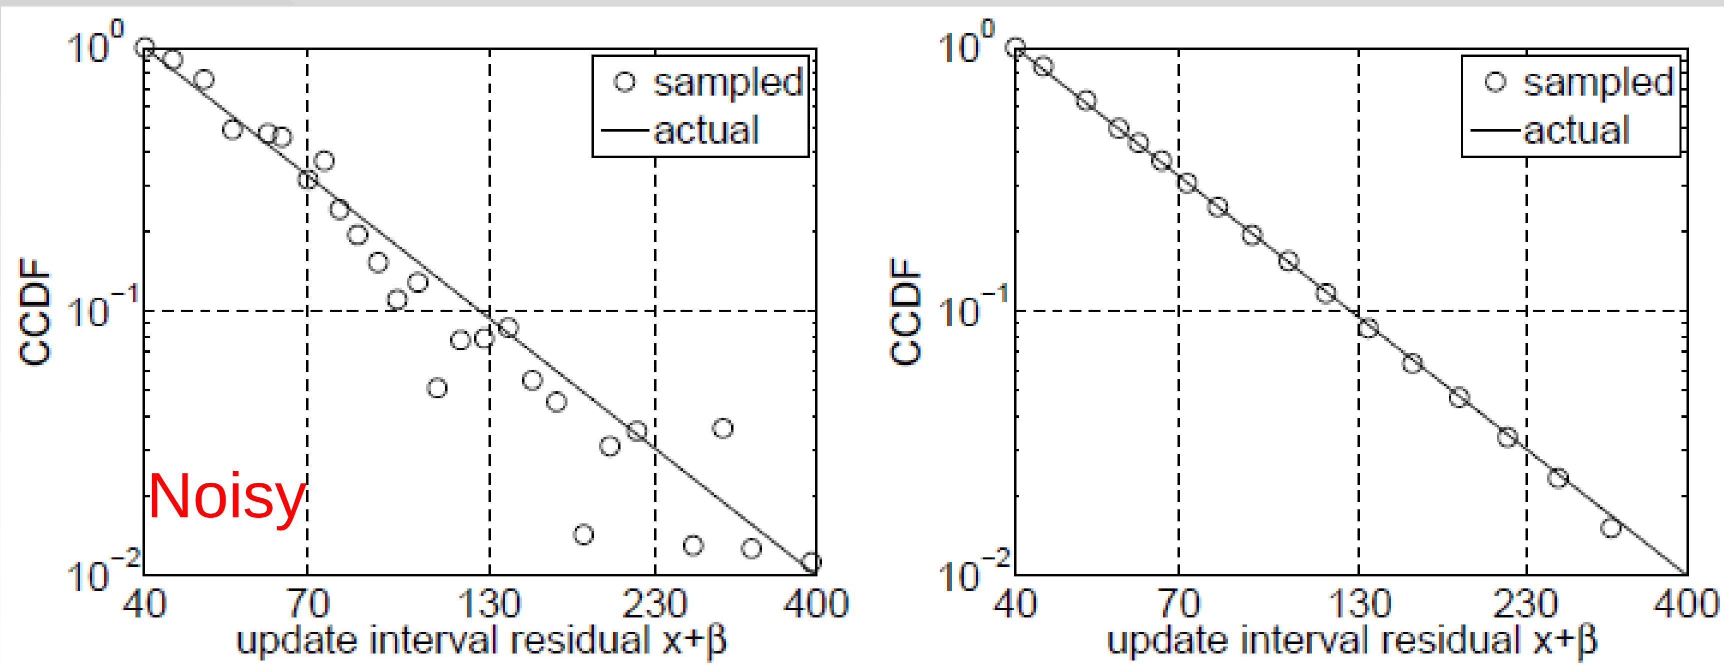

M_6 Chameleon

Experiments

- Passive recovery of $F_U(x)$ with 10K samples
  - Non-concave EM produce a bunch of random points that cannot be a CDF function

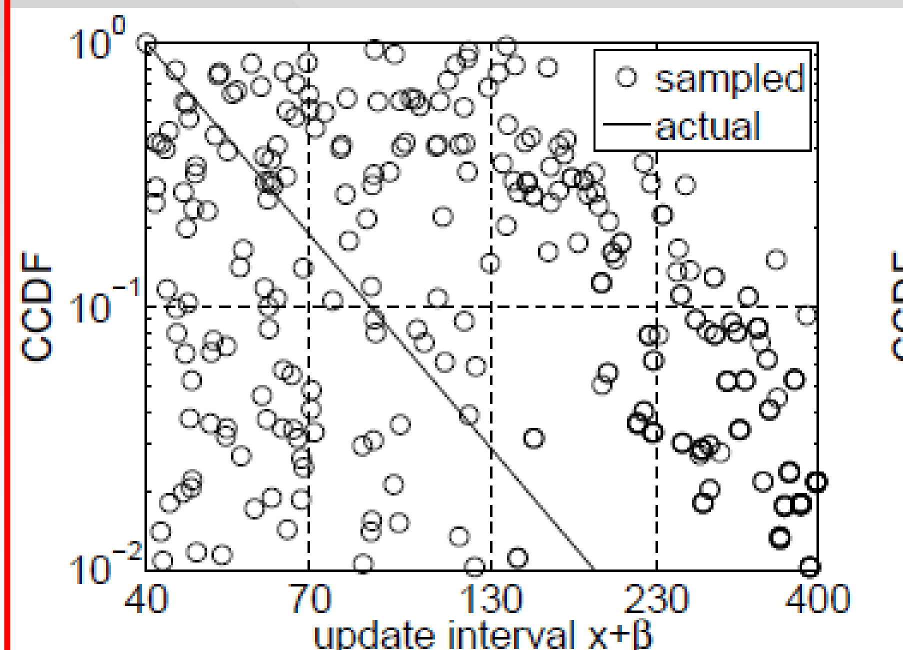

(a) non-concave EM ($\epsilon = 10^{-5}$)

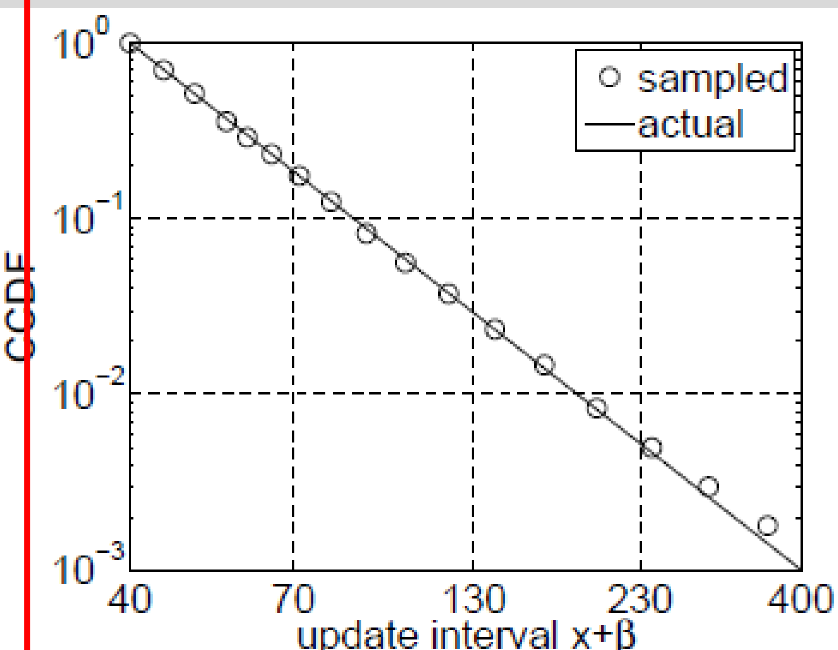

(b) concave EM ($\epsilon = 10^{-5}$)

Experiments

- Active recovery of $G_U(x)$ with 10K samples
  - Pareto and uniform $F_U(x)$ with $E[U] = 20$ sec

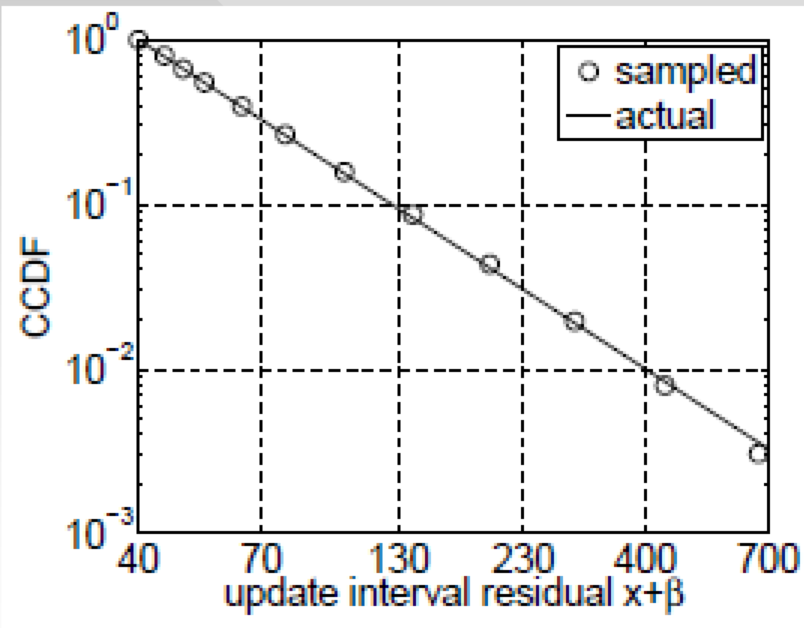

(a) Pareto U

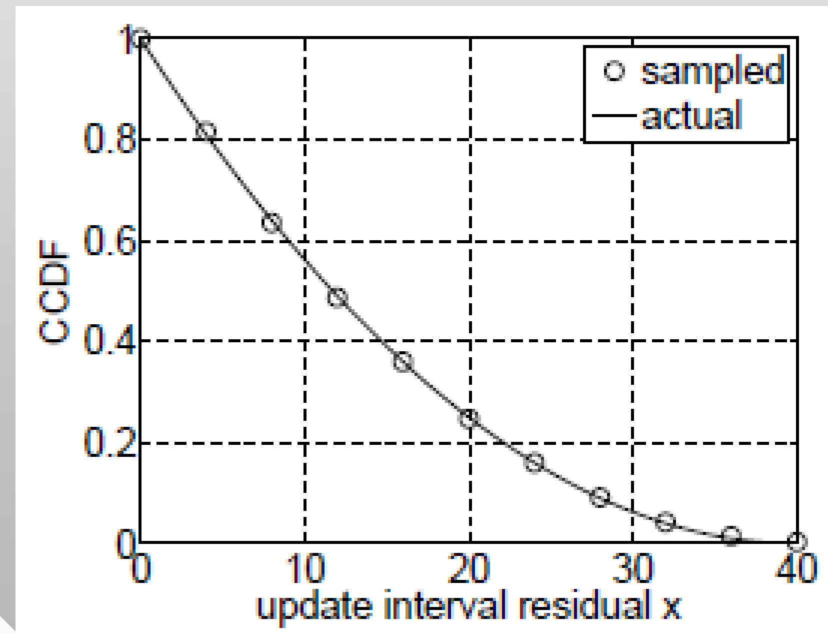

(c) uniform U

Experiments

- Estimation of complex parameters
  - Although Chameleon operates with more information, Shark comes pretty close to matching its accuracy

n	Pareto U			
	Shark		Chameleon	
	p	f	p	f
10^2	6.5%	5.8%	6.3%	5.6%
10^3	1.9%	1.8%	1.9%	1.7%
10^4	0.5%	0.4%	0.7%	0.6%

Uniform U			
Shark		Chameleon	
p	f	p	f
5.0%	4.8%	4.0%	3.4%
1.7%	1.5%	1.2%	1.1%
0.5%	0.4%	0.5%	0.5%

Thank you!

Questions?