Modeling Residual-Geometric Flow Sampling

Xiaoming Wang

Joint work with Xiaoyong Li and Dmitri Loguinov

Amazon.com Inc., Seattle, WA

April 13th, 2011

Agenda

- Introduction
- Underlying model of residual sampling
- Analysis of existing estimators
- Proposal of new estimators
- Performance evaluation
- Conclusion

Introduction

- Traffic monitoring is an important topic for today's Internet
  - Security, accounting, traffic engineering
- It has become challenging as Internet grew in scale and complexity
- In this talk, we focus on two problems in the general area of measuring flow sizes
  - Determining the number of packets of elephant flows
  - Recovering the distribution of flow sizes

Related Work

- Packet sampling
  - Sampled NetFlow (Cisco)
  - Adaptive NetFlow (Estan, SIGCOMM'04)
  - Sketch-guided sampling (Kumar, INFOCOM'06)
  - Adaptive non-linear sampling (Hu, INFOCOM'08)
- Flow sampling
  - Sample-and-hold (Estan, SIGCOMM'02)
  - Flow thinning (Hohn, IMC'03)
  - Smart sampling (Duffield, IMC'03/SIGMETRICS'03)
  - Flow slicing (Kompella, IMC'05)

Analysis of Underlying Model

- Our talk is based on the sampling method proposed by sample-and-hold (Estan, SIGCOMM'02)
- We call this method by Residual-Geometric Sampling (RGS) due to two reasons:
  - This belongs to the class of residual-sampling techniques (Wang, INFOCOM'07/P2P'09)
  - It can be modeled by a geometric process
- Our analysis of RGS covers two goals:
  - Providing a unifying analytical model
  - Understanding the properties of samples it collects

Analysis of Underlying Model 2

- How does RGS work?
  - For a sequence of packets traversing a router, it checks each packet's flow id x in some RAM table
  - If x is found, its counter is incremented by 1
  - Otherwise, an entry is created for x with probability p and this packet is discarded with probability $1 - p$
- The state of a flow can be modeled by a simple geometric process

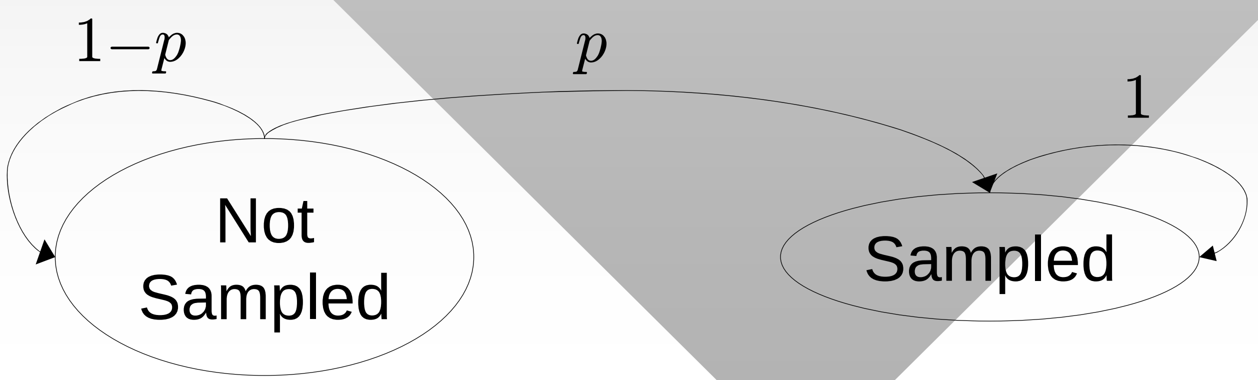

Analysis of Underlying Model 3

- We need several definitions:
  - Assume that flow sizes are i.i.d
  - Given a random flow with size L, define geometric age A_L the number of packets discarded from the front
  - Define geometric residual R_L the final counter value
- A flow of size 9 is not sampled until the 4th packet

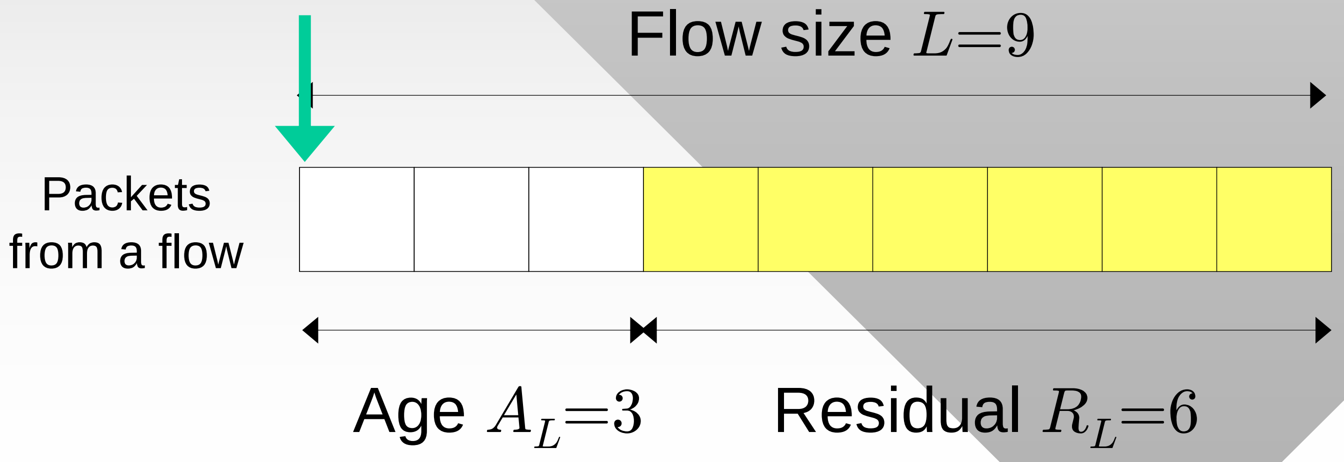

Analysis of Underlying Model 4

- Assume flow size L has a PMF f_i:

$$f_i = P(L = i)$$

- *Lemma 1*: Probability p_s of a flow being selected by RGS is:

$$p_s = 1 - \sum_{i=1}^{\infty} f_i (1-p)^i$$

- *Lemma 2*: PMF h_i of geometric residual R_L can be expressed as:

$$h_i = \frac{p \sum_{j=i}^{\infty} f_j (1-p)^{j-i}}{p_s}$$

Agenda

- Introduction
- Underlying model of residual sampling
- Analysis of existing estimators
- Proposal of new estimators
- Performance evaluation
- Conclusion

Previous Method – Single-Flow Usage

- Prior work on RGS (Estan, SIGCOMM'02) suggested following estimator of single-flow size:

$$e(R_l) = R_l - 1 + 1/p$$

- *Theorem 1*: For given size l, the expected value of estimator $e(R_l)$ is:

$$E[e(R_l)] = \frac{l}{1-(1-p)^l}$$

- It tends to overestimate the original flow size by a factor of up to $1/p$

Simulations – Estimated Size

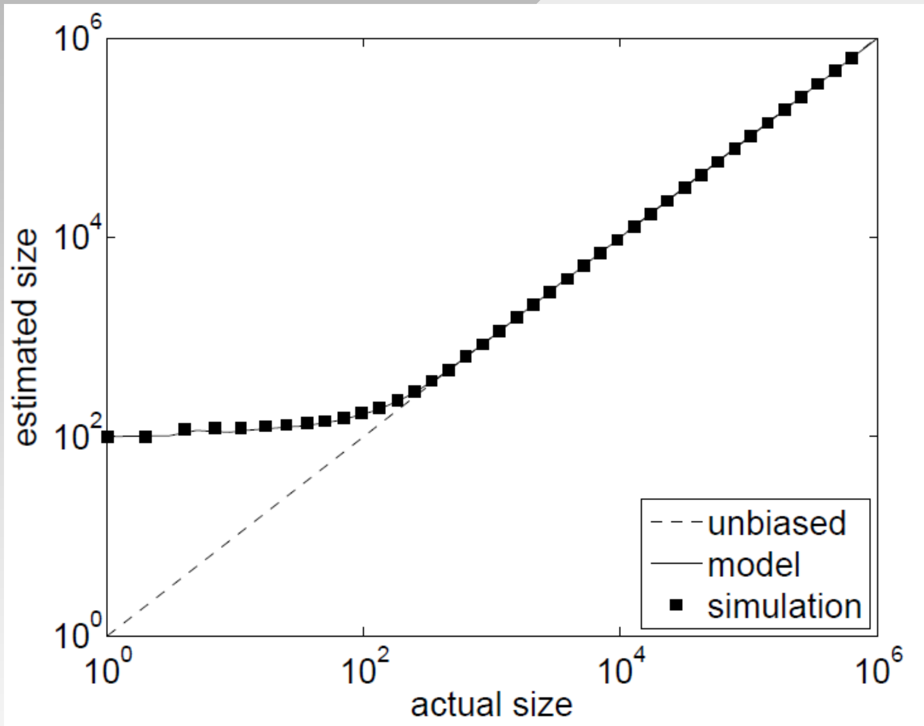

p=0.01

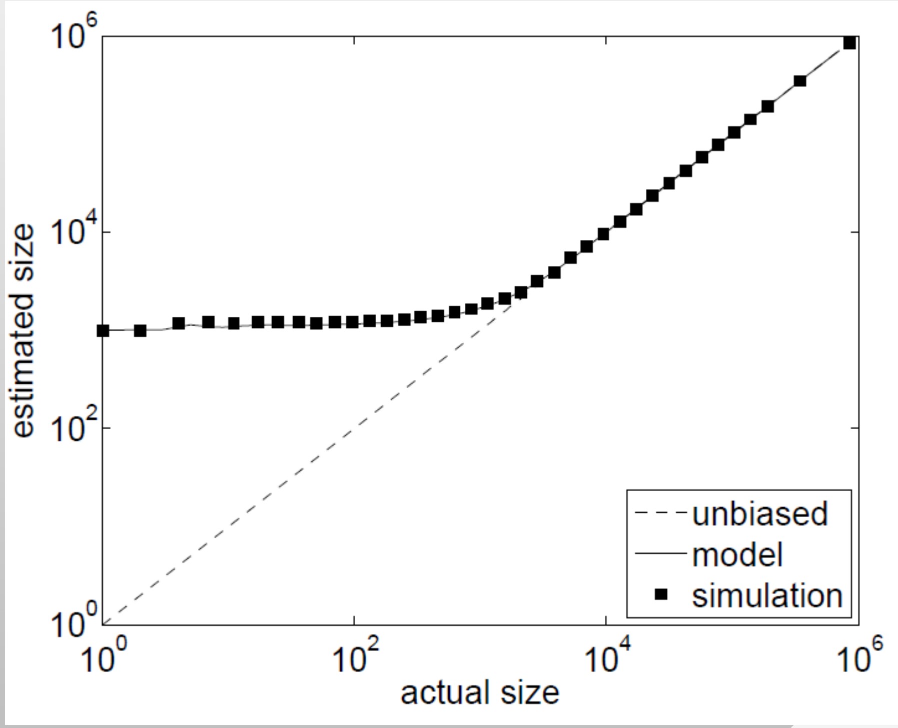

p=0.001

Previous Method – Single-Flow Usage 2

- Quantifying the error of individual values $e(R_l)$ in estimating flow size l
  - Relative Root Mean Square Error (RRMSE)

$$\delta_l = \sqrt{E[(Y_l - 1)^2]}$$

where $Y_l = e(R_l)/l$ is relative error

- *Theorem 2*: RRMSE of the existing RGS estimator:

$$\delta_l = \sqrt{\frac{1-p-l(l-1)p^2(1-p)^l-(1-p)^{l+1}}{l^2p^2(1-(1-p)^l)}}$$

Simulations – Relative RMSE

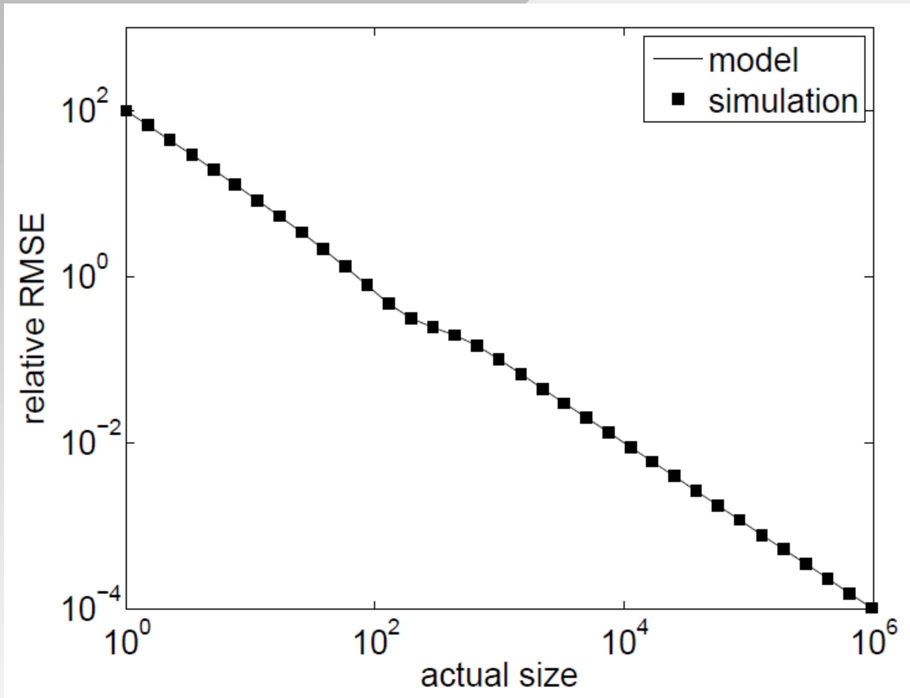

$p=0.01$

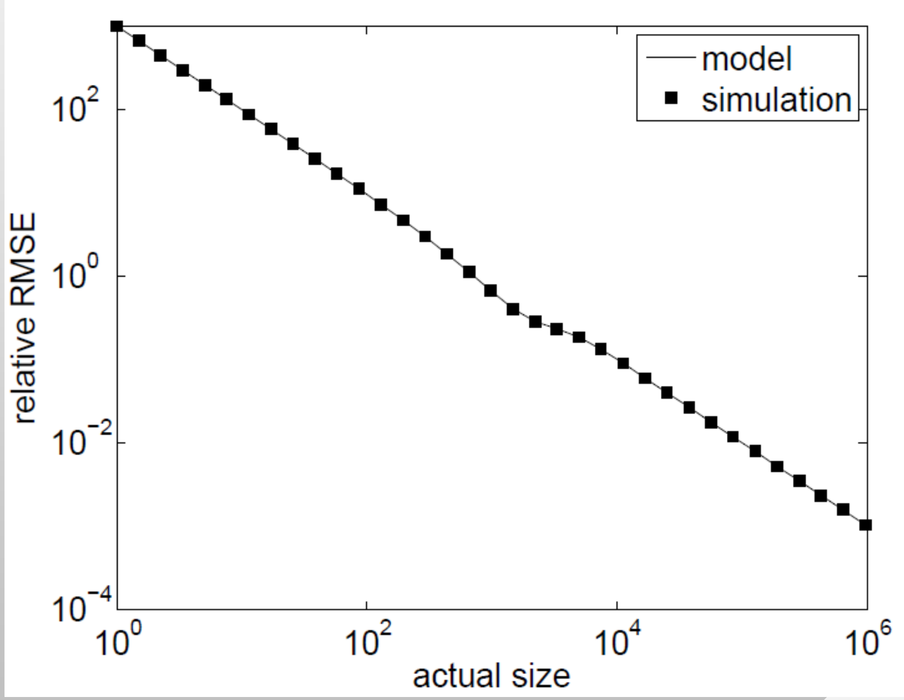

$p=0.001$

Previous Method – Flow-Size Distribution

- Consider PMF q_i of $e(R_L)$ and compare it with f_i

$$q_i = P(e(R_L) = i)$$

- *Theorem 3*: PMF of flow sizes estimated from $e(R_L)$ is:

$$q_i = \frac{\sum_{j=y(i)}^{\infty} f_j (1-p)^{j-y(i)} p}{p_s}$$

  - where p_s is the probability of a flow is selected and

$$y(i) = \lceil i + 1 - 1/p \rceil$$

- The estimated distribution q_i is quite different from actual f_i

Simulations – Flow Size Distribution

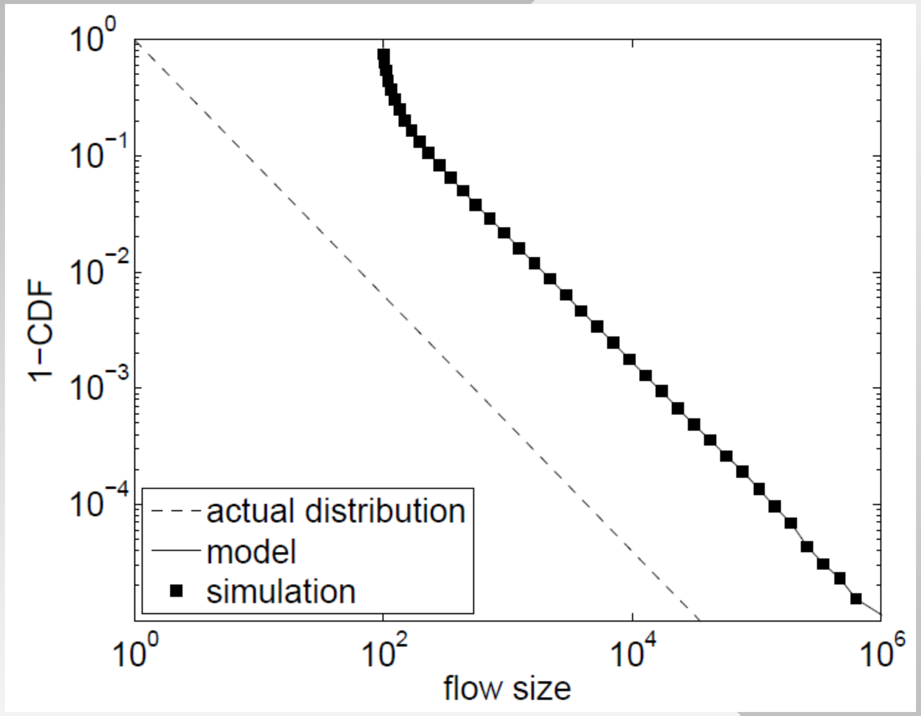

$p=0.01$

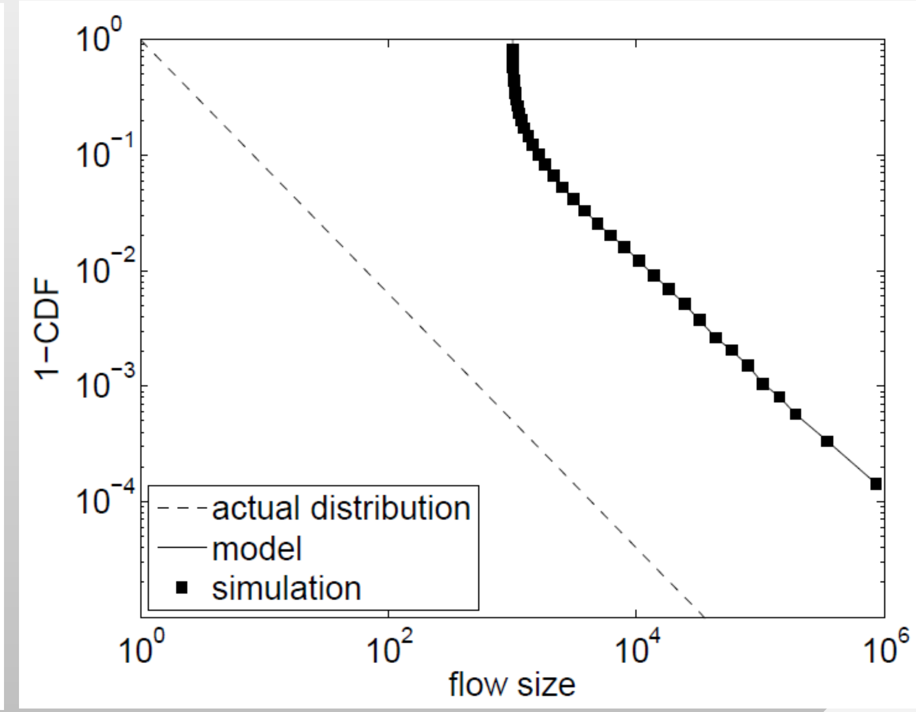

$p=0.001$

Agenda

- Introduction
- Underlying model of residual sampling
- Analysis of existing estimators
- Proposal of new estimators
- Performance evaluation
- Conclusion

URGE – Single-Flow Usage

- For single-flow size, we propose following estimator:

same different

$$\hat{e}(R_L) = R_L - 1 + 1/p - \frac{(1-p)^{R_L}}{p}$$

- *Lemma 3*: $\hat{e}(R_L)$ is unbiased for any flow size l
- *Theorem 4*: RRMSE of $\hat{e}(R_L)$ as:

$$\hat{\delta}_l = \sqrt{\frac{1-p+lp(p-2)(1-p)^l-(1-p)^{2l+1}}{l^2p^2(1-(1-p)^l)}}$$

URGE Simulations – Single-Flow Usage

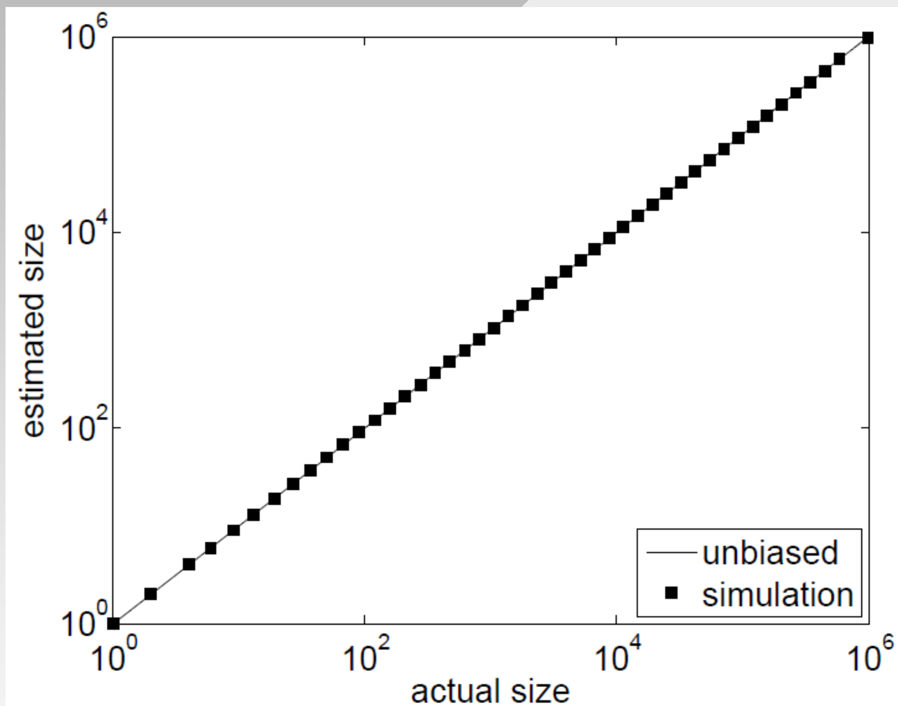

p=0.01

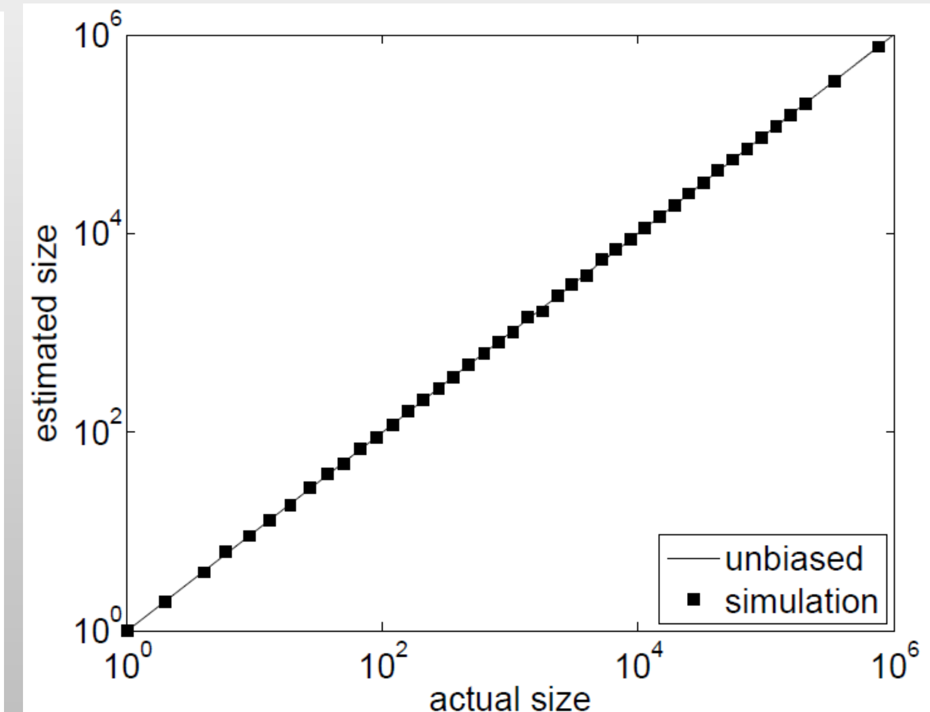

p=0.001

URGE Simulations – Relative Error

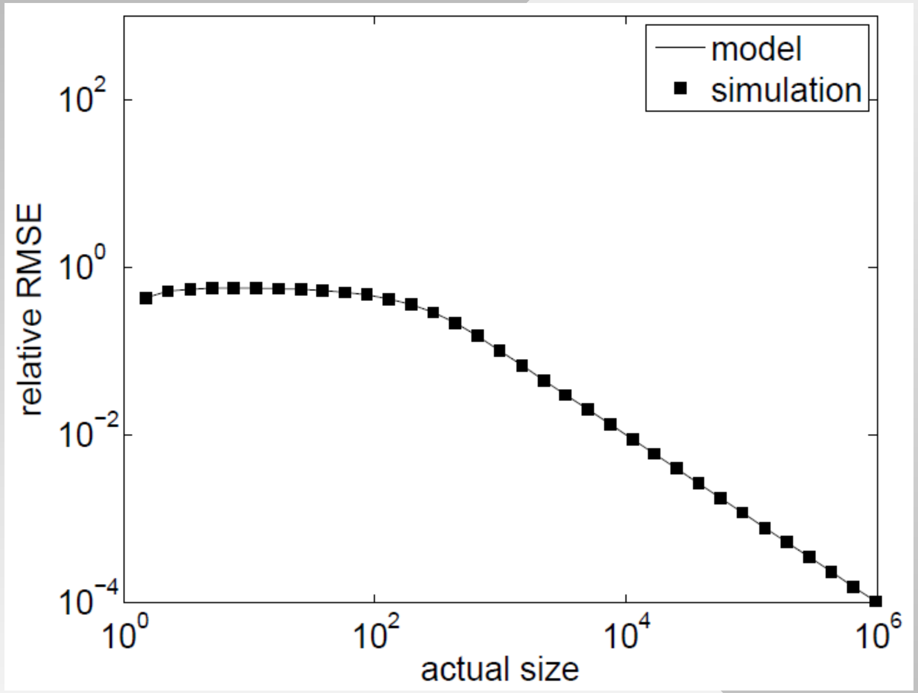

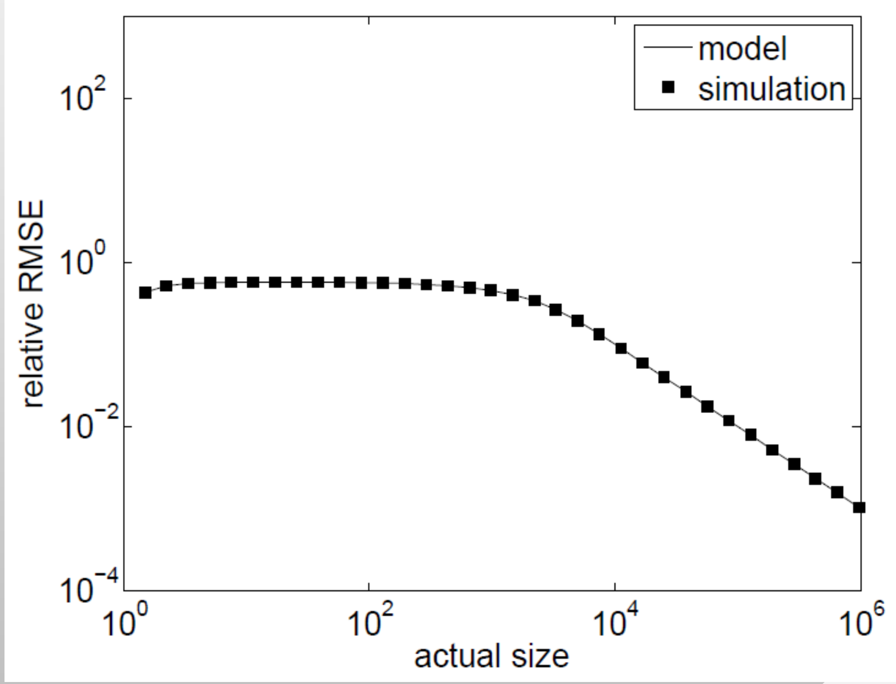

p=0.01

p=0.001

URGE – Flow Size Distribution

- *Lemma 5*: The flow size distribution f_i can be expressed using PMF of geometric residual h_i as:

$$f_i = \frac{h_i-(1-p)h_{i+1}}{p+(1-p)h_1}$$

- For flow size distribution, we propose following estimator:

$$\tilde{q}_i = \frac{M_i-(1-p)M_{i+1}}{Mp+(1-p)M_1}$$

- *Corollary 2*: Estimator $\{\tilde{q}_i\}$ is asymptotically unbiased, that is, $\{\tilde{q}_i\}$ converges in probability to f_i as the number M of sampled flows $\to \infty$

Computer Science, Texas A&M University

URGE Simulations – Flow Size Dist.

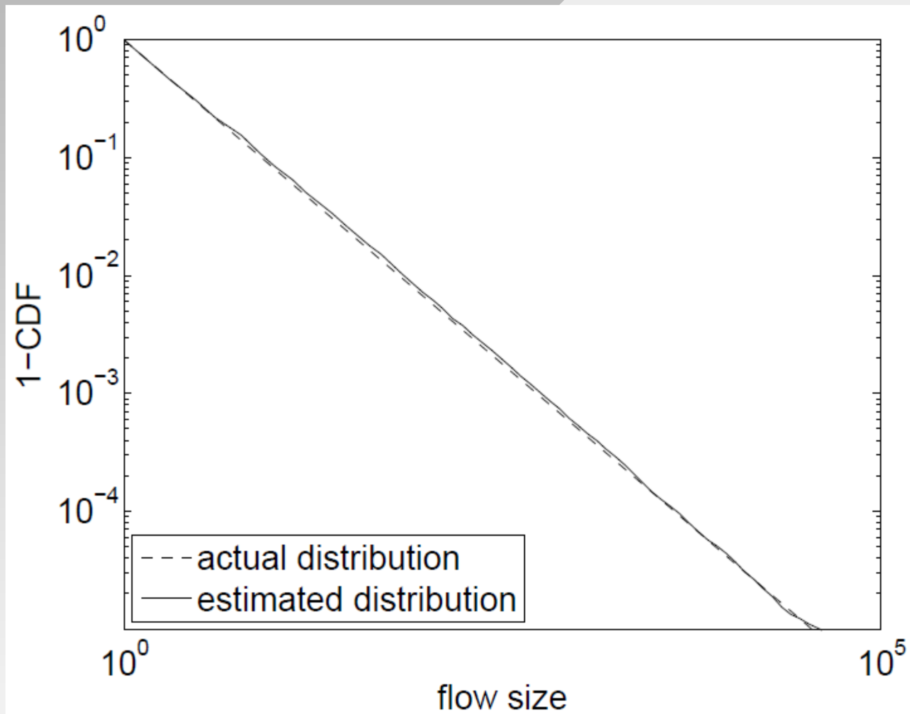

p=0.01

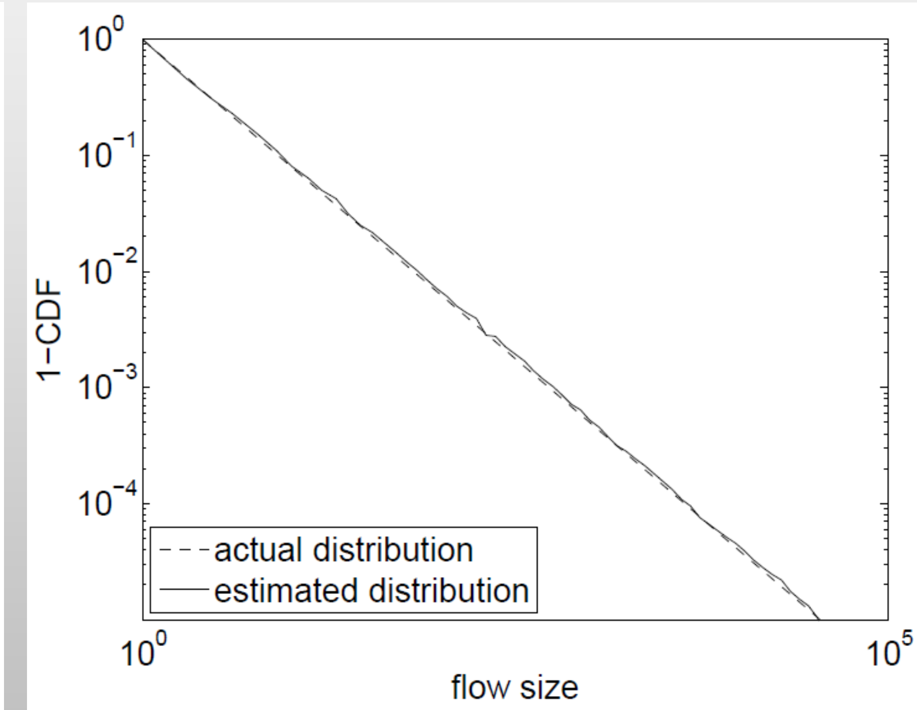

p=0.001

Computer Science, Texas A&M University

URGE Convergence

- We next examine the effect of sample size M on the convergence of estimator

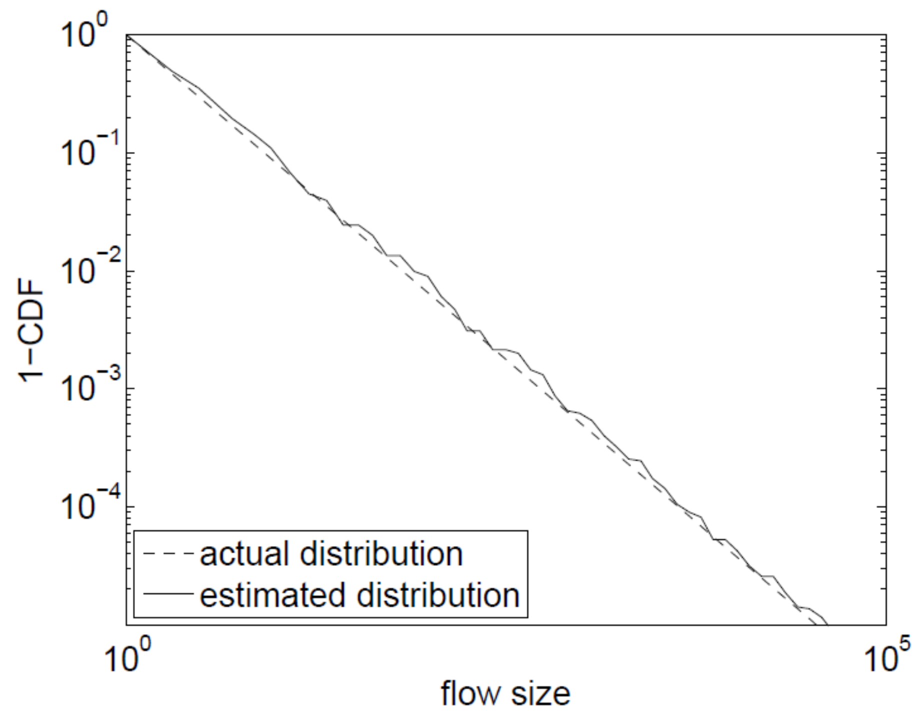

p=10^{-4}, M=3,090

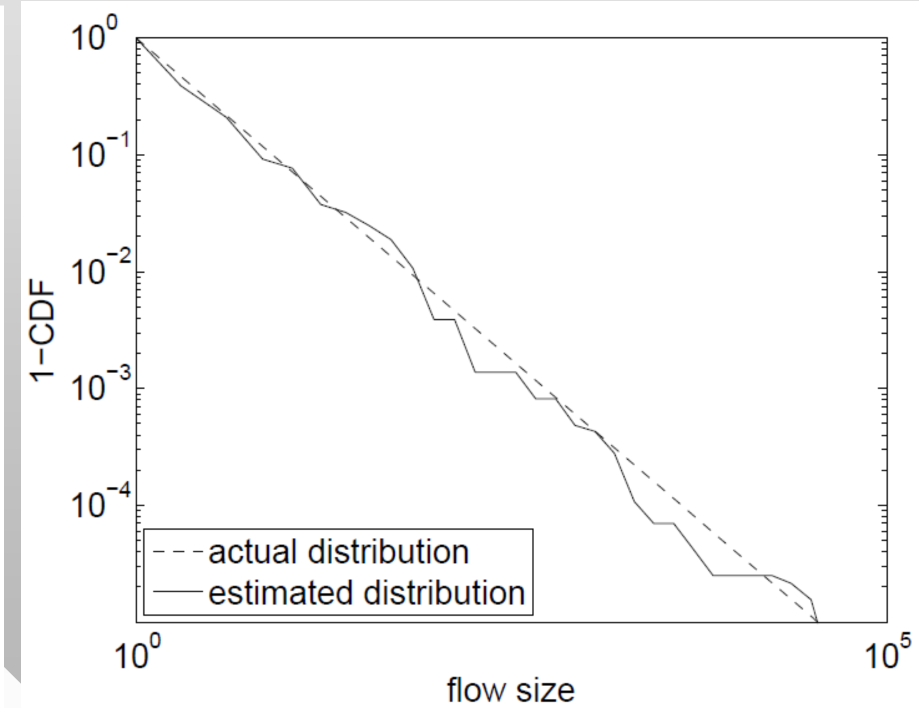

p=10^{-5}, M=337

URGE Convergence 2

- *Theorem 5:* For small constants η and ξ with probability $1-\xi$, following holds for $j \in [1, i+1]$

$$|\tilde{h}_j - h_j| \leq \eta h_j$$

if sample size M is no less than:

$$M \geq \frac{(1-h_i)}{h_i \eta^2} \left(\Phi^{-1}\left(1-\xi/2\right)\right)^2$$

where $\Phi(x)$ is the CDF of the standard Gaussian distribution $N(0,1)$

Performance Evaluation

- We applied our estimation algorithm to the traces collected by NLANR and CAIDA
  - All of them confirm the accuracy of URGE
- As example, we show our experiment on dataset FRG, collected from a gigabit link between UCSD and Abilene

Performance Evaluation – Usage

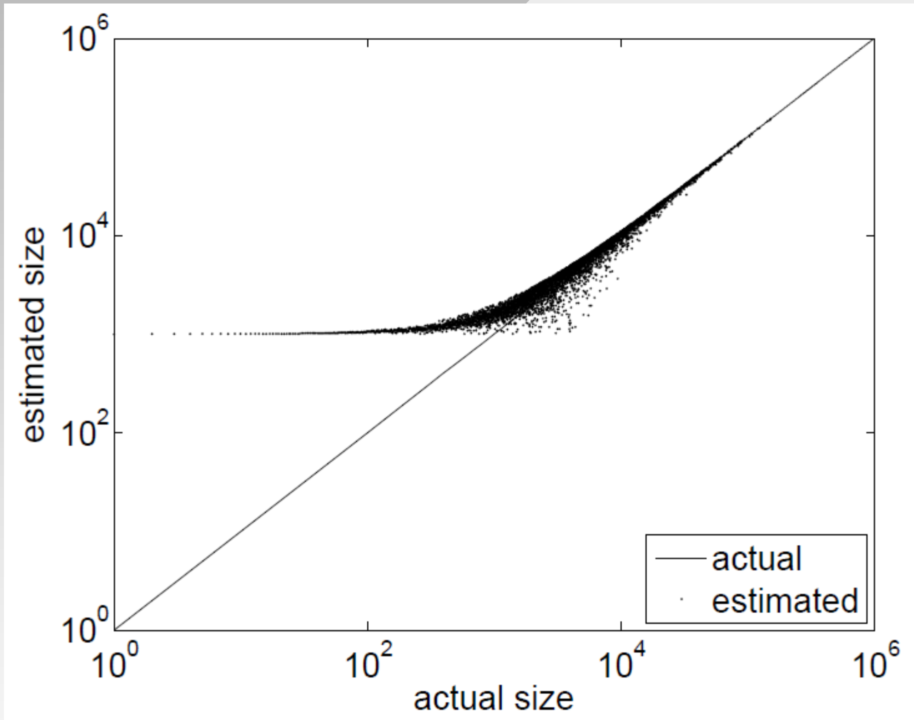

Previous method

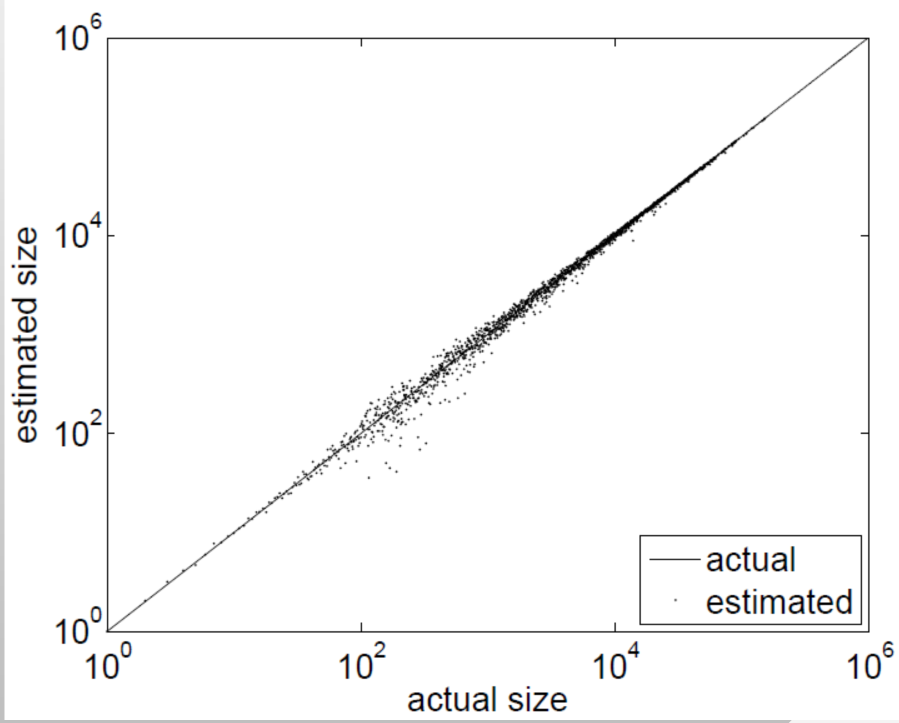

URGE

Performance Evaluation – Distribution

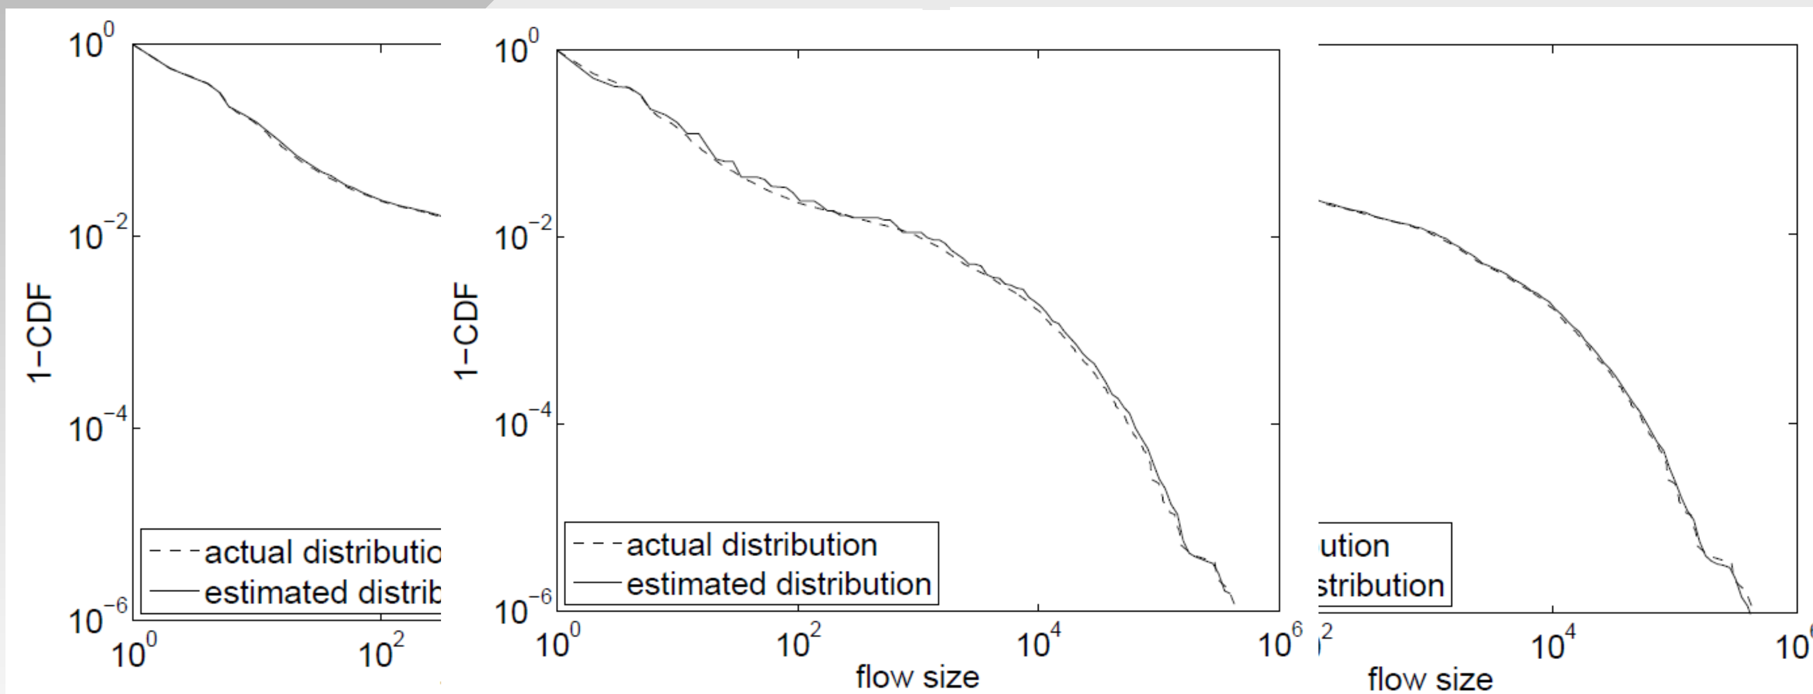

p=0.01 p=0.0001 p=0.001

Conclusion

- We proposed a novel modeling framework for analyzing residual sampling
  - Proved that previous estimators based on RGS had certain bias
- We also developed a novel set of unbiased estimators
  - Verified them both in simulations and on Internet traces
- Results show that the proposed method provides an accurate and scalable solution to Internet traffic monitoring

The End

- Thanks!