Predicting Resource Usage and Estimation Accuracy in an IP Flow Measurement Collection Infrastructure

Nick Duffield, Carsten Lund

AT&T Labs Research, Florham Park

http://www.research.att.com/projects/flowsamp

Background

- Service providers collect enormous amounts of usage data
  - Vast numbers of elementary data
    - packet headers, flow summaries
  - Used for network planning, engineering, control, accounting
  - Many applications need aggregate byte usage, differentiated by e.g. flow key
- Need costly measurement infrastructure to accommodate data
  - At measurement point
  - Transmission from measurement point to collector
  - Data storage
  - Data processing
- Working within constraints of equipment available today
  - No flexibility to (quickly) insert new functionality into routers
  - Have flexibility in rest of measurement infrastructure

Approach

- Statistician's reflex:
  - Sample! Usually simple to implement, quick to execute
  - Allows retention of fine grained detail
    - Compare aggregation, filtering
  - Maintains ability to satisfy ad-hoc queries
    - No need to predefine aggregations
- Questions
  - What is the best sampling strategy given the data, applications?
  - How does one form estimates of byte usage for applications?
  - What are trade-offs between sampling rate and estimation accuracy?
  - How to dimension the measurement infrastructure?
  - How to attribute accuracy to usage estimates?

Sampling and Estimation

- Framework:
  - Regard data elements as deterministic quantities to be estimated
  - Robust: does not utilize statistical model for data during estimation
- Sampling
  - Suppose usage element of size y is sampled with probability p
  - E.g packet of y bytes, or flow with total size y bytes
- If sampled, then renormalize usage estimate as y/p
- Example:
  - If you sample 1 out of N packets ($p=1/N$)
  - Then estimate number of packets in population by N * number sampled packets
- Yields unbiased estimate of actual usage y
  - Division by p compensates against chance not to be sampled
  - Formally, write estimate as random variable $\hat{y} = w\, y / p$
  - w is random indicator of sampling ($w = 1$ if data sampled, 0 if not)
  - $\text{Prob}[w = 1] = p$, hence $\hat{y}$ is unbiased: $E[\hat{y}] = y$

Setting: Flow Measurement Collection Infrastructure

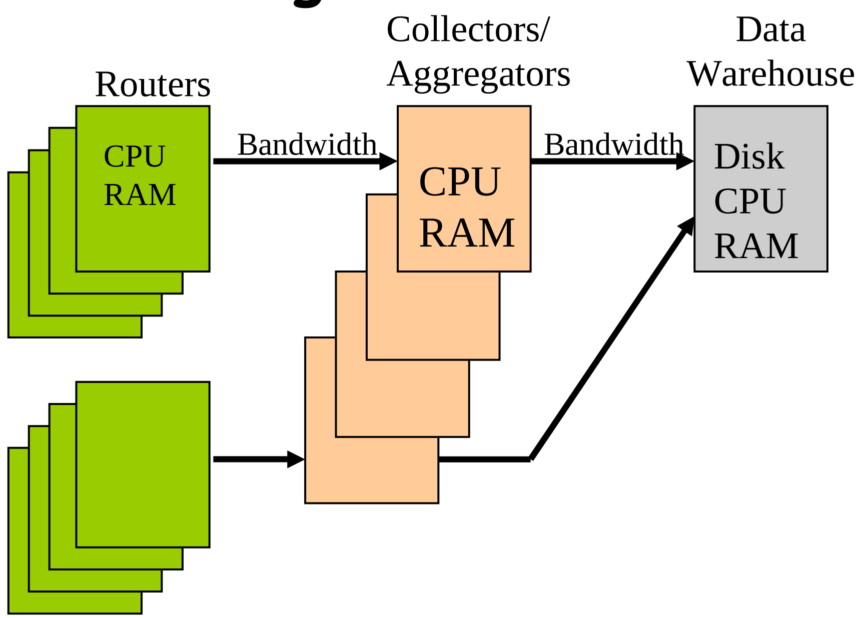

Single router could produce 10's of GBytes of flow records per hour, without sampling

- Three Types of Sampling
  - Routers:
    - packet stream sampled prior to formation of NetFlow records
  - Transmission:
    - flow records possibly lost in transmission (UDP)
  - Collector and Warehouse:
    - size-dependent sampling of flow records (Smart Sampling)
    - potentially multiple levels of smart sampling and aggregation

Three Sampling Steps Composed

- Packet sampling:
  - Uniform packet sampling probability 1/N
- NetFlow loss in transmission:
  - Assume uniform independent transmission with probability q
- Smart Sampling
  - Flow report of size x sampled with probability $p_z(x) = \min\{ 1 , x/z \}$
    - Parameter z is called sampling threshold
  - Large flows ($x \geq z$) always selected
  - Small flows ($x < z$) selected with probability proportional to size

Internet Traffic Flows

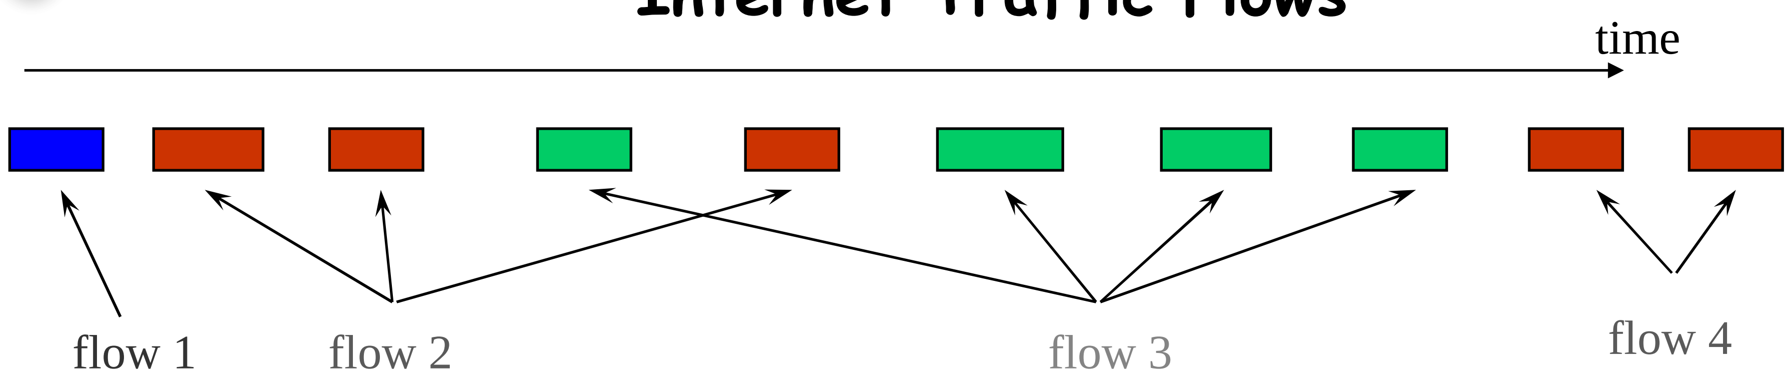

- Internet flows
  - set of packets with common property, observed in some time period
- Common property
  - key: built from header fields (e.g. src/dst address, TCP/UDP ports)
- Flow termination criteria
  - interpacket (inactive) timeout
  - protocol signals (e.g. TCP FIN, RST)
  - ageing (active timeout), flushing, …
- Flow records
  - reports of measured flows exported from routers
  - E.g., flow key, flow packets/bytes, first/last packet time, router state
- Measured flow semantics
  - artificial, may capture appl. transactions if good start/termination criteria

Three Sampling Steps Composed

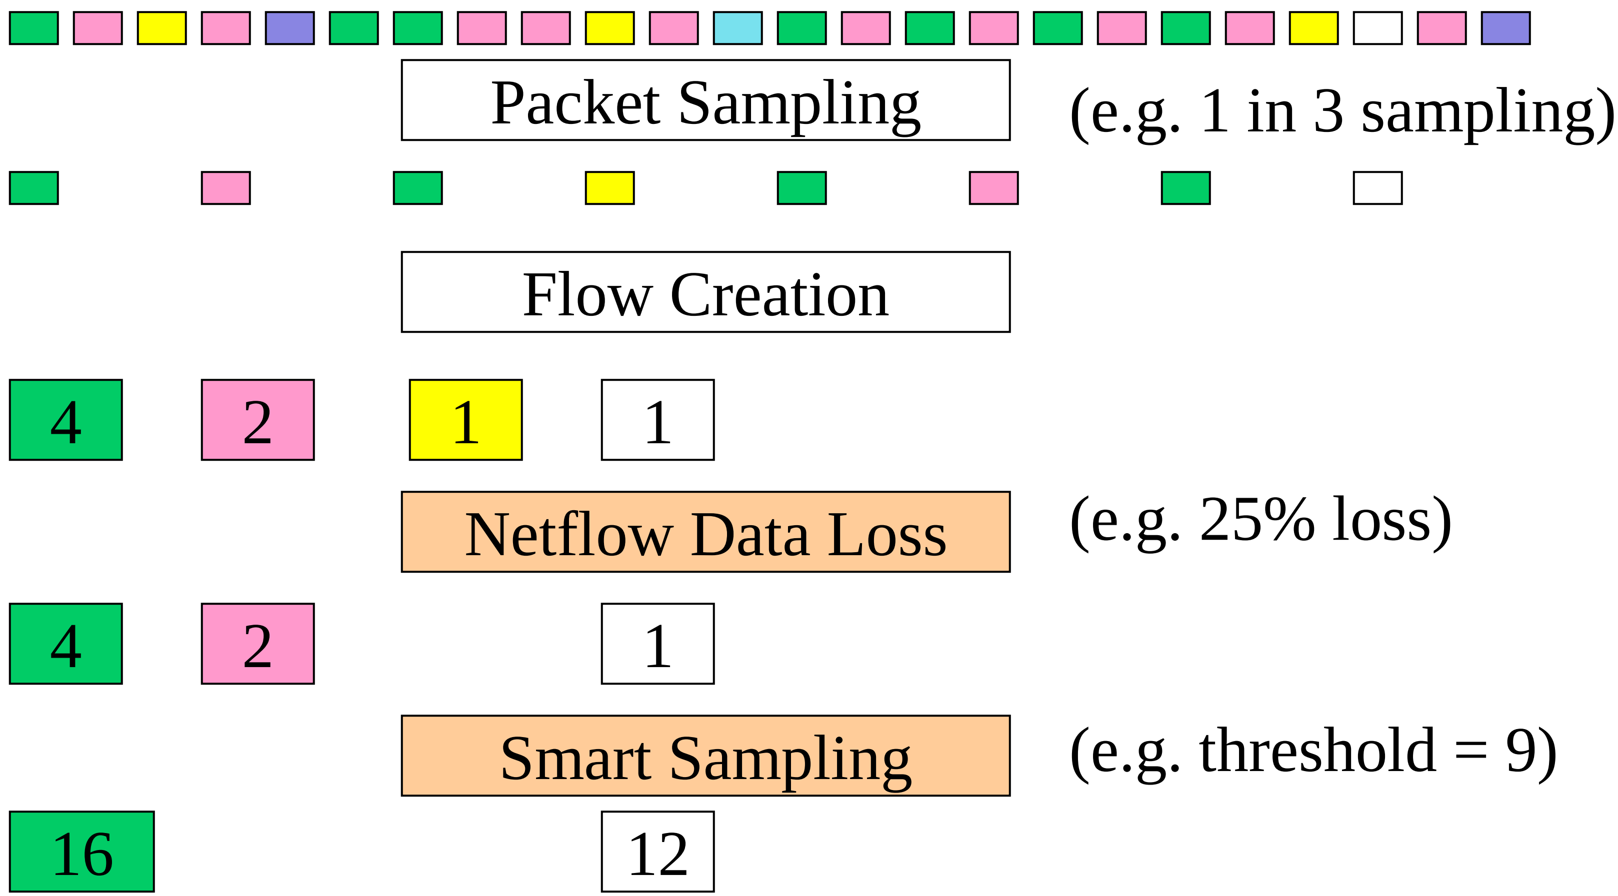

= max{9,4*3}/0.75 = max{9,1*3}/0.75

duffield@research.att.com

AT&T

Why Smart Sampling for Flow Records?

- Targeted towards estimation of byte usage
- Flow lengths have heavy tailed distribution
- Usage estimates very sensitive to omission of large flow
  - Uniform sampling:
    - high estimate variance for large flows
  - Smart sampling: select all large flows (size $x \geq z$)
    - no estimation variance for large flows
- Optimality of flow sampling with probabilities p_z
  - Let X' be estimate of usage X from samples
  - N = mean number of samples taken in estimating X
  - Set cost $C = \mathrm{Var}(X') + z^2 N$
    - Consider as a function of flow sampling probability $p(x)$
  - Use of p_z minimizes the cost C

Program

- Bounding usage estimation variance for combined sampling
  - How to compose variances for chained sampling
  - Want to reliably estimate large contributions to usage
  - Conclusion: need to avoid transmission loss
- Attaching reliability to usage estimates as they are made
  - Estimating the variance of the estimators
- Dimensioning the measurement infrastructure
  - Avoid transmission loss of flow records (need sufficient bandwidth)
  - Enable aggregation of flow records (need sufficient cache space)

Estimation Variance, and its Estimation

AT&T

Chained Estimates and Their Variances

- Multiple stages of sampling
  - Our case: sampling packets, then flow loss, then smart sampling, …
- Successive sampling estimates $y_0, y_1, \ldots, y_n, \ldots$
  - y_0 = actual usage
  - n^{th} sampling stage yields y_n from y_{n-1} ($y_n = 0$ if not sampled)
  - $y_n = w_n \; y_{n-1} / p_n$
    - p_n : probability to be survive sampling stage n, *given* survival to stage n-1
    - w_n: random indicator of stage n sampling (1 if sampled, 0 if not). $P[w_{n=1}] = p_n$
- Conditional unbiasedness (a.k.a. *Martingale* property)
  - $E[y_n | y_{n-1}] = y_{n-1}$
- Conditioning recursion for variances
  - $Var(y_n) = Var(y_{n-1}) + E[Var(y_n | y_{n-1})]$
  - Second term = incremental variance due to n^{th} estimation in chain

Bounding Estimation Variance

- Setup:
  - flow i, size $x_i = \sum_j b_{ij}$ where b_{ij} = size of packet j in flow i, total usage $X = \sum_i x_i$,
- Three stage sampling (packet, flow loss, smart flow) yields estimate X'
- Use chaining formula to calculate Var(X')
  - Result depends in complex way on details of b_{ij}
  - Against our philosophy to exploit distributional details!
- Instead: simple upper bound:
  - $\mathrm{Var}(X') \leq X\,(z + (1-q)\,x_{max} + (N-1)\,b_{max})/q$
  - Bound is tight (i.e. can approach equality in some cases)

duffield@research.att.com

Bounding Estimation Variance (2)

- Variance in estimating usage X
  - Upper bound: $X\,(z + (1-q)\,x_{max} + (N-1)\,b_{max)})/q$
- Bad News: if x_{max} is significant proportion of X
  - Standard Error does not vanish for large X
  - Unless $q = 1$: no transmission loss
  - So need to arrange for sufficient transmission bandwidth
- Good News: if no transmission loss ($q = 1$)
  - Standard Error bounded above by $((z + (N-1)b_{max})/X)^{1/2}$
  - Vanishes for large X: more reliable estimation of larger components
- Design criteria
  - N bounded below for router (need time for flow cache lookup)
    - E.g. N = 100
  - Roughly equal variance from packet and flow sampling if $z = (N-1)b_{max}$
    - Can have z up to $(N-1)b_{max}$without introducing much greater variance

duffield@research.att.com

Bounding Estimation Variance (3)

- No loss: $q = 1$: $(\text{Standard\_Error}(X'))^2 \leq z X + b_{max} (N-1) X$
- Example: $N = 100$, $z = 10^6$
- Data: ~2M raw, unsampled NetFlow records

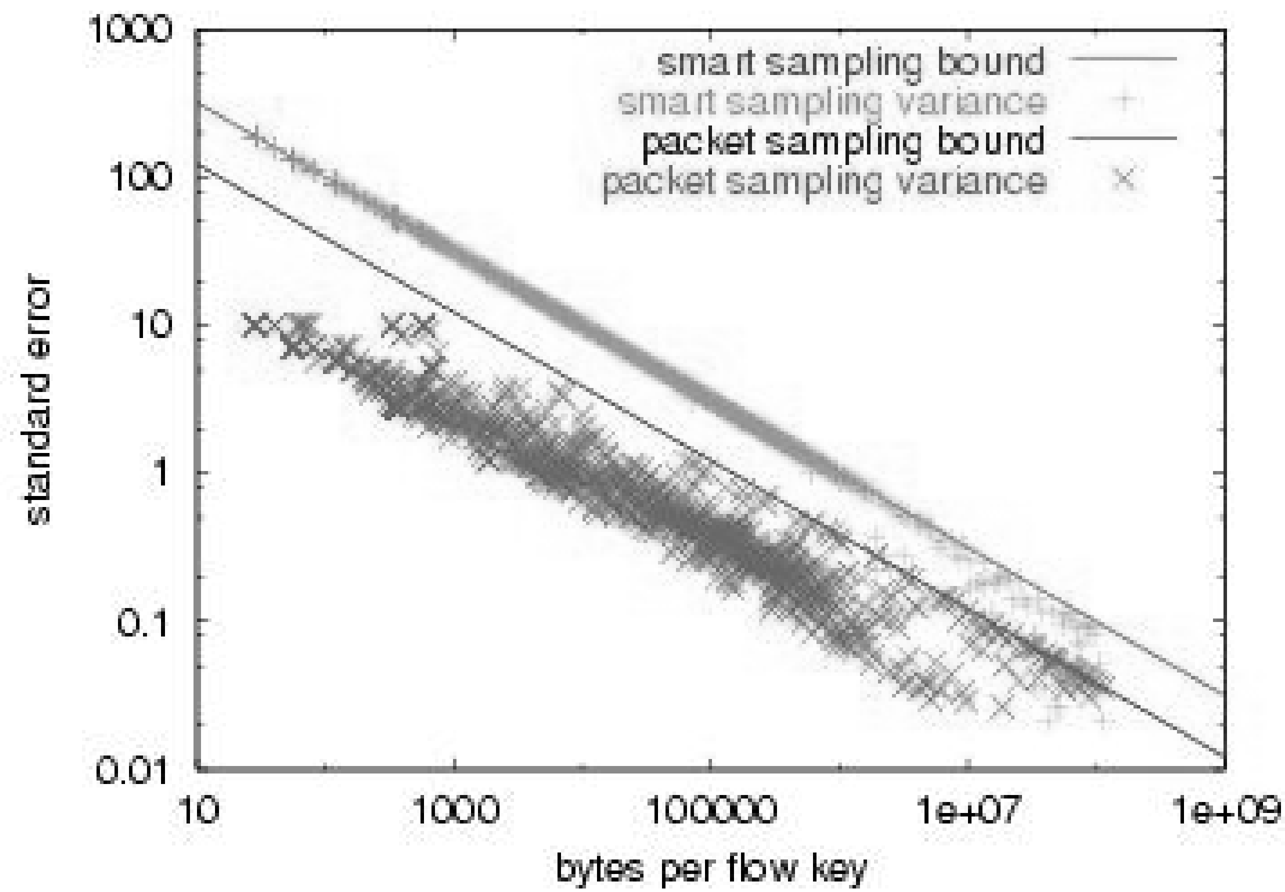

duffield@research.att.com

Estimating the Variance of Usage Estimates

- Want to calibrate usage estimates for accuracy
- Suppose $Var(X') = v(X)$ some function v
  - Can we just estimate variance $Var(X')$ by $v(X')$?
  - No: biased, since does not compensate for chance not to be sampled
- Chained Variance Estimates
  - Recall $y_n = w_n x_{n-1} / p_n$ and $Var(y_n) = Var(y_{n-1}) + E[Var(y_n|y_{n-1})]$
- Unbiased estimation of variance increment:
  - $v_n = w_n y^2_{n-1} (1-p_n)/p^2_n$ obeys $E[v_n \mid y_{n-1}] = Var(y_n|y_{n-1})$
- Operation in measurement infrastructure:
  - Associate a variance with each object throughout sampling chain
  - Prior to n^{th} sampling stage we know y_{n-1}, p_n
  - If object survives n^{th} sampling stage, add $y^2_{n-1} (1-p_n)/p^2_n$ to variance
  - If aggregating objects, associate with sum of component variances

Dimensioning Bandwidth for Flow Records

AT&T

Dimensioning Bandwidth for Flow Records

- Motivation:
  - Allocate sufficient bandwidth to avoid loss of flow records
  - In transmission (otherwise high estimation variance)
  - And at collector
- Design basis
  - Set of flow records (n_i packets, b_i bytes, t_i duration)
  - Flows collected over time period τ
- Sampling Rates
  - Rate of production of packet sampled flow statistics at router
  - Rate of production of smart sampled flow statistics at collector
- Aim:
  - Obtain relatively simple bounds and approximations
    - as function of flow characteristics
  - Without need to simulate sampling

Flow splitting under packet sampling

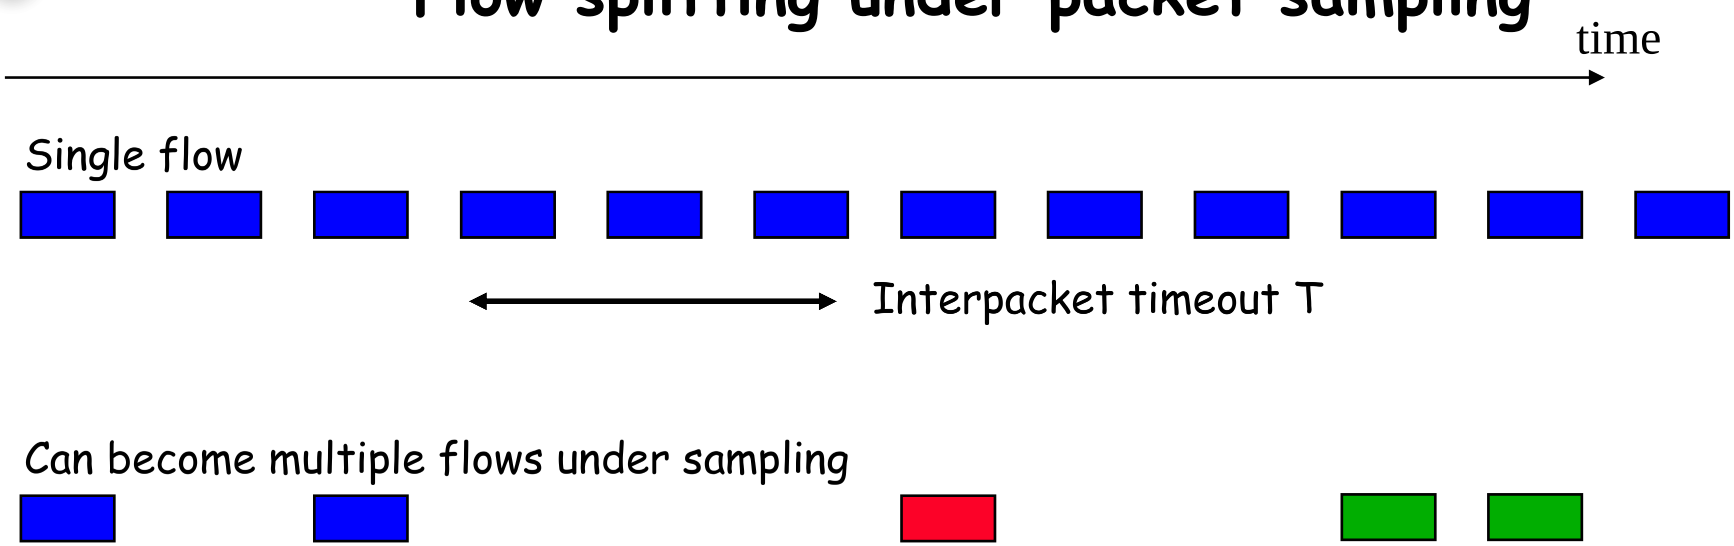

- Sampling increases interpacket times
- Flow splitting when interpacket time exceed active timeout
- Flows vulnerable to splitting: call these **sparse**
  - flows with many packets, not too fast packet rate
  - e.g. streaming, p2p applications, VoIP

Model of Flows and Splitting

- Flow
  - n packets, duration t
  - model packets as randomly distributed within duration t
- Sampling
  - period N, active timeout T
  - Packets selected independently
- Flow split occurs
  - When time between successive sampled packets exceeds T
- Each flow (n,t) yields random number of packet sampled flows
- Theorem: mean number of packet sampled flows is
  - $f(n,t;N,T) = 1 + (1+(\kappa-1)/N)^{n-1} ((\kappa(n-1)+1)/N - 1)$ with $\kappa = \max\{0, 1-T/t\}$
- Estimate mean transmission rate of flow statistics
  - $R = \tau^{-1} \sum_i f(n_i,t_i;N,T)$

AT&T

Rate of Smart Sampled Flow Records

- Smart sampling flow records with byte threshold z
- Simple Upper Bound
  - $R_s \le \min\{ R, B/z \}$ where B is total byte usage represented in flows
- Less Simple Estimate
  - $R_s \approx \tau^{-1} \sum_i \min\{ f(n_i, t_i; N, T), b_i/z \}$
  - Reasoning:
    - suppose b bytes are evenly split amongst f sampled flows
    - average # smart sampled flows is f $\min\{1, b/(fz)\} = \min\{f, b, z\}$

Rate of Smart Sampled Flow Records (2)

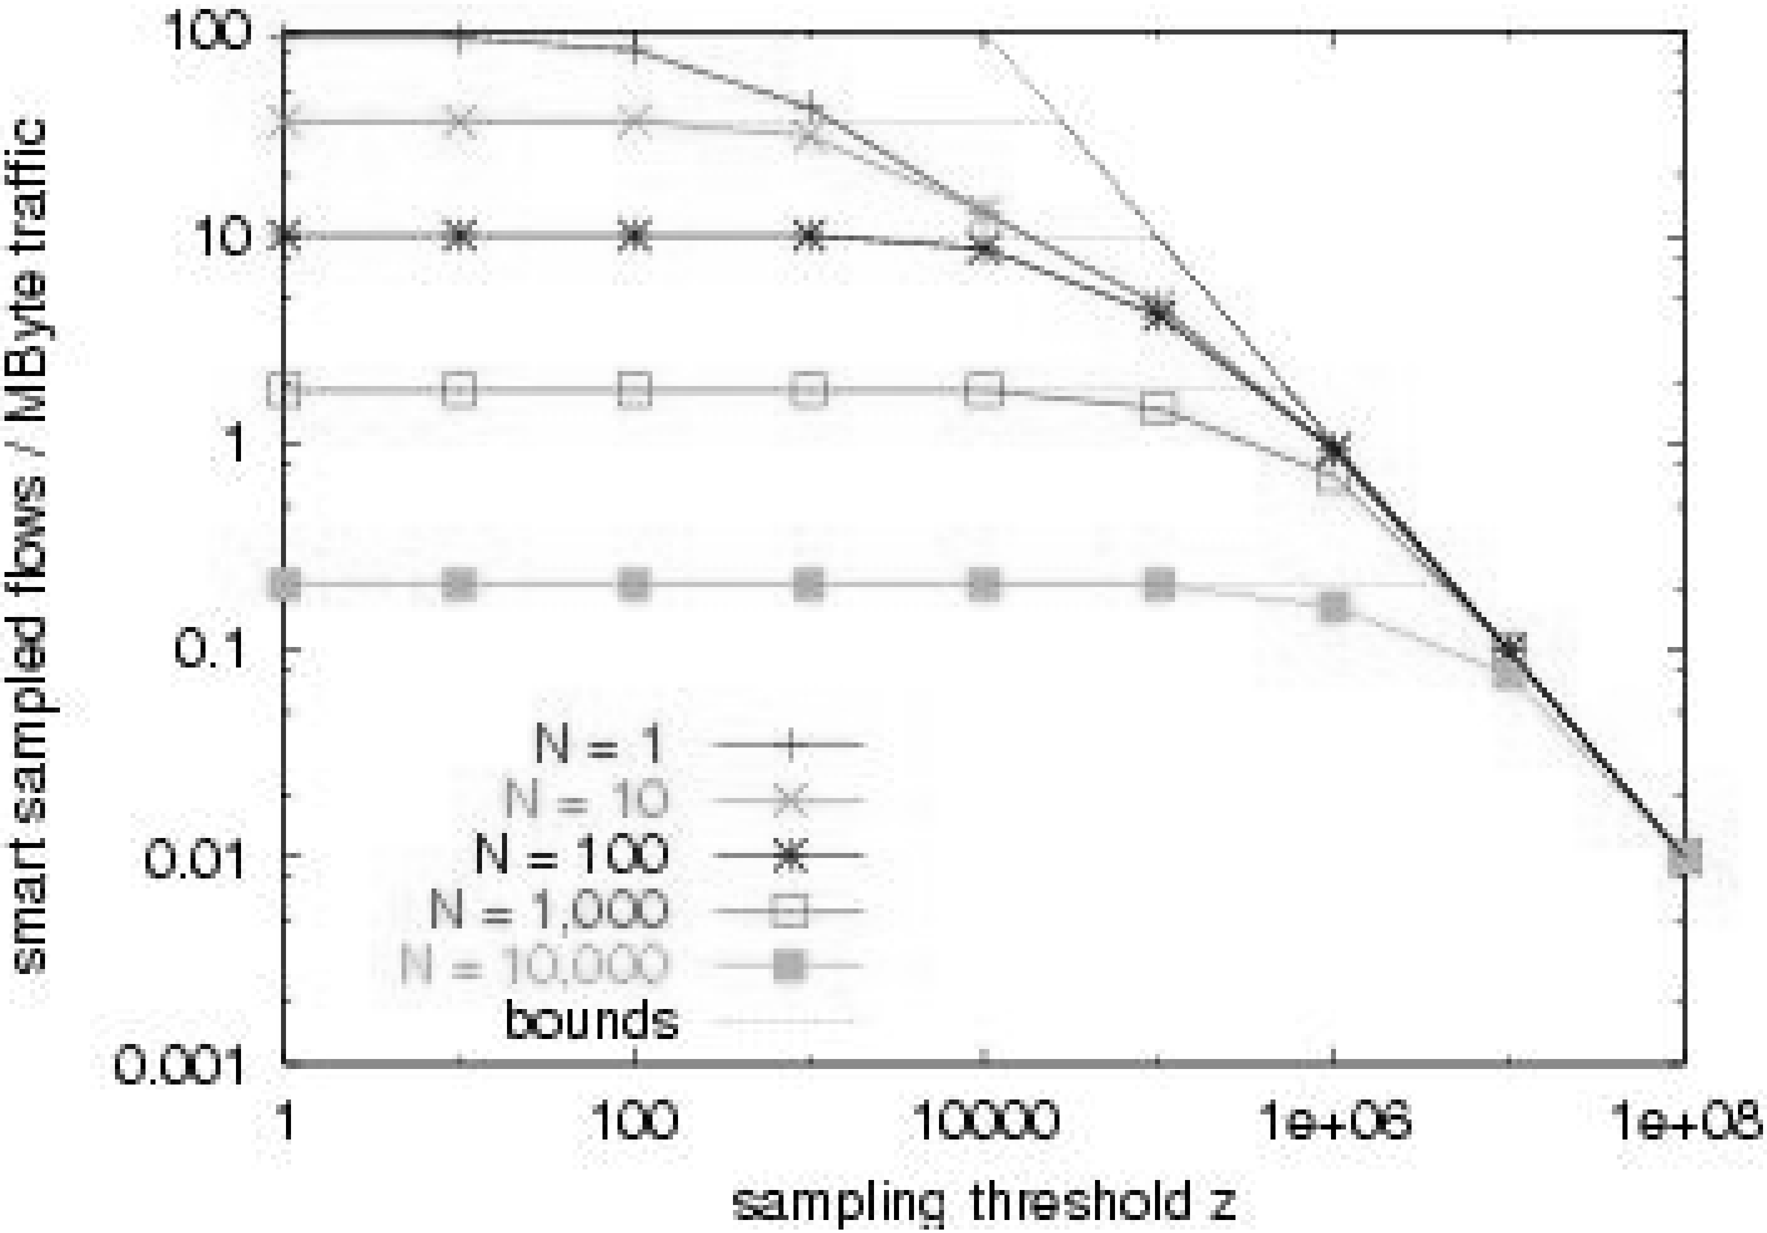

Dimensioning Aggregation Resources at Collector

- Example: want to aggregate keys over BGP prefix
- Need: rate $R_{s,agg}$ at which new aggregate keys are created from smart sampled flow records
- Simple upper bound:
  - $R_{s,agg} \le \min \{ R_{agg} , B/z \}$
    - where R_{agg} = rate of new aggregate keys in non-smart sampled flows
- Less simple approximation
  - Let I_k be set of flows wit aggregate key k

$$R_{s,agg} = \tau^{-1} \sum_k \left(1 - \prod_{i \in I_k} (1 - q_i) \right)$$

  - Where q_i is probability that at least one sampled flow arising from flow (n_i, b_i, t_i) survives smart sampling
  - Approximation: write $f(n_i,t_i) = f_i$

$$q_i \approx \begin{cases} f_i\, p_z(b_i/f_i) : & \text{if } f_i < 1 \\ 1 - (1 - p_z(b_i/f_i))^{f_i} : & \text{if } f_i \ge 1 \end{cases}$$

Dimensioning Aggregation Resources at Collector (2)

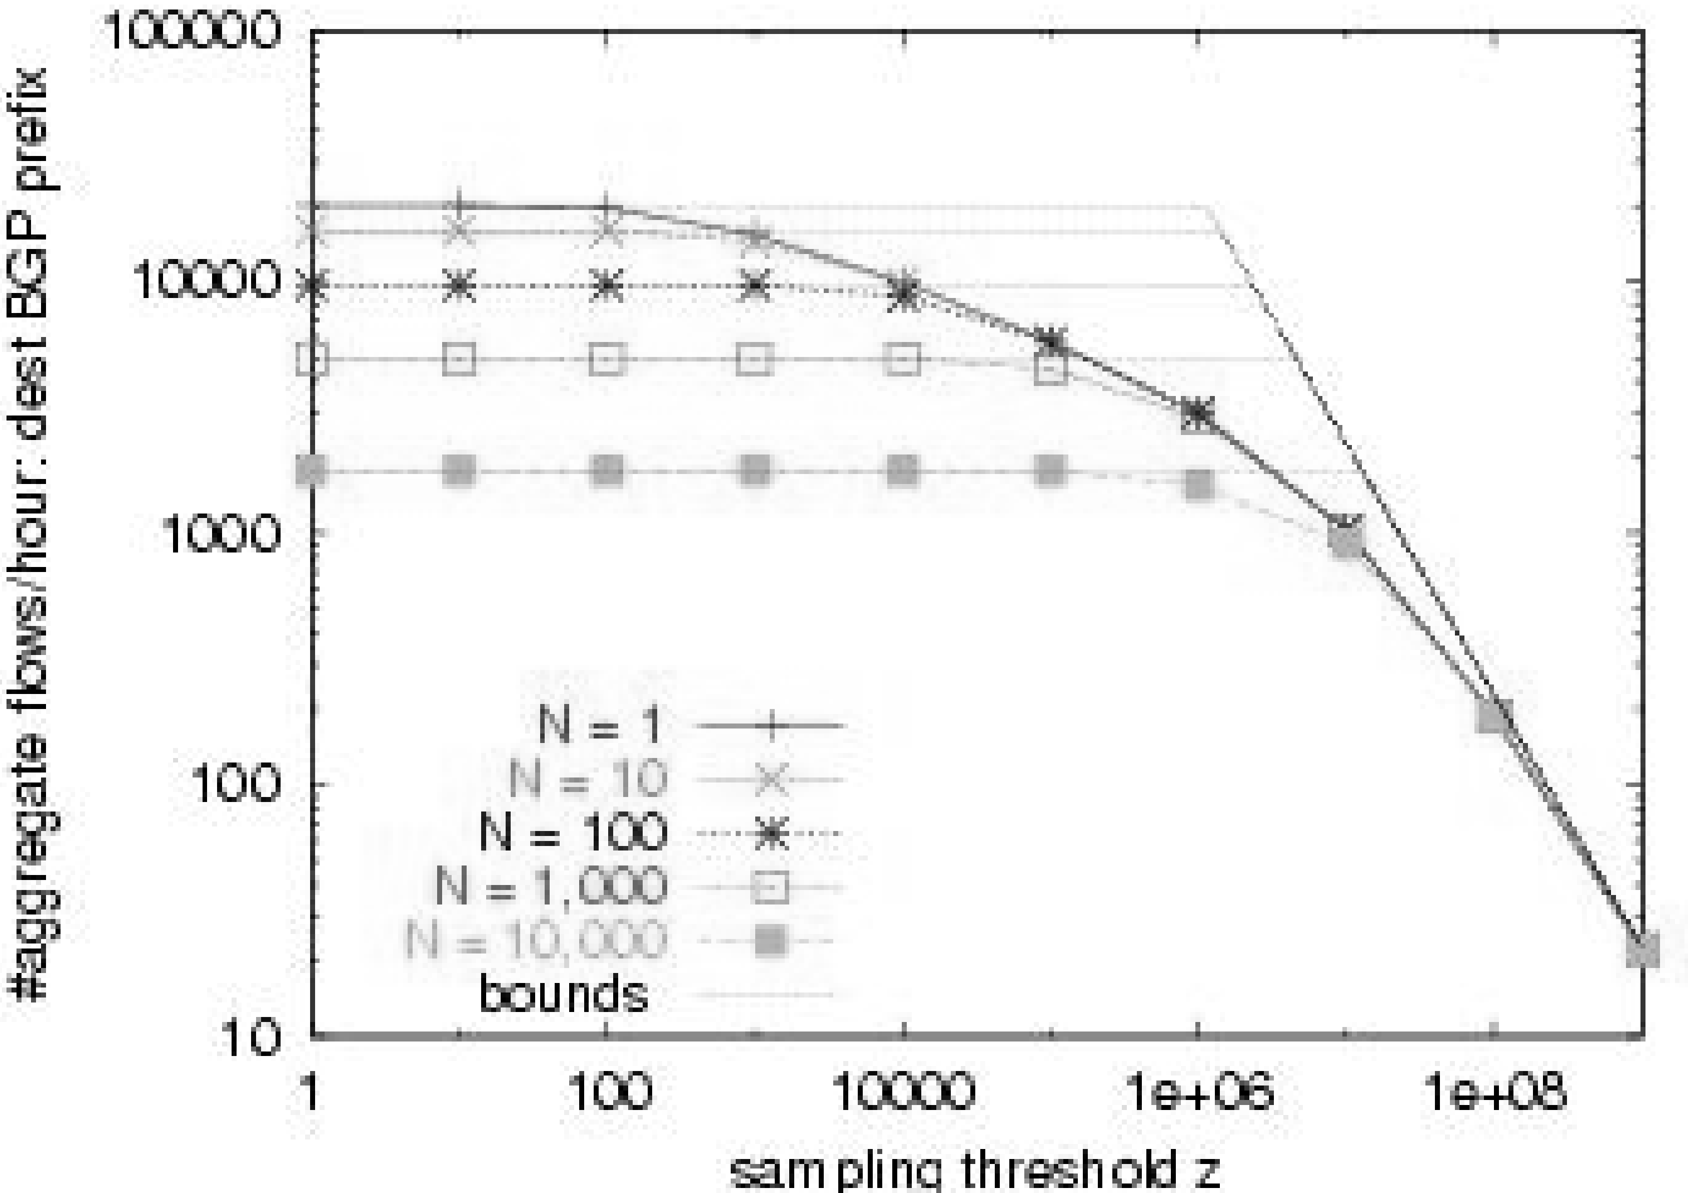

Summary

- Setting
  - Packet and flow sampling in measurement collection infastructure
  - Trade off usage estimation variance, resource usage
- Usage Estimator Variance
  - Need to avoid flow loss in transmission
  - Standard error is estimating usage X bounded by $(z X + b_{max}(N-1) X)^{1/2}$
  - Estimating estimator variance online:
    - Maintain variance for each object, increment if sampled
- Resource usage
  - Transmission resources (need to dimension for no loss)
  - Aggregation resources
  - Want quick formula to predict resource usage w/o simulation
    - from trace or model of flow characteristics
  - Simple bounds, also approximations
- http://www.research.att.com/projects/flowsamp