On the Predictability of Large Transfer TCP Throughput

Qi He
Constantine Dovrolis
Mostafa Ammar

College of Computing
Georgia Institute of Technology

Outline

- TCP throughput prediction: problem statement and motivation
- Formula-Based (FB) prediction
  - A formula-based predictor
  - Types of FB prediction errors
  - Experimental evaluation
- History-Based (HB) prediction
  - Typical history-based predictors
  - Dealing with outliers and level shifts
  - Experimental evaluation
- Predictability factors
  - What makes some paths less predictable than others?

Problem Statement and Motivation

- Objective:
  - *Predict the throughput of a bulk TCP transfer on a given path*
- Motivation:
  - Server selection
  - Overlay/multi-homed routing
  - Load balancing
  - Grid computing
  - P2P downloading

Constraints and Assumptions

- Prediction is needed before the start of transfer
- Performing “test TCP transfer” just for prediction is too intrusive/slow
- Measuring certain “lightweight path characteristics” (e.g., loss rate or RTT) is *not* intrusive

Two Classes of TCP Throughput Predictors

Prediction Method	Basis	Inputs	Advantages	Issues
Formula Based (FB)	Analytical model for TCP throughput	Estimates of path's RTT and loss rate	No previous transfers required	Prediction accuracy?
History Based (HB)	Time series forecasting theory	History of previous TCP transfers on the same path	Prediction based on actual TCP transfers	Prediction accuracy?

Main Contributions

- Evaluate prediction accuracy for FB and HB predictions
  - FB can be significantly inaccurate, especially for congestion-limited flows
  - HB is quite accurate even with simple linear predictors and sporadic previous samples
- Explain major causes of prediction errors in terms of underlying network and TCP behavior
  - Focus on cause-effect relations, rather than black box evaluation
- Study effects of path properties and transfer characteristics on prediction accuracy
  - Load, degree of multiplexing
  - Receiver window, transfer frequency

Outline

- TCP throughput prediction: problem statement and motivation
- Formula-Based (FB) prediction
  - A formula-based predictor
  - Types of FB prediction errors
  - Experimental evaluation
- History-Based (HB) prediction
  - Typical history-based predictors
  - Dealing with outliers and level shifts
  - Experimental evaluation
- Predictability factors
  - What makes some paths less predictable than others?

TCP Throughput Model

- Analytical model of the expected TCP throughput R as a function of several path characteristics

$$R = f\,(T, p) \qquad\qquad (p > 0)$$

  - T, p: RTT and loss rate experienced during the flow
  - We use PFTK model by Padhye et. al (Sigcomm '98)

$$R = \min\left(\frac{M}{T\sqrt{\frac{2bp}{3}} + T_0 \min\left(1, \sqrt{\frac{3bp}{8}}\right) p (1 + 32p^2)}, \frac{W}{T} \right), p > 0$$

M: path MTU (Maximum Transfer Unit)
W: TCP maximum congestion window
T_0: TCP retransmission timeout
b: segments released per new ACK

An FB Predictor

- Measure loss rate p’, RTT T’ before the target flow starts
  - Typical measurement: periodic probing, e.g., Ping
- Apply T’ and p’ to the throughput equation

$$\hat{R} = f(T', p') \qquad (p' > 0)$$

- With the PFTK model…

$$\hat{R} = \min\left(\frac{M}{T'\sqrt{\frac{2bp'}{3}} + T_0 \min\left(1, \sqrt{\frac{3bp'}{8}}\right) p'(1+32p'^2)}, \frac{W}{T'}\right), \quad p' > 0$$

An FB Predictor

- Measure loss rate p′, RTT T′ before the target flow starts
  - Typical measurement: periodic probing, e.g., Ping
- Apply T′ and p′ to the throughput equation

$$\hat{R} = f(T', p') \qquad (p' > 0)$$

- With the PFTK model…

$$\hat{R} = \begin{cases} \min\left(\dfrac{M}{T'\sqrt{\dfrac{2bp'}{3}} + T_0 \min\left(1, \sqrt{\dfrac{3bp'}{8}}\right) p(1+32p'^2)}, \dfrac{W}{T'}\right), & \text{if } p' > 0 \\ \min\left(A', \dfrac{W}{T'}\right), & \text{if } p' = 0 \end{cases}$$

A′: available bandwidth estimation

Potential Issues with FB Prediction

- Differences between T′ and T, p′ and p

	T′, p′	T, p
Temporal:	before flow	during flow
Sampling:	periodic probing	TCP “sampling”

Issue	Effect
Additional load of the target flow may increase *T, p*	Overestimate throughput
Adaptive and bursty TCP sampling vs. non-adaptive periodic sampling	Underestimate or overestimate throughput

Evaluation of Throughput Prediction Accuracy

- Epoch:
  - IPerf for TCP transfers, pathload for available bandwidth, ping (interval: 100ms, pkt size: 41bytes) for RTT & loss rate

One Measurement Epoch

Evaluation of Throughput Prediction Accuracy

- Epoch:
  - IPerf for TCP transfers, pathload for available bandwidth, ping (interval: 100ms, pkt size: 41bytes) for RTT & loss rate

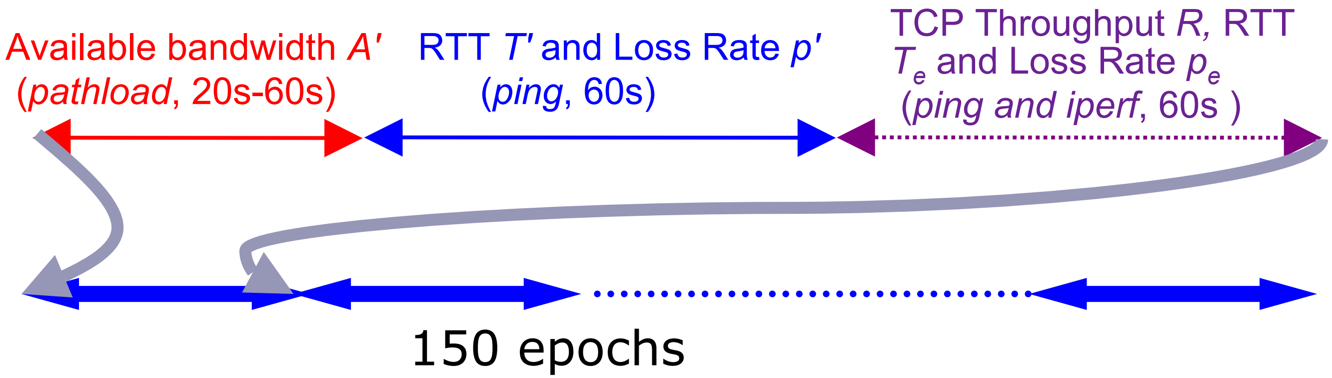

- Each trace consists of 150 consecutive epochs
- We used 35 Internet paths; 7 traces on each path; hosts in US, Europe, Korea

Prediction Error Metrics

- Relative Error

$$E = \frac{\hat{R} - R}{\min(\hat{R}, R)}$$

$\hat{R}$ =(1/w) R, and $\hat{R}$ = wR both have: |E|=w-1

e.g., $\hat{R}$ = ½ R, and $\hat{R}$ = 2 R both have: |E|=1

- Root Mean Square Relative Error

$$RMSRE = \sqrt{\frac{1}{n}\sum_{i=1}^{n} E_i^2}$$

CDF of FB Prediction Error

- Overestimation by >100% (E>1) for 40% of the measurements
- Dominance of overestimation errors (E>0)
  - Prevalent occurrences of T′ < T and p′ < p

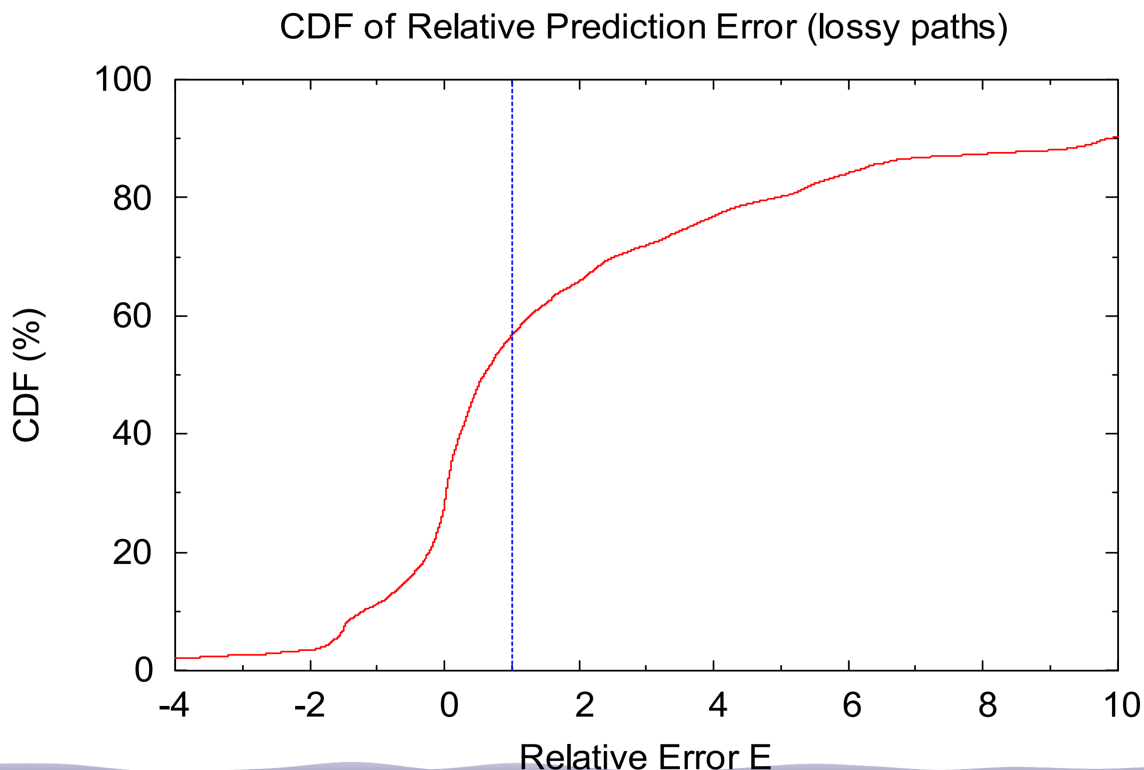

CDF of FB Prediction Error

- Overestimation by >100% (E>1) for 40% of the measurements
- Dominance of overestimation errors (E>0)
  - Prevalent occurrences of T′ < T and p′ < p

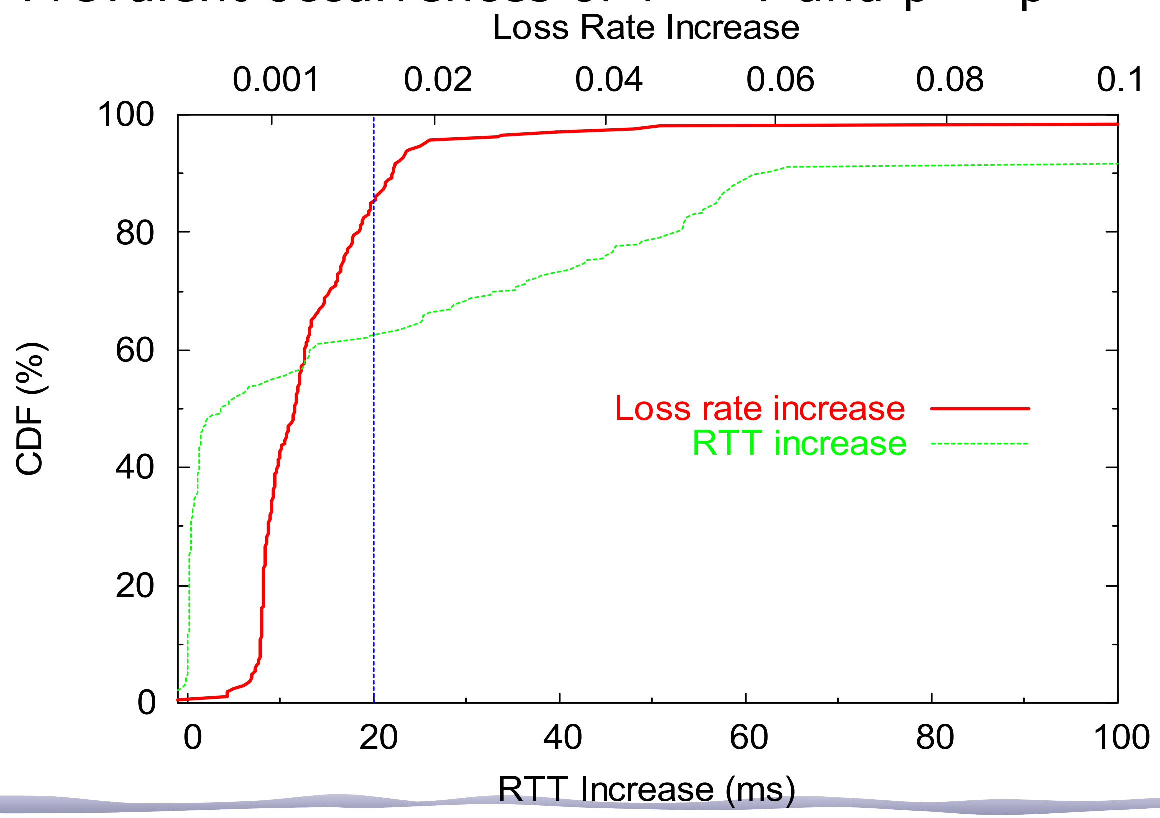

Errors Due to Sampling Differences

- Prediction using Ping RTT & loss rate measurements during target flow
- Prediction errors are still significant, but overestimation & underestimation are almost symmetric

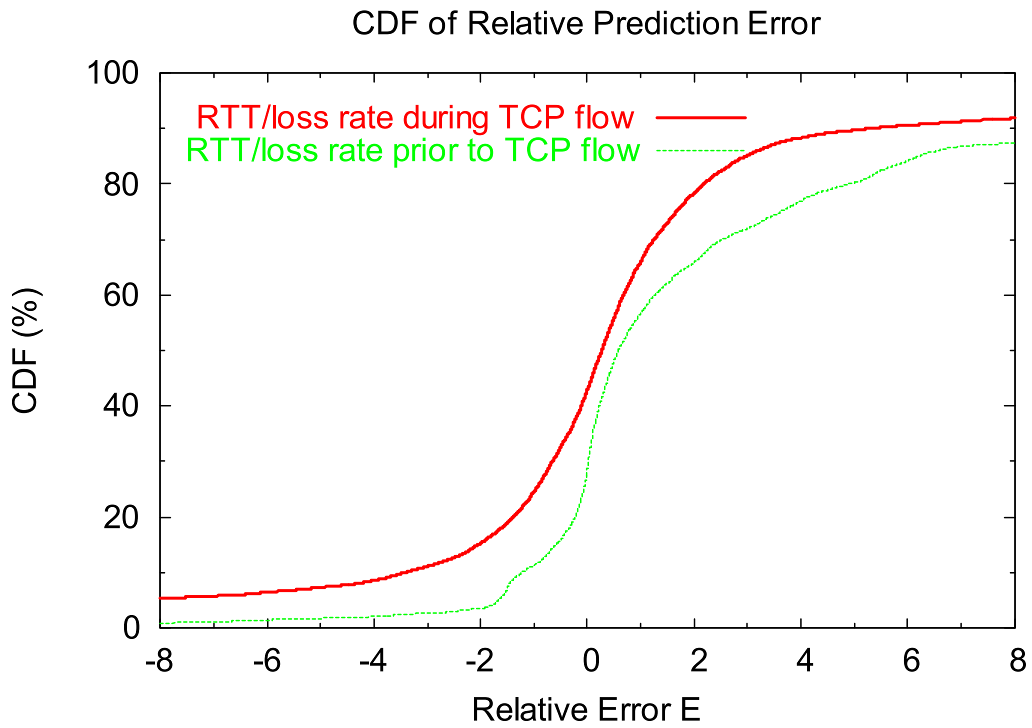

Prediction Accuracy vs. Actual Throughput

- Large errors are more common in lower-throughput paths
- Explanation: in a congested path, slight load increase causes large loss rate increase

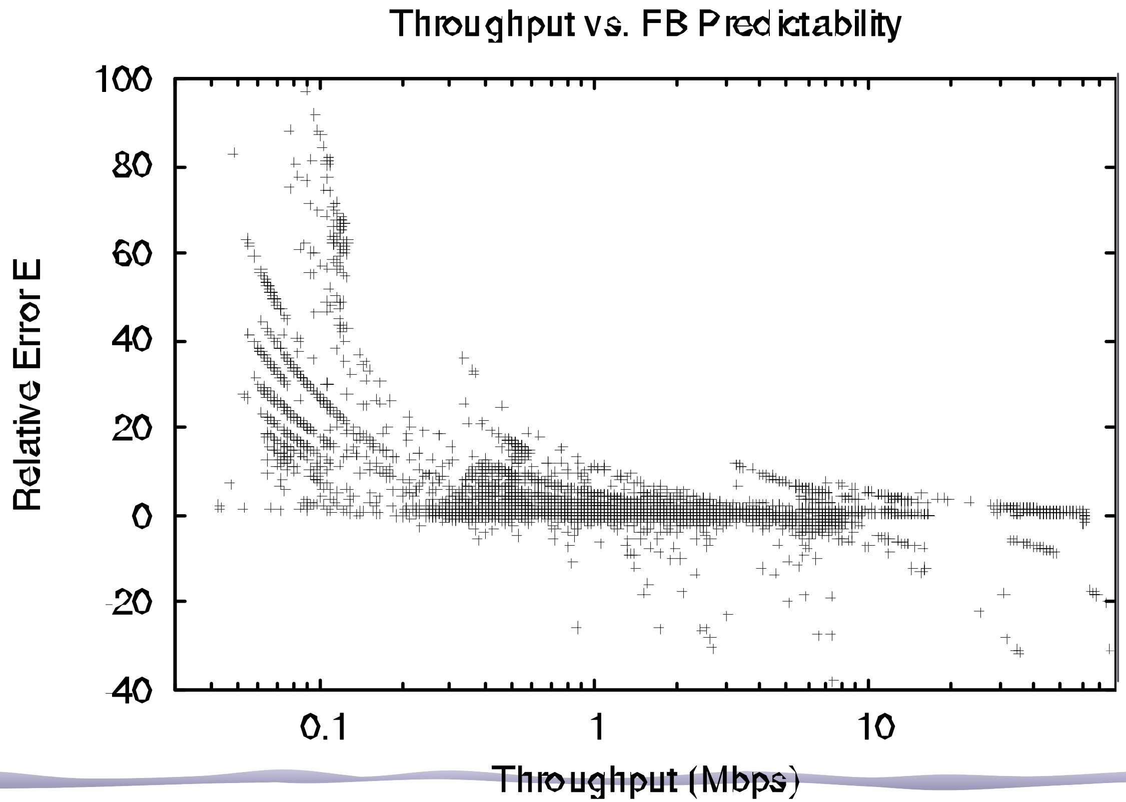

Window-limited Flows

- Throughput is more predictable for window-limited TCP flows
- Explanation: window-limited flows do not saturate path's bottleneck

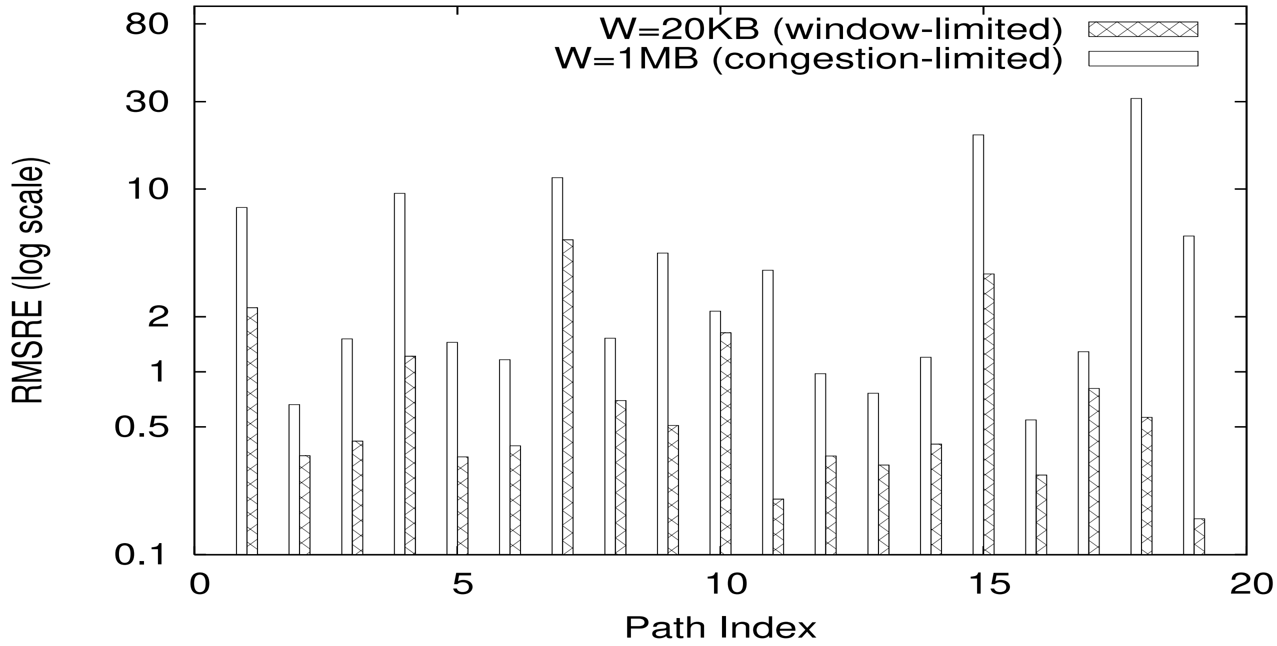

Outline

- TCP throughput prediction: problem statement
- Formula-Based (FB) prediction
  - A formula-based predictor
  - Types of FB prediction errors
  - Experimental evaluation
- History-Based (HB) prediction
  - Typical history-based predictors
  - Dealing with outliers and level shifts
  - Experimental evaluation
- Predictability factors
  - What makes some paths less predictable than others?

History-Based Prediction

- General one-step forecasting problem

$$\hat{R}_n = f(R_1, R_2, \ldots, R_{n-1})$$

- We only consider simple linear predictors
  - Moving Average (MA)

$$\hat{X}_{i+1} = \frac{1}{n}\sum_{k=i-n+1}^{i} X_k$$

  - Exponentially Weighted Moving Average (EWMA)

$$\hat{X}_{i+1} = \alpha X_i + (1-\alpha)\hat{X}_i$$

  - Non-seasonal Holt-Winters (HW)
    - An EWMA variation that captures the time series trend

$$\hat{X}_i^f = \hat{X}_i^s + \hat{X}_i^t$$

$$\hat{X}_{i+1}^s = \alpha X_i + (1-\alpha)\hat{X}_i^f \qquad \text{--smoothing}$$

$$\hat{X}_{i+1}^t = \beta(\hat{X}_i^s - \hat{X}_{i-1}^s) + (1-\beta)\hat{X}_{i-1}^t \qquad \text{--trend}$$

Level Shifts (LS) and Outliers (OL)

- Why are LS and OL undesirable?
  - Cause large prediction errors and differences among predictors; complicate the analysis of HB predictability
- Dealing with LS and OL is more important than choosing among predictors
  - Actions: ignore OL, restart predictor upon LS

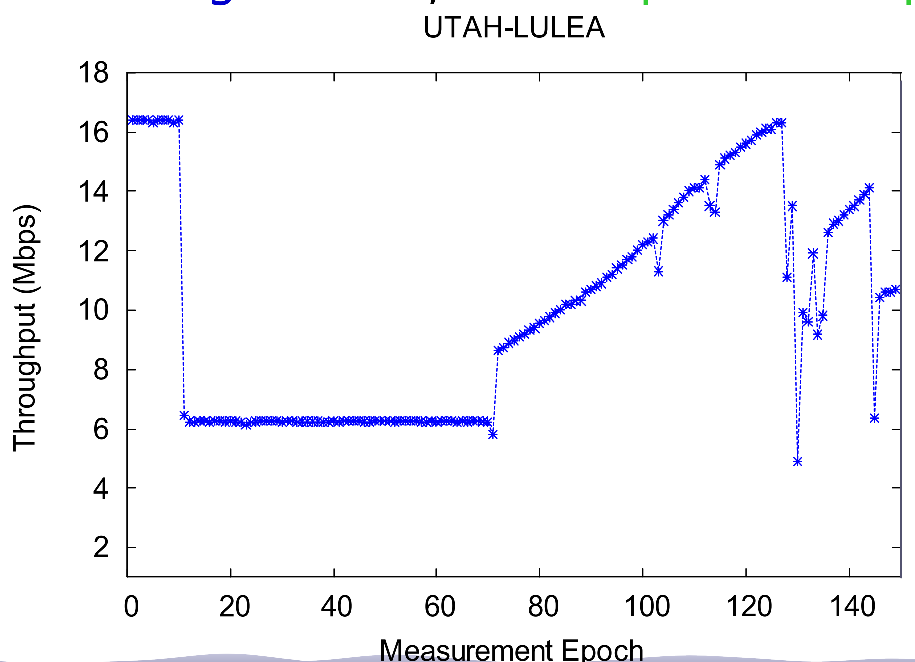

Overall HB Prediction Accuracy

- HB prediction is much more accurate than FB prediction
- 90% of traces have RMSREs < 0.4 (with LS/OL detection)
- With LS/OL detections, the choices of predictor and of predictor parameters make little difference

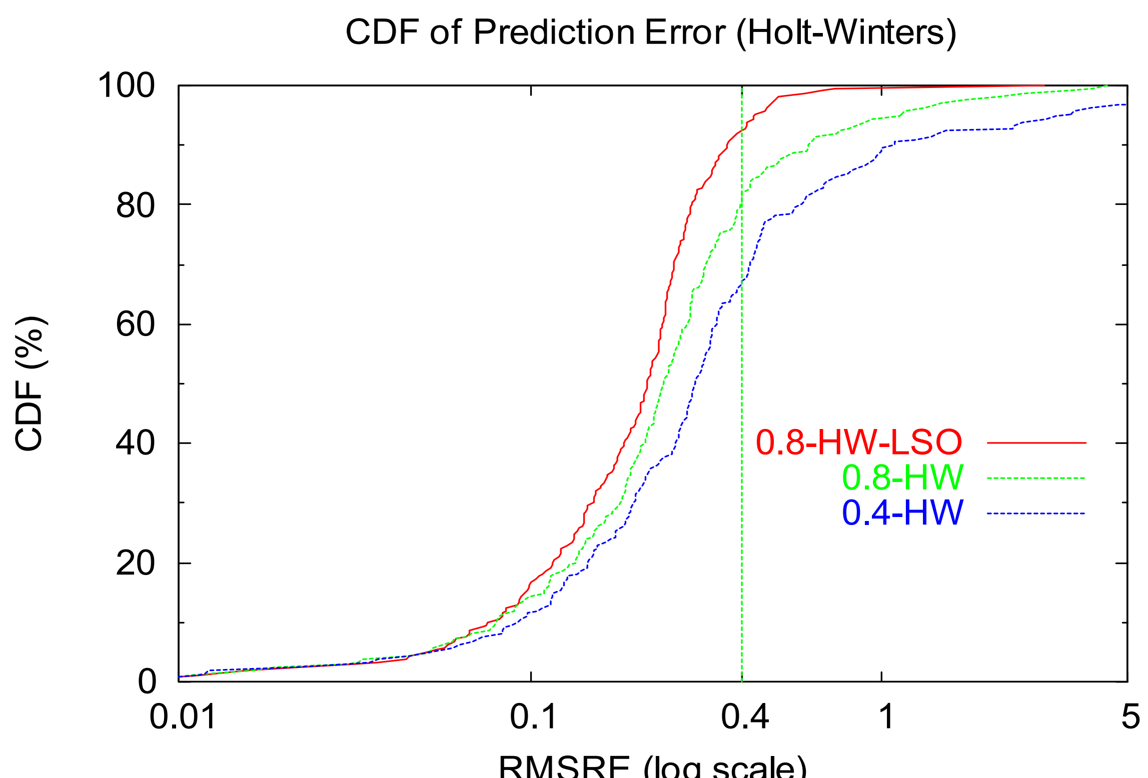

Effect of Measurement Frequency

- Longer measurement period does not degrade accuracy significantly
  - Even with single transfer every 24 minutes, RMSRE is below 0.4 in 75% of the traces

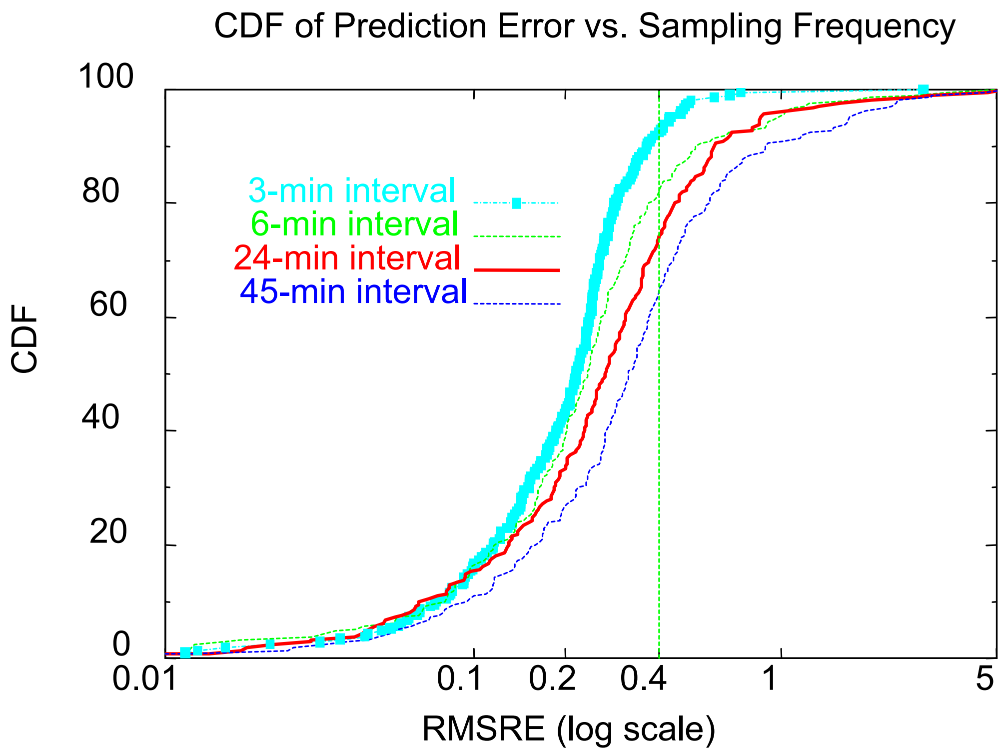

Outline

- TCP throughput prediction: problem statement
- Formula-Based (FB) prediction
  - A formula-based predictor
  - Types of FB prediction errors
  - Empirical evaluation
- History-Based (HB) prediction
  - Typical history-based predictors
  - Dealing with outliers and level shifts
  - Empirical evaluation
- Predictability factors
  - What makes some paths less predictable than others?

What makes throughput more predictable on some paths than on others?

- Factors examined:
  - Link utilization
  - Degree of statistical multiplexing
- Approach:
  - Analyze the Coefficient of Variation (CoV) of the marginal distribution of TCP throughput
    - CoV $\propto$ time series prediction error

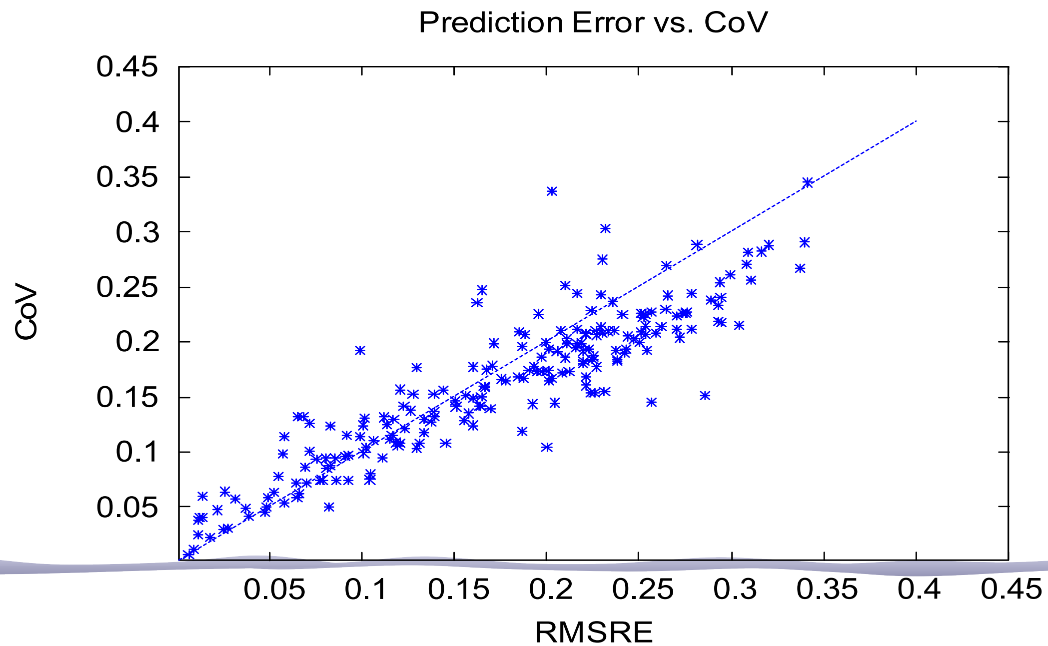

Impact of Load (congestion-limited flow)

- Model: Processor Sharing server (C) with Poisson session arrivals
  - Flow arrival rate: λ, avg flow size: θ
  - Offered load: $\rho = \frac{\lambda\theta}{C}$
  - Per-flow rate: $r(N) = C / N$
  - Distribution of # sessions:

$$\pi(N) = \rho^N (1-\rho)$$

$$CoV[r(N)] = \frac{(1-\rho)\log(1-\rho)^2 + \rho \cdot L(2,\rho)}{(1-\rho)\log(1-\rho)^2}$$

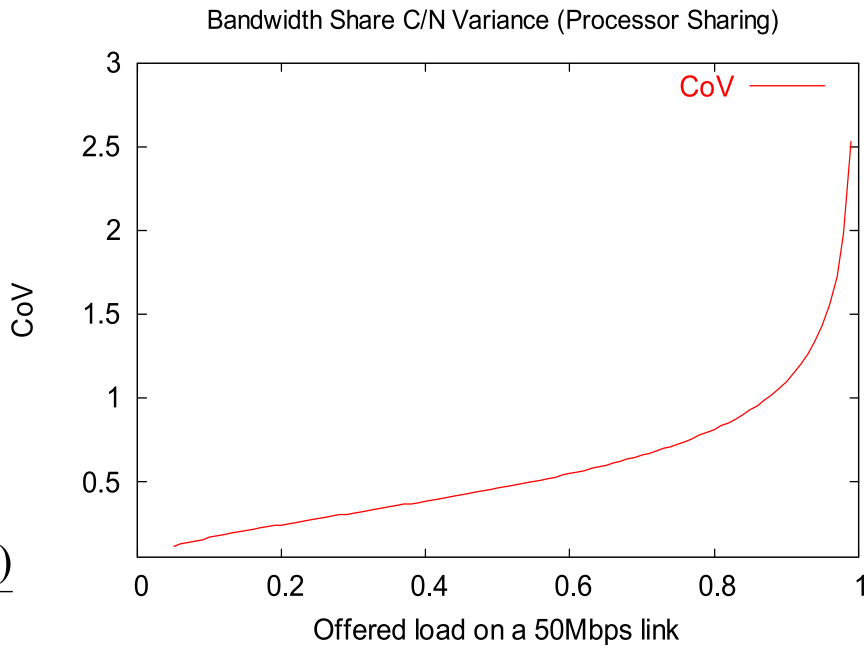

- CoV of per-session throughput increases with the offered load
- So, relative prediction error increases with offered load

Impact of Degree of Multiplexing

- Consider the avail-bw *A* at non-congested Processor Sharing server (*C*)
  - Traffic model: *N* homogenous flows with rate limit: *r*, flow arrival rate: *λ* (Poisson), avg flow size: *θ*

$$CoV[A] = CoV[C - Y] = \frac{1}{\sqrt{E[N]}} \frac{\rho \cdot C}{C(1-\rho)}$$

- Conclusion: provided that utilization remains constant, CoV of available bandwidth decreases as number of flows increases
  - So, we expect lower prediction error as number of flows increases

Conclusions

- FB prediction for congestion-limited TCP flows can cause major errors
  - Main reason: loss rate and RTT increase due to target flow
- HB prediction is much more accurate
  - Even with very simple predictors and sporadic previous transfers
- Path HB-predictability depends on load and degree of multiplexing at bottleneck link
  - Hardest-to-predict paths: heavily utilized bottleneck link, loaded with just a few flows