

Internet traffic constancy and predictability

Zhongtang Cai

Talk Outline

- Paper discussed: On the Constancy of Internet Path Properties. Yin Zhang, Nick Duffield, Vern Paxson, Scott Shenker.
- Motivation
- Three notions of constancy
 - Mathematical
 - Operational
 - Predictive
- Constancy of three Internet path properties
 - Packet loss
 - Packet delays
 - Throughput
- Conclusions

Motivation

- Interests in network measurement
 - Mathematical modeling
 - Operational procedures
 - Adaptive applications
- Measurements are most valuable when the relevant network properties exhibit *constancy*
 - Constancy: *holds steady and does not change*

Mathematical Constancy

- Mathematical Constancy
 - A dataset is *mathematically steady* if it can be described with a mathematical model.
 - Simplest form: IID – independent and identically distributed
 - *Key: finding the appropriate model*
- Examples
 - Mathematical constancy
 - Session arrivals are well described by a fix-rate Poisson process over time scales of 10 minutes to an hour [PF95]
 - Mathematical non-constancy
 - Session arrivals over larger time scales

Operational Constancy

- Operational constancy
 - A dataset is *operationally steady* if the quantities of interest remain within bounds considered operationally equivalent
 - *Key: whether an application cares about the changes*
- Examples
 - Operationally but not mathematically steady
 - Loss rate remained constant at 10% for 30 minutes and then abruptly changed to 10.1% for the next 30 minutes.

Predictive Constancy

- Predictive constancy
 - A dataset is *predictively steady* if past measurements allow one to reasonably predict future characteristics
 - *Key: how well changes can be tracked*
- Examples
 - Mathematically but not predictively steady
 - IID processes are generally impossible to predict well
 - Neither mathematically nor operationally steady, but highly predictable
 - E.g. RTT

Analysis Method

- Mathematical constancy
 - Identify change-points and partition a timeseries into change-free regions (CFR)
 - Test for IID within each CFR
- Operational constancy
 - Define operational categories based on requirements of real applications
- Predictive constancy
 - Evaluate the performance of commonly used estimators
 - Exponentially Weighted Moving Average (EWMA)
 - Moving Average (MA)
 - Moving Average with S-shaped Weights (SMA)

Predictive Constancy of Loss Rate

- Estimators

- MA, SMA, EWMA

MA Moving Average

$$s_{HtL} = \frac{\sum_{i=1}^M Y_{Ht-iL}}{M}, \quad M \geq 1$$

SMA S-shaped Moving Average


$$s_{HtL} = \frac{\sum_{i=1}^M w_i Y_{Ht-iL}}{\sum_{i=1}^M w_i}, \quad M \geq 1$$

EWMA Exponentially Weighted Moving Average

$$s_{HtL} = Y_{HtL} + \lambda(1-L)s_{Ht-1L} \quad @0, 1D.$$

Testing for Change-Points

- Identify a candidate change-point using CUSUM(CUMulative SUM)



- Use CUSUM and bootstrapping to detect changes(CP/Bootstrap)
 - Analyze ranks – resistant to the presence of outlier
 - Find a candidate change-point
 - Use bootstrap analysis

Testing for Change-Points - CP/RankOrder

- Analyze ranks – resistant to the presence of outlier

For $\{x_i\}_{i=1, 2, \dots, n}$

r_i : rank of x_i

- Find a candidate change-point

$$s_i = \sum_{j=1}^i r_j$$

$$\bar{s}_i = i \lfloor (n+1) \hat{\epsilon} \rfloor$$

$$s_i^{\text{oe}} = \left\| s_i - \bar{s}_i \right\|,$$

Candidate change - point at i , s.t.

$$s_i^{\text{oe}} > s_{i-1}^{\text{oe}}, \text{ where } 1 \leq i \leq n, i \neq 1$$

Testing for Change-Points - CP/RankOrder(Cont'd)

- Bootstrap analysis

A. Let $S_{\text{diff}} = S_{\text{max}} - S_{\text{min}}$, where

$$S_{\text{max}} = \max_{i=1, \dots, n} |S_i|$$

$$S_{\text{min}} = \min_{i=1, \dots, n} |S_i|$$

B. Generate bootstrap sample : $x_1^k, x_2^k, \dots, x_n^k, 1 \leq k \leq M$.

Sampling with replacement

C. Calculate $S_{\text{diff}}^k, 1 \leq k \leq M$

$$Y_k = \begin{cases} 1 & \text{if } S_{\text{diff}}^k < S_{\text{diff}} \\ 0 & \text{if } S_{\text{diff}}^k \geq S_{\text{diff}} \end{cases}$$

$$X = \sum_{k=1}^M Y_k$$

Change - point at i with confidence Level = $100 \frac{X}{M} \%$

Measurement Environment

- Two basic types of measurements
 - Poisson packet streams (for loss and delay)
 - Payload: 64 or 256 bytes; rate: 10 or 20 Hz; duration: 1 Hour.
 - Poisson intervals
 - Bi-directional measurements → RTT
 - TCP transfers (for throughput)
 - 1 MB transfer every minute for a 5-hour period
- Measurement infrastructure
 - NIMI: National Internet Measurement Infrastructure
 - Patterned after Paxson's Network Probe Daemon (NPD),
 - 35-50 hosts
 - ~75% in USA; the rest in 6 countries
 - Well-connected: mainly academic and research institute


Datasets Description

- Two main sets of data
 - Winter 1999-2000 (W_1)
 - Winter 2000-2001 (W_2)

Dataset	# NIMI sites	# packet traces	# packets	# thruput traces	# transfers
W_1	31	2,375	140M	58	16,900
W_2	49	1,602	113M	111	31,700
$W_1 + W_2$	49	3,977	253M	169	48,600

Individual Loss vs. Loss Episodes

- Traditional approach – look at individual losses [Bo93,Mu94,Pa99,YMKT99].
 - Correlation reported on time scales below 200-1000 ms
- Our approach – consider *loss episodes*
 - Loss episode: a series of consecutive packets that are lost
 - Loss episode process – the time series indicating when a loss episode occurs
 - Can be constructed by collapsing loss episodes and the non-lost packet that follows them into a single point.



Source of Correlation in the Loss Process

- Many traces become consistent with IID when we consider the loss episode process
- IID-independent identically distributed - **Box-Ljung test**

Box - Ljung statistic Q_k

$$Q_k = n \sum_{i=1}^k \frac{r_i^2}{n-i}, \text{ where } r_i \text{ is the autocorrelation}$$

$$r_i = \frac{\text{Cov}(x_{t+k}, x_t)}{\sqrt{\text{Var}(x_{t+k}) \text{Var}(x_t)}}$$

Under null hypothesis: x_t are independent Gaussian RV,

Q_k converges to χ^2 distribution.

If $Q_k \geq \chi^2_{1-\alpha}$ - quantile of χ^2 distribution, reject.

If $Q_k < \chi^2_{1-\alpha}$ - accept.

Source of Correlation in the Loss Process(Cont'd)

- Individual Loss vs. Loss Episodes

Traces consistent with IID	
Loss	Episode
27%	64%

- Loss process is not IID. Loss episode process can be viewed as IID

Poisson Nature of Loss Episodes within CFRs

- Independence of loss episodes within change-free regions (CFRs)


IID CFRs	IID traces
88%	64%

- Exponential distribution of interarrivals within change-free regions
 - 85% CFRs have exponential interarrivals

Loss episodes are well modeled as homogeneous Poisson process within change-free regions.

Poisson Nature of Loss Episodes within CFRs(Cont'd)

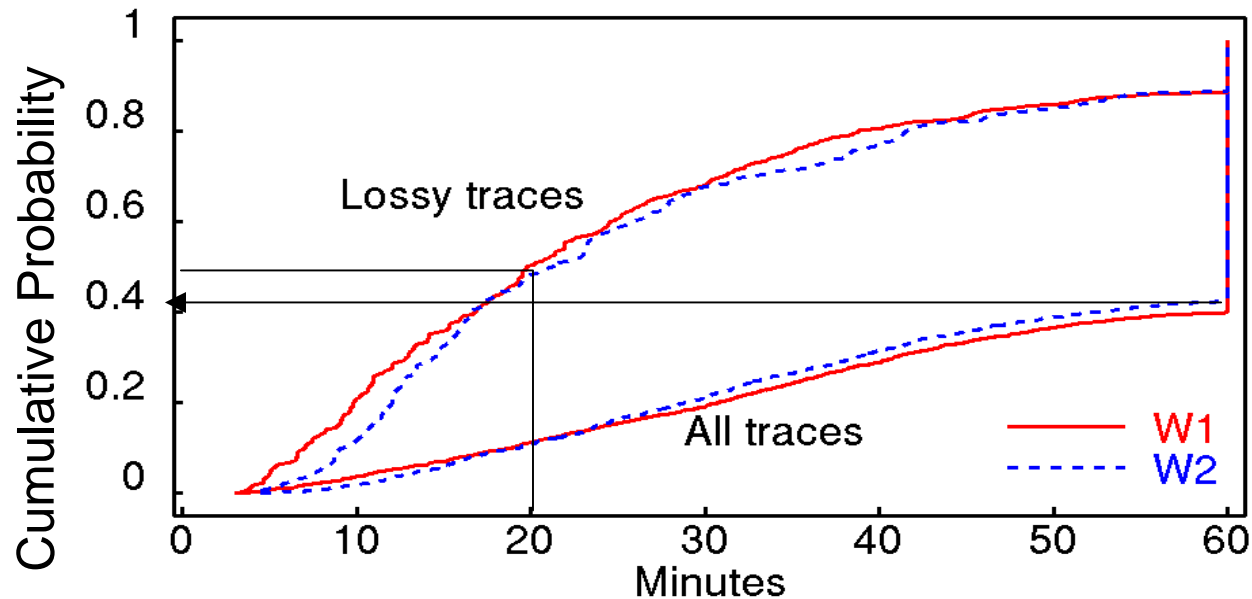
- Exponential distribution of interarrivals within change-free regions



- X: length of the loss-free periods(loss episode inter-arrival time)
- Y: CDF
- Argues strongly for Poisson loss episode arrivals

Mathematical Constancy of Loss Episode Process

Size of Largest Change Free Regions



- X: size of the largest CFR found for each trace
 - CFR: Change Free Regions. (Change-point test)
- “Lossy” traces are traces with overall loss rate over 1%
 - Only 50% with largest CFR > 20 mins
 - Higher loss rate makes the loss episode process less steady
- All traces: more than half of the traces are steady over full hour

Operational Constancy of Loss Rate

- Loss rate categories
 - 0-0.5%, 0.5-2%, 2-5%, 5-10%, 10-20%, 20+%
- Probabilities of observing a steady interval of **50 or more** minutes

Interval to calculate loss rate:	Type	Prob.
1 min	Episode	71%
	Loss	57%
10 sec	Episode	25%
	Loss	22%


■ .

Mathematical vs. Operational

- Categorize traces as “steady” or “not steady”
 - whether a trace has a 20-minute steady region

M: Mathematically steady

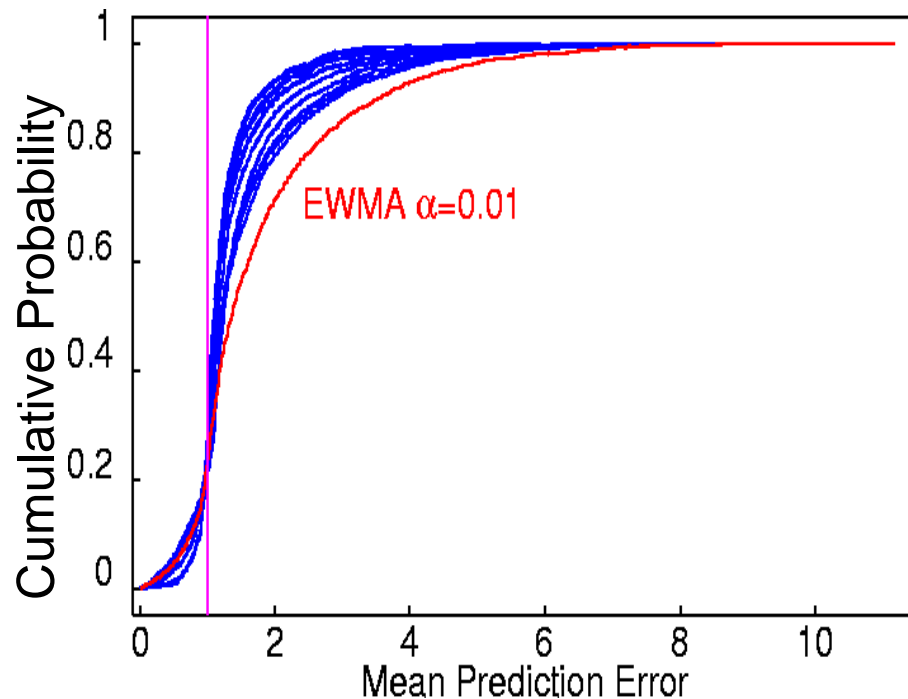
O: Operationally steady



Set	Interval	
	1 min	10 sec
M O	6-9%	11%
M O	6-15%	37-45%
M O	2-5%	0.1%
M O	74-83%	44-52%

Operational constancy of packet loss coincides with mathematical constancy on big time scales (e.g. 1 min), but not so well on medium time scales (e.g. 10 sec).

Predictive Constancy of Loss Rate



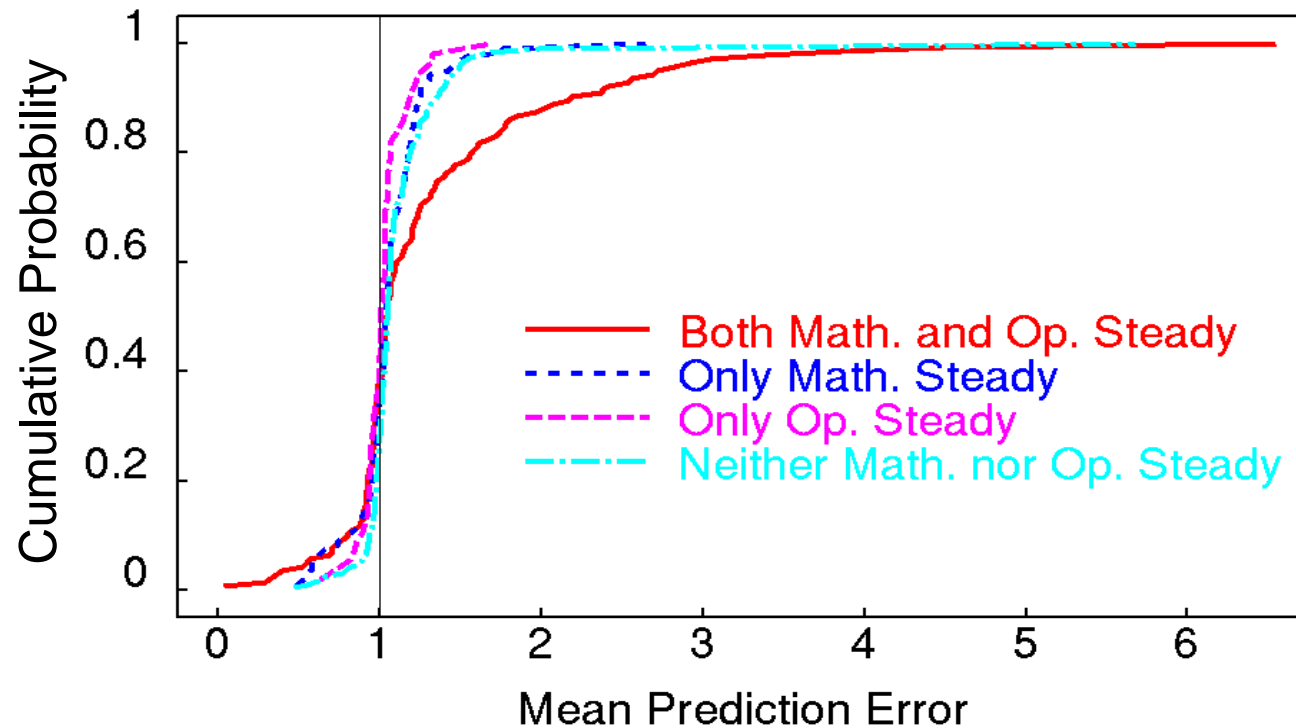
- What to predict?
 - The total time of next loss free run
- Estimators
 - EWMA, MA, SMA
- Mean prediction error
 - $E [| \log (\text{predicted} / \text{actual}) |]$

The parameters don't matter, nor does the averaging scheme.

Effects of Mathematical and Operational Constancy on Prediction

- How predictive constancy is related with mathematical constancy and operational constancy.
 - Aim only to understand the coarse grained relationship
 - Consider a trace mathematically steady if it has a maximum CFR of at least 20 mins.
 - Consider a trace operationally steady if it stays within a particular loss region for at least 20 mins

Effects of Mathematical and Operational Constancy on Prediction (cont'd)



Effects of Mathematical and Operational Constancy on Prediction (cont'd)

- Quality of the predictor is virtually unchanged if we have
 - MO
 - $M\bar{O}$
 - $\bar{M}\bar{O}$
- Prediction performance is the worst for traces that are both mathematically and operationally steady
 - Loss episode process resembles an IID process
 - no significant short-term variations
 - recent samples provide no help in predicting the next event

Delay Constancy

- Mathematical constancy
 - Delay “spikes”
 - A spike is identified when
 - $R' \geq \max\{ K \cdot R, 250\text{ms} \}$ ($K = 2$ or 4)
 - where
 - R' is the new RTT measurement;
 - R is the previous non-spike RTT measurement;
 - The spike episode process is well described as Poisson within CFRs
 - Body of RTT distribution
 - Good agreement (90-92%) with IID within CFRs

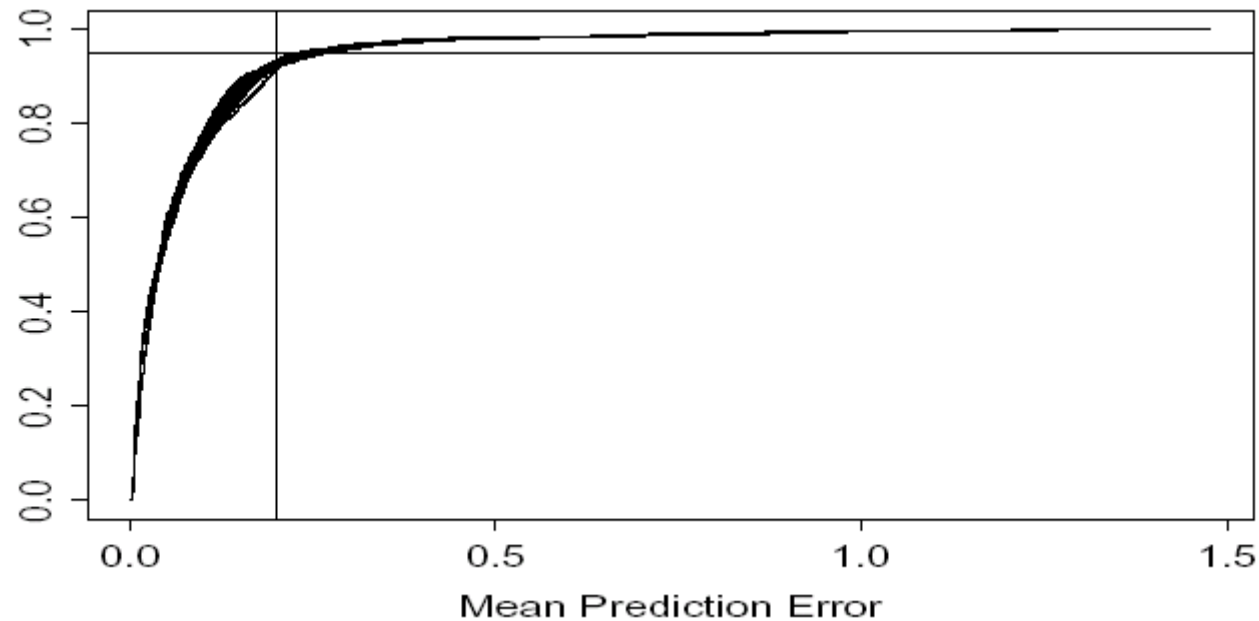
Delay Constancy (cont'd)

- Operational constancy
 - Operational categories
 - 0-0.1sec, 0.1-0.2sec, 0.2-0.3sec, 0.3-0.8sec, 0.8+sec
 - Based on ITU Recommendation G.114
 - No operational constancy
 - Over 50% traces have max steady regions under 10 min;
 - 80% are under 20 minutes

Delay Constancy (cont'd)

■ Predictive constancy

- All estimators perform similar
- Highly predictable in general (whether including RTT spikes or not)

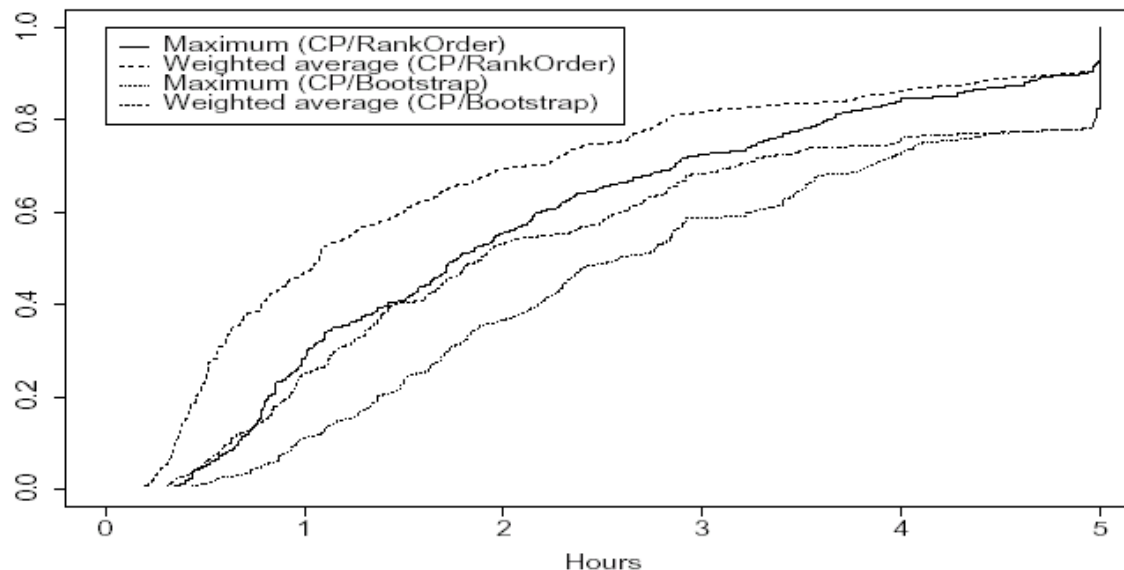


CDFs of the mean error for a large number of delay predictors.

Throughput Constancy

■ Mathematical constancy


- Apply change-point analysis to the mean of the series of per-minute throughput measurements.



CDF of maximum and weighted average CFRs for throughput achieved transferring 1 MB using TCP. X: length of maximum CFR or weighted average of length of CFR (in hours). Y: Cumulative distribution function.

Throughput Constancy

- Operational constancy
 - Categorize based on $p = \text{Max_Throughput} / \text{Min_Throughput}$.
 - Distribution of length of period in which bw stays in the region of $p < a$.



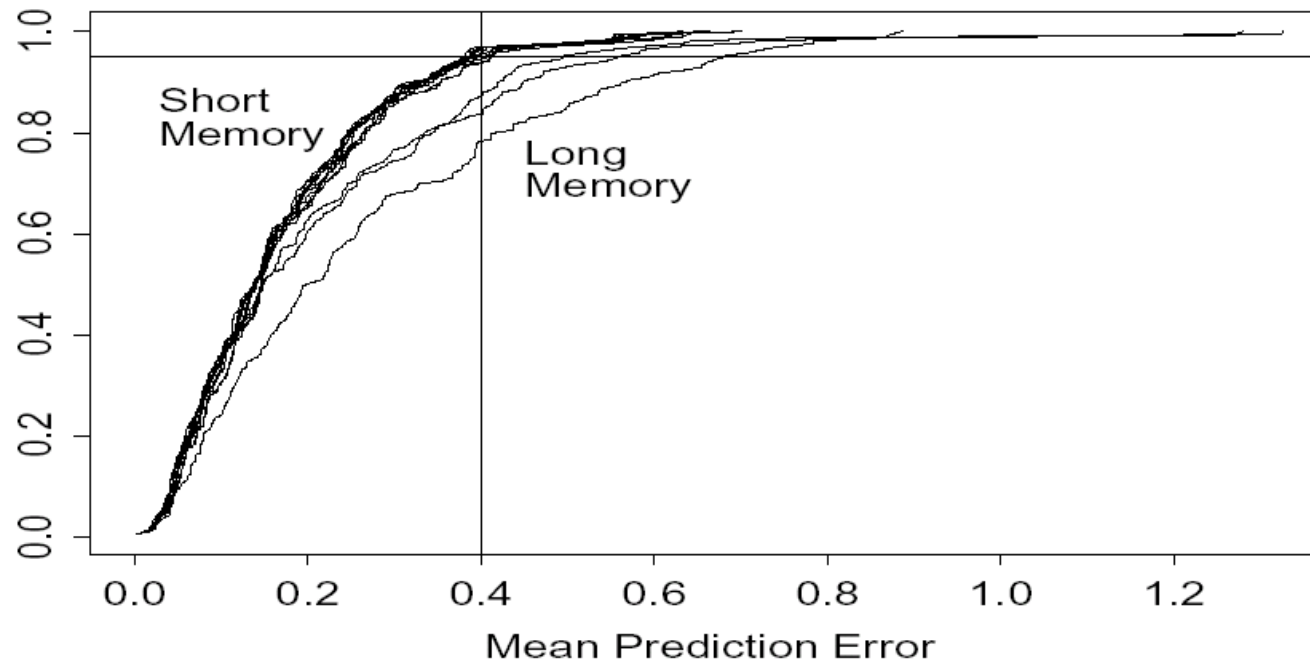
Throughput Constancy

- Mathematical Constancy vs. Operational Constancy
 - No simple relationship between mathematical constancy and operational constancy due to there is a wide range as $p = 1.2 \sim 10$.

p	$\hat{M}^{\wedge}O$	$M^{\wedge}O$	$\hat{M}O$	MO
1.2	53%	39%	2.4%	5.9%
10	3.6 %	1.2 %	51.5 %	43.8 %

Throughput Constancy

- Predictive constancy
 - All estimators perform very similar
 - Fairly good predictive constancy



Conclusions

- Three kinds of constancy:
Mathematical, Operational, Predictive
- Three key Internet path properties
 - IID works surprisingly well
 - It's important to find the appropriate model.
 - Different classes of predictors frequently used in networking produced very similar error levels
 - One can generally count on constancy on at least the time scales of minutes
 - This gives the time scales for caching path parameters