On the Constancy of Internet Path Properties

Yin Zhang
Nick Duffield
AT&T Labs – Research
{yzhang,duffield}@research.att.com

Vern Paxson
Scott Shenker
ACIRI
{vern,shenker}@aciri.org

ACM SIGCOMM Internet Measurement Workshop

November, 2001

Talk Outline

- Motivation
- Three notions of constancy
  - Mathematical
  - Operational
  - Predictive
- Constancy of three Internet path properties
  - Packet loss
  - Packet delays
  - Throughput
- Conclusions

Motivation

- Recent surge of interest 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*
  - We will also use the term *steady*, when use of “constancy” would prove grammatically awkward

Mathematical Constancy

- Mathematical Constancy
  - A dataset is *mathematically steady* if it can be described with a single time-invariant 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 10s of 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.
  - Mathematically but not operationally steady
    - Bimodal loss process with high degree of correlation

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 Methodology

- 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)

Testing for Change-Points

- Identify a candidate change-point using CUSUM

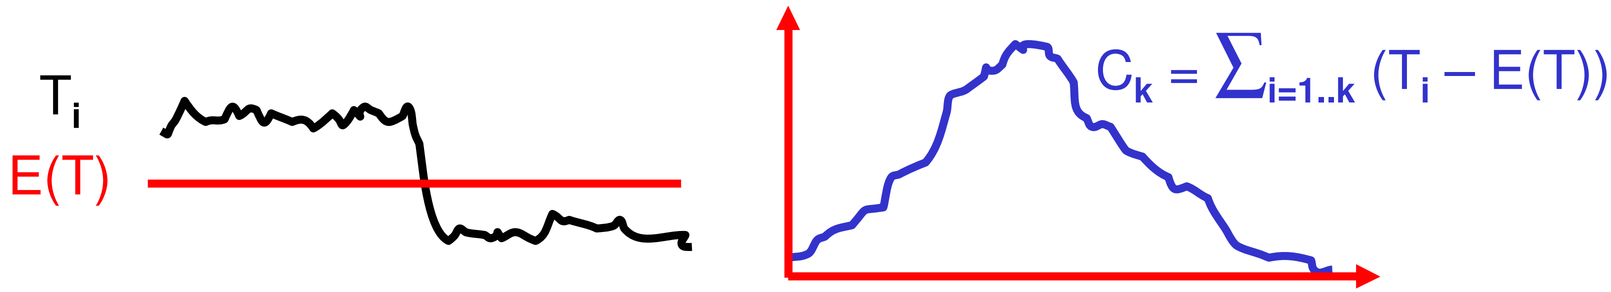

- Apply a statistical test to determine whether the change is significant
  - *CP/RankOrder*:
    - Based on Fligner-Policello Robust Rank-Order Test [SC88]
  - *CP/Bootstrap*:
    - Based on bootstrap analysis
- Binary segmentation for multiple change-points
  - Need to re-compute the significance levels

Measurement Methodology

- 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 → unbiased time averages [Wo82]
    - 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
    - 35-50 hosts
    - ~75% in USA; the rest in 6 countries
    - Well-connected: mainly academic and laboratory sites

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.

loss process 0 0 1 0 1 1 1 0 0 1 0 0 0

episode process 0 0 1 1 0 1 0 0

Source of Correlation in the Loss Process

- Many traces become consistent with IID when we consider the loss episode process

Time scale	Traces consistent with IID	
	Loss	Episode
Up to 0.5-1 sec	27%	64%
Up to 5-10 sec	25%	55%

Correlation in the loss process is often due to back-to-back losses, rather than intervals over which loss rates become elevated and "nearby" but not consecutive packets are lost.

Poisson Nature of Loss Episodes within CFRs

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

Time scale	IID CFRs	IID traces
Up to 0.5-1 sec	88%	64%
Up to 5-10 sec	86%	55%

- 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.

Mathematical Constancy of Loss Episode Process

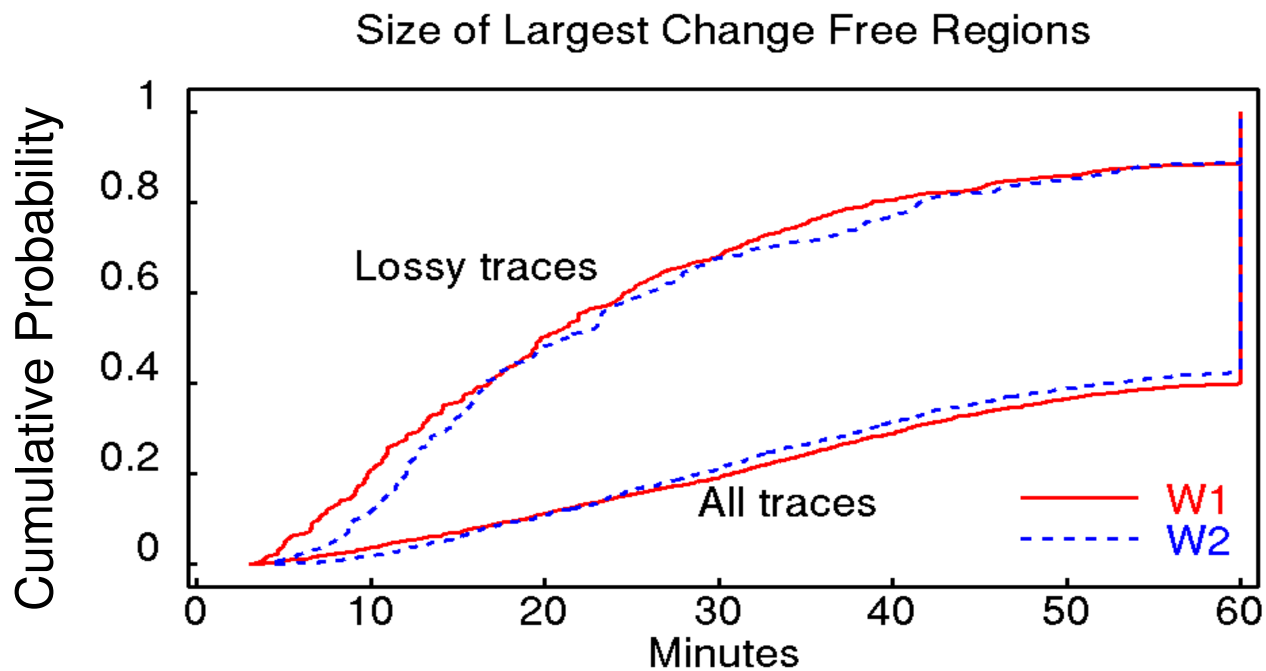

- Change-point test: *CP/RankOrder*
- "Lossy" traces are traces with overall loss rate over 1%

Higher loss rate makes the loss episode process less steady

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	Type	Prob.
1 min	Episode	71%
	Loss	57%
10 sec	Episode	25%
	Loss	22%

- There is little difference in the size of steady intervals of 50 or less minutes.

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

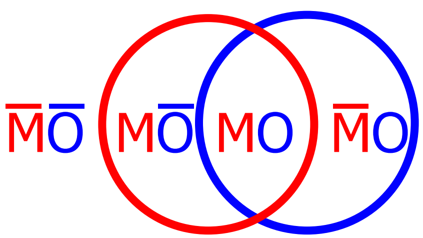

Set	Interval	
	1 min	10 sec
$\overline{M}\overline{O}$	6-9%	11%
$M\overline{O}$	6-15%	37-45%
$\overline{M}O$	2-5%	0.1%
MO	74-83%	44-52%

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

Predictive Constancy of Loss Rate

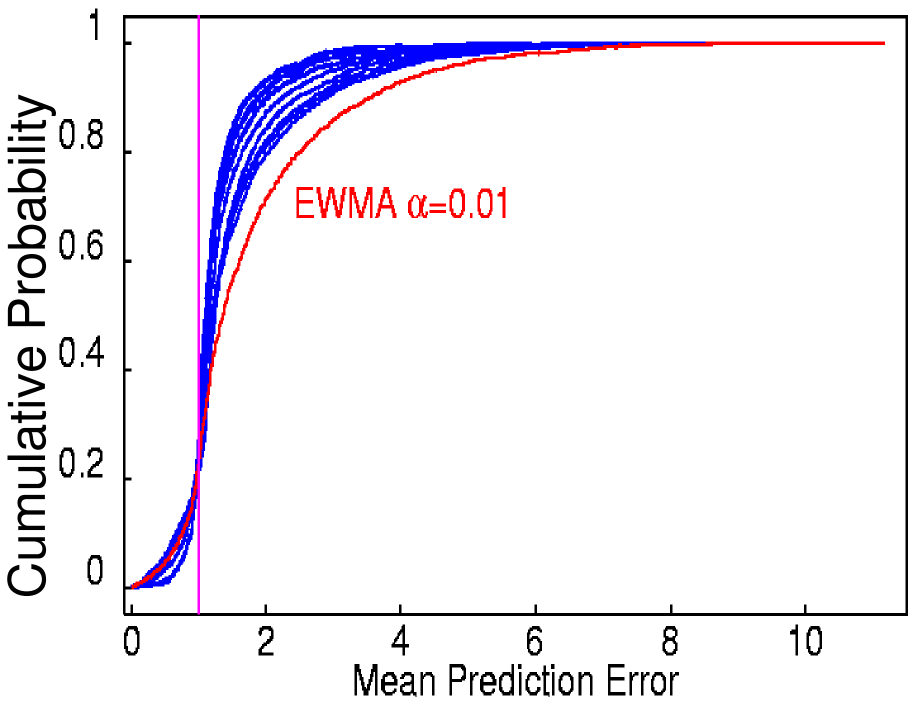

- What to predict?
  - The length of next loss free run
    - Used in TFRC [FHPW00]
- Estimators
  - EWMA, MA, SMA
- Mean prediction error

  E [| log (predicted / actual) |]

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

Effects of Mathematical and Operational Constancy on Prediction

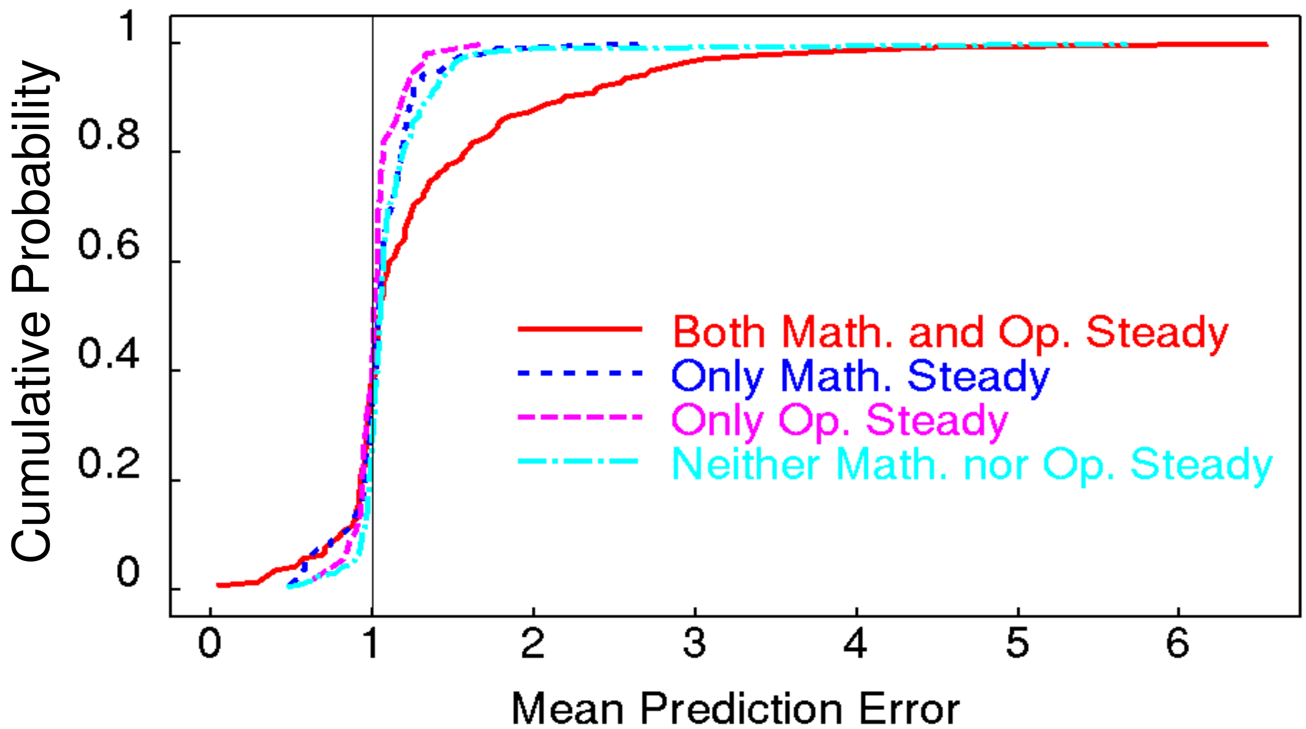

Prediction performance is the worst for traces that are both mathematically and operationally steady

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 (Median, IQR)
    - Overall, less steady than loss
    - 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
  - Not operationally steady
    - Over 50% traces have max steady regions under 10 min;
    - 80% are under 20 minutes
- Predictive constancy
  - All estimators perform similar
  - Highly predictable in general

Throughput Constancy

- Mathematical constancy
  - 90% of time in CFRs longer than 20 min
  - Good agreement (92%) with IID within CFRs
- Operational constancy
  - There is a wide range
- Predictive constancy
  - All estimators perform very similar
  - Estimators with long memory perform poorly

Conclusions

- Our work sheds light on the current degree of constancy found in 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
    - What really matters is whether you adapt, not how you adapt.
  - One can generally count on constancy on at least the time scales of minutes
    - This gives the time scales for caching path parameters
- We have developed a set of concepts and tools to understand different aspects of constancy
  - Applicable even when the traffic condition changes

Acknowledgments

- Andrew Adams
- Matt Mathis
- Jamshid Mahdavi
- Lee Breslau
- Mark Allman
- NIMI volunteers