

Evaluating Predictors of Congestion Collapse in Communication Networks

Christopher Dabrowski

Kevin Mills

National Institute of Standards & Technology



Paper Contributions

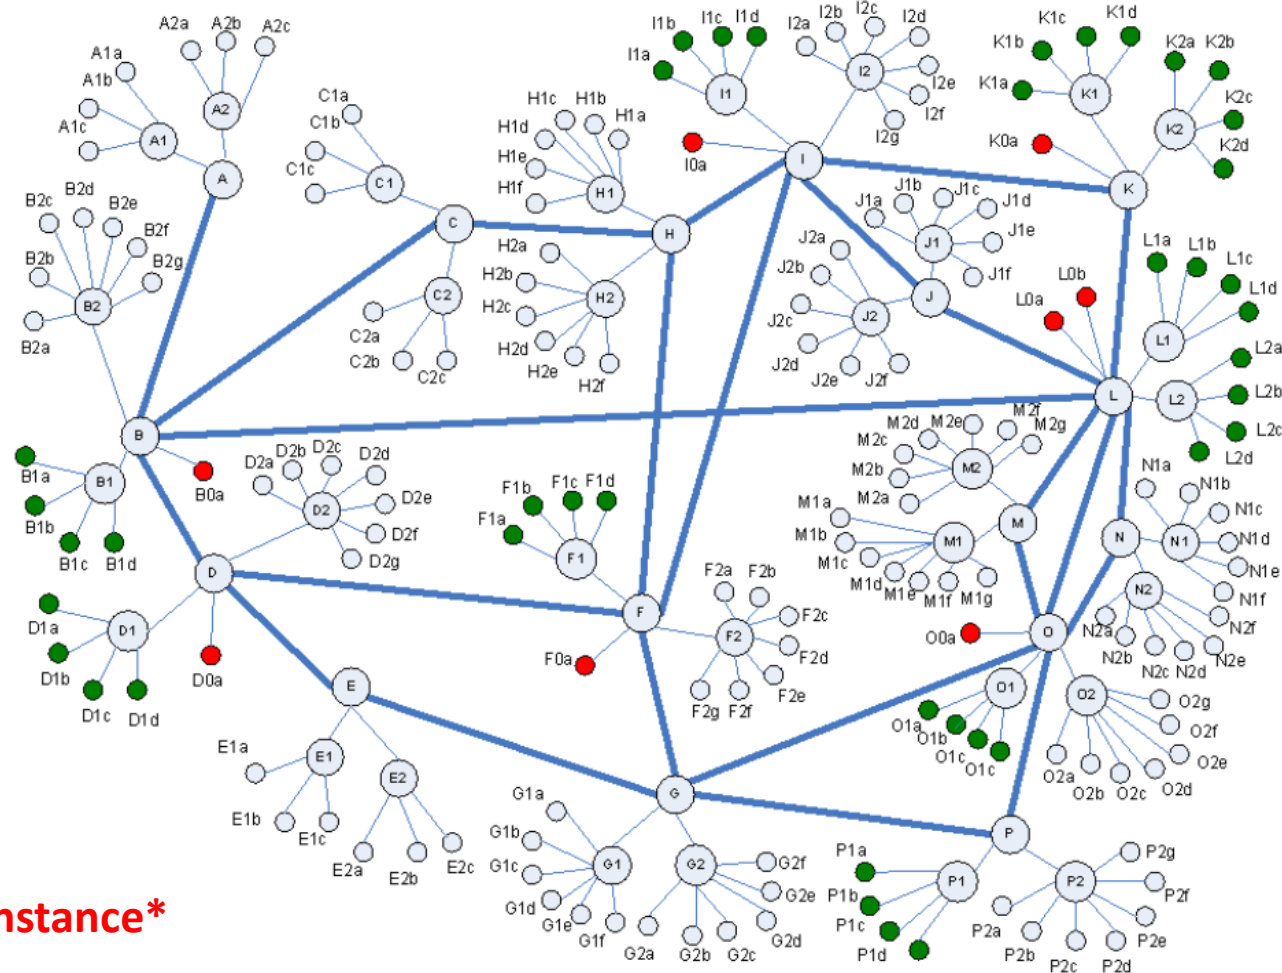
- Proposes *5 run-time predictors* of congestion collapse
- Defines *simulation-based method* to evaluate predictors
- Applies method to characterize and *compare predictors*
- Shows *simulation realism affects predictor performance*

Evaluated Congestion Predictors in Simulated Internet-Protocol Network

- 218 routers in 3 Tiers: 16 core, 32 PoP, and 170 edge
- Tier 4 (not shown): 257K sources/receivers
- Bush-Meyer buffer sizing
- TCP (or UDP) packet-injection procedures
- SPF routing - hops
- Tiered router speeds

Tier	Speed
Core	2S
PoP	S/4
Edge: Normal	S/4/10
Edge: Fast	2S/4/10
Edge: Very Fast	S/4

gray
green
red

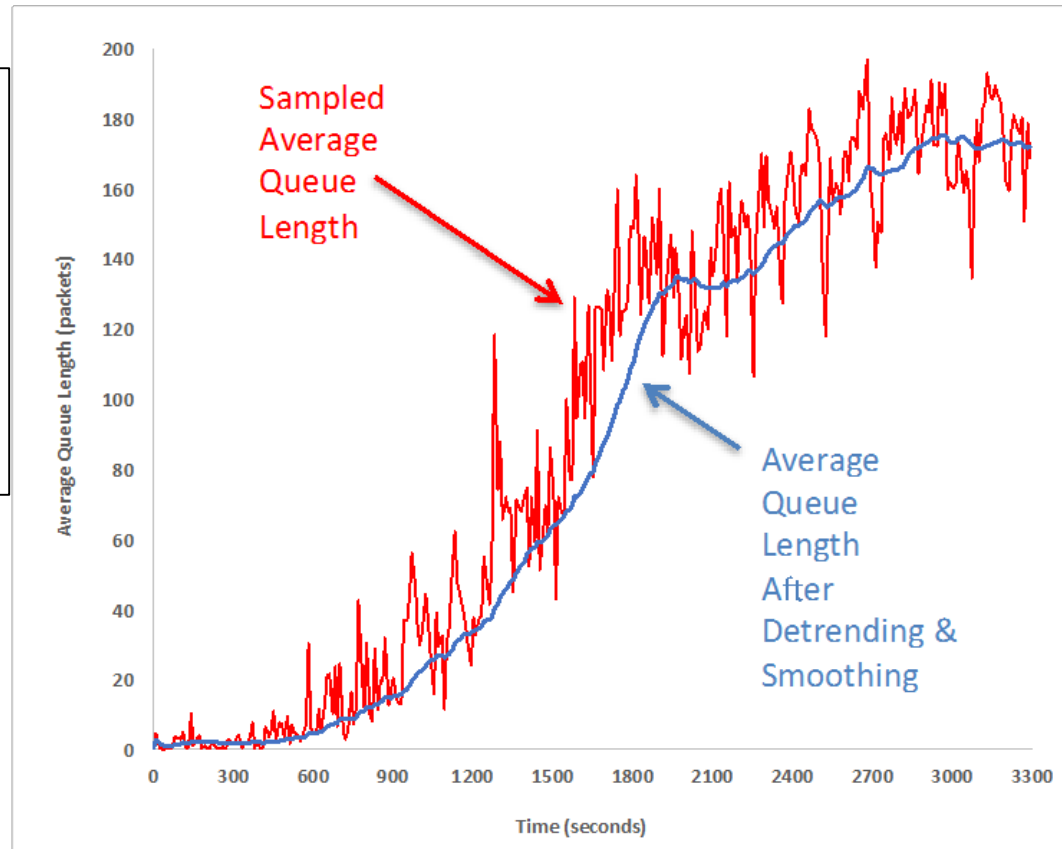


Each router runs predictor instance

Predictors Analyze Sampled, Smoothed Router Queue-Length Time Series

1. Queue sampled every 200 ms
2. Statistic (μ or v) computed over 50 samples (i.e., in each 10-s slot)
3. Queue under normal traffic subtracted
4. Samples smoothed using Nadaraya-Watson algorithm

Predictors analyze sliding window, W , of size 300 s, i.e., 30 10-s slots, with a 10-s advance, over smoothed queue-length statistic time series



Predictor Definition

GENERAL PROCEDURE FOR ANY PREDICTOR:

$$\sum_{i=1}^W [\mathbf{P}(Q_i) \geq P_{Threshold}] \geq P_{AlertLimit}$$

THREE PARAMETERS SPECIALIZE EACH PREDICTOR:

$(\mathbf{P}, P_{Threshold}, P_{AlertLimit}) \rightarrow$

- $\mathbf{P} \equiv$ predictor-specific **Function**
- $P_{Threshold} \equiv$ predictor-specific **Threshold**
- $P_{AlertLimit} \equiv$ predictor-specific **Alert Limit**

Predictors

Predictors require **11.9 us (TH) to 23.3 us (AC)** per sliding window (i.e., every 10 s) on a 3.4 GHz processor

AC – autocorrelation ($\alpha(1)$, 0.7, 21)

$$a(1) = \frac{E[(Q_i - \mu(Q))(Q_{i-1} - \mu(Q))]}{\sigma^2(Q)}$$

MATLAB `xcorr()`, `xcov()`, or statistical autocorrelation function

VR – variance (v , $\mu(v) + |3\sigma(v)|$: *steady load*, 21)

$$v = \frac{\sum_{i=1}^N |Q_i - \mu(Q)|^2}{(N - 1)}$$

Variance beyond $P_{Threshold}$ should occur by chance < 1% of the time (also tried $\mu(v) + |\sigma(v)|$)

TH – threshold (Q_j , $0.25(C - \mu(Q))$: *steady load*, 1)
C is buffer capacity and $\mu(Q)$ is mean queue length under steady load

GP – growth persistence ($Q_j \mid Q_j > Q_{j-1}$, $0.25(C - \mu(Q))$: *steady load*, 21)

GR – growth rate ((estimated $Q_{j-1} + s_{j-1}) + s_j$, $0.25(C - \mu(Q))$: *steady load*, 21))

$$s = \frac{\sum_{i=1}^W (i - (W + 1)/2)(Q_i - \mu(Q))}{\sum_{i=1}^W (i - (W + 1)/2)^2}$$

Predictor **implementation costs** in processor kilocycles/window;
↓ = better; best cost in **bold**: TH cheapest; AC most expensive

Predictor	Sample/ Detrend	Smooth	Compute Predictor	Decide Alert	Total↓
AC	0.98	25.86	51.90	0.41	79.14
VR	36.10	26.58	3.21	0.36	66.26
TH	0.98	25.86	0	13.61	40.45
GP	0.98	25.86	0	13.95	40.78
GR	0.98	25.86	1.19	12.93	40.95

Experiment

Predictor	Network Model	Traffic Scenario
Autocorrelation	Realistic TCP	Increasing Load
Variance		
Threshold	Realistic UDP	Steady Load
Growth		
Growth Rate	Abstract	

Evaluate each predictor under 6 conditions


Abstract Network Model :=

- flatten topology
- assign routers identical speeds
- remove sources & receivers
- infinite buffers
- packets injected at rate p , source routers chosen uniformly
- destination routers chosen uniformly
- SPF routing - hops

Each router raises **overload exception**, O , when queue length $\geq 70\% C$

Steady Load := p is 10 packets/s throughout time
Increasing Load := p increases by 10 packet/s every 10 s after warm up

Time Line :=



TCP & UDP sources inject flows (i.e., packet streams) with **flow injection probability** := $p/s \times f$, where s is # sources and f is mean flow size (350 packets here)

A := predictor alert; O := overload exception; t := time

true positive (tp) := $A_t < O_t$

true negative (tn) := $\neg A \wedge \neg O$

false positive (fp) := $A \wedge \neg O$

false negative (fn) := $\neg A \wedge O$

late positive (lp) := $O_t < A_t$

Performance

Metrics

measured for each router; then averaged

accuracy := $(tp + tn) /$

$(tp + tn + fp + fn + lp)$

latency := $O_t - A_t \mid tp$

persistence := $\frac{\int_{A_t}^{O_t} A}{(O_t - A_t)} \mid tp$

Successful Outcomes under Increasing Load

True-Positive & True-Negative Rates

	TPR ↑			TNR ↑			TPR + TNR		
	TCP	UDP	ABS	TCP	UDP	ABS	TCP	UDP	ABS
AC	0.913	0.963	0.131	0	0	0.805	0.913	0.963	0.936
VR	0.862	0.963	0.133	0	0	0	0.862	0.963	0.133
TH	0.723	0.75	0.128	0.254	0.25	0.859	0.977	1	0.986
GP	0.634	0.74	0	0.259	0.26	0.853	0.894	1	0.853
GR	0.686	0.75	0	0.254	0.25	0.853	0.94	1	0.853

- ~85% of routers in **Abstract Network** do not reach overload, as congestion spreads outward slowly from routers with high centrality (**unrealistic**)
- **Threshold** predictor performs **best**
- **Variance** predictor performs **worst**, no true negatives in Abstract Network

Mean Warning Times & Persistence

	Mean Latency (min)			Mean Persistence ↑		
	TCP	UDP	ABS	TCP	UDP	ABS
AC	10.73	22.89	17.52	1	1	0.992
VR	10.48	21.72	24.33	0.994	1	0.991
TH	7.67	15.77	0.58	1	1	1
GP	4.61	11.86	—	0.972	0.99	—
GR	5.63	13.53	—	1	1	—

All predictors quite **persistent**

- Under realistic network models, **average warning times** from ~ 5 to 23 mins
- **Variance** and **Autocorrelation** gave **most warning time** because the other predictors transition more slowly to alerts (guarded by $0.25(C - \mu(Q))$: *steady load*)
- **TCP Network** yielded **least warning time** because congestion-control procedures slow the rate of increasing congestion
- **TCP Network warning time exceeded 15 mins** for **12%** (GP) to **23%** (VR) of routers

Erroneous Outcomes under Increasing Load

Error Rates ↓ is better

	Error Rate (1 – Con)			Mean Error Rate
	TCP	UDP	ABS	
AC	0.087	0.037	0.074	0.06
VR	0.138	0.037	0.867	0.347
TH	0.023	0	0.014	0.012
GP	0.106	0	0.147	0.084
GR	0.06	0	0.147	0.069

- **Threshold** predictor performs **best** (1.2% error)
- **Variance** predictor performance **unacceptable**
- **Other** predictor error rates **okay** (< 10% error)

Error Rates Decomposed: **False-Positive Rate;**
False-Negative Rate; Late-Positive Rate

	FPR ↓			FNR ↓			LPR ↓		
	TCP	UDP	ABS	TCP	UDP	ABS	TCP	UDP	ABS
AC	0.087	0.037	0.058	0	0	0	0	0	0.016
VR	0.12	0.037	0.85	0	0	0	0.018	0	0.017
TH	0.023	0	0	0	0	0	0	0	0.014
GP	0.023	0	0	0	0	0	0.083	0	0.147
GR	0.023	0	0	0	0	0	0.037	0	0.147

- **No** predictors made **false negatives**
- **Autocorrelation** and **Variance** predictors mainly **false-positive errors**
- **Growth Persistence** and **Growth Rate** predictors mainly **late-positive errors**
- **False-positive rate** (2.3%) for **Threshold**, **Growth Persistence**, and **Growth Rate** predictors encompassed **5 edge routers** that had reached 64% of C and **would** have reached **overload** given more time

False-Positive Rates under Steady Load

	FPR ↓		
	TCP	UDP	ABS
AC	0.995	0.927	0.064
VR	0.362	0.45	0.949
TH	0	0	0
GP	0	0	0
GR	0	0	0

- Three predictors gave no false positives
- For realistic network models, Autocorrelation predictor issued predominately false positives, as queue time series highly self-similar
- Variance predictor issued too many false positives in all network models
- Adding guard condition – $0.25(C - \mu(Q))$: *steady load*) – to Autocorrelation and Variance predictors eliminated false positives, but at the cost of reduced warning times (results not shown)

Predictor Accuracy across Load Scenarios

	Accuracy ↑		
	TCP	UDP	ABS
AC	0.459	0.518	0.936
VR	0.75	0.757	0.092
TH	0.989	1	0.993
GP	0.947	1	0.927
GR	0.97	1	0.993

- Three predictors perfectly accurate in UDP Network
- For TCP Network, Threshold predictor most accurate
- Threshold predictor gives 3-4 mins more warning time than Growth Persistence and Growth Rate predictors (which must await alerts in 21/30 slots)
- Autocorrelation and Variance predictors unreliable
- Threshold predictor most suited to investigate further in emulated and real networks
- Abstract Network models should not be used to evaluate congestion predictors

Future Work

- Verify predictor results in emulated and real networks
- Explore parameter space of all predictors to determine relative influence of parameters and to select optimal parameter values
- Investigate predictor performance in more complex traffic scenarios, e.g., cycles of increasing and decreasing load
- Define and investigate additional predictors

For more research results, see:

<https://www.nist.gov/programs-projects/measurement-science-complex-information-systems>