UiO : University of Oslo

# Progress on Practical Shared Bottleneck Detection for Coupled Congestion Control

David Hayes (UiO)
{ Simone Ferlin (SRL) and Michael Welzl (UiO) }

R / T E
REDUCING INTERNET TRANSPORT LATENCY

SEVENTH FRAMEWORK PROGRAMME

# Background

## Problem

- Flows traversing different paths through a network may share a common congested link — a bottleneck
- Detecting which flows share a bottleneck and coupling their congestion control can provide performance advantages.

## SBD design objectives

- Reliable
- Practical (not CPU nor network intensive)
- Small numbers of bottlenecks ( $< 10$)
- Timely stable bottleneck detection ( $< 10\,s$)

# Shared Bottleneck Detection

## What does it rely on?

- flows that share a bottleneck are similar in a measurable way

## Why is it hard?

- delay and loss measurements include “noise” from rest of the path
- delay and loss at the bottleneck is noisy – each packet sees a different queueing delay
- different path delays cause time correlations to be lost or degraded at the measurement point

# Classic cross correlation techniques

## Pairwise flow cross correlation of delay samples

- delay signal is noisy
  - filter
- delay distribution is often skewed
  - sophisticated filter
- different path delays
  - incrementally shift and cross correlate to find lag of maximum correlation.

# The delay signal

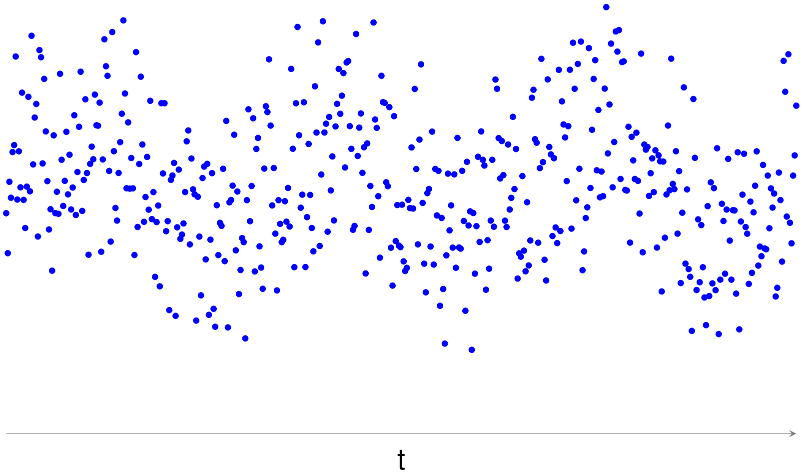

# Dealing with the signal noise

- Remove half the noise from other links by using OWD instead of RTT
- Only using difference statistics
  - removes queueing delay estimate errors due to inaccurate estimate of $OWD_{min}$
- Mitigate lag and sample noise by:
  - relatively large statistic gathering periods
  - relax thresholds (no need to distinguish between 1000 bottlenecks)
  - use multiple measures

# Summary statistics

- Skewness in OWD
  - an estimate using 2 counters
- Variance in OWD
  - estimated using PDV (RFC 5481)
- Key frequency of OWD at the bottleneck link
  - estimated based on significant mean crossings

# Grouping overview

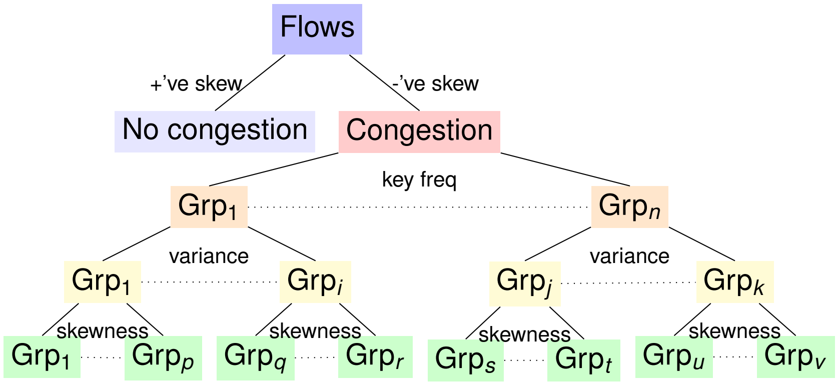

# Simulation tests

## Objectives

- Test with a known “ground truth”
- Simulations can allow us to look at worse than real scenarios.

But, real network experiments are in progress will be discussed before the end.

# Example NS2 Simulation

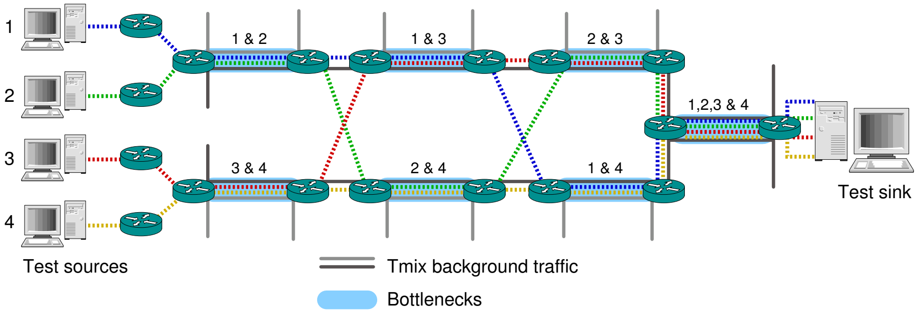

# Notes on results

- Decisions made every 300 ms, but based on 6 – 15 s statistics.
- Decision "points" are large for legibility, but it can tend to magnify errors.
- Results illustrates what can and can't be done.

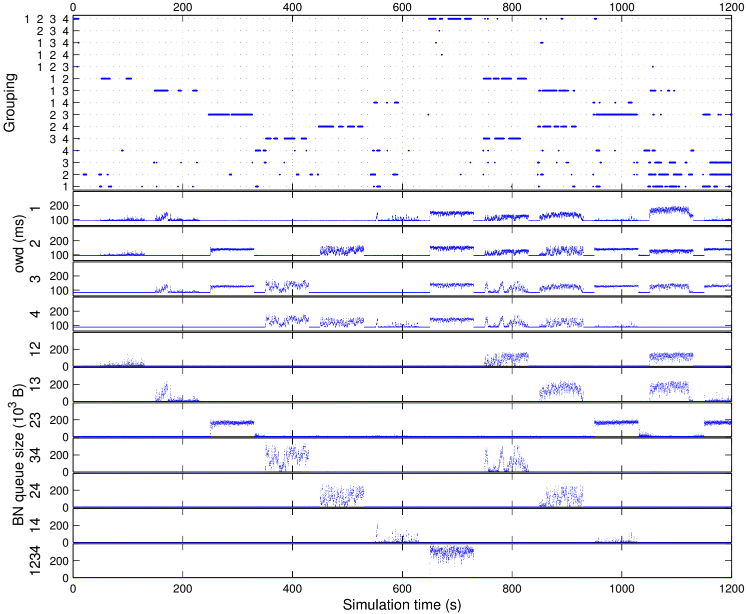


# Simulation results

Different queue sizes, link owd std 2.5.
Note bottleneck queue sizes have been subsampled (1:350) and OWDs (1:20)

# Real network experiments

- Bottleneck “ground truth” cannot be known with 100 % certainty.
  - Find thinnest link using STAB
  - Load thinnest link with distant internet sources to create known bottleneck
- What are we testing?
  - Robustness in unpredictable “real” environments

# Real network experiments (in progress)

Work with Simone Ferlin

[ simula . research laboratory ]

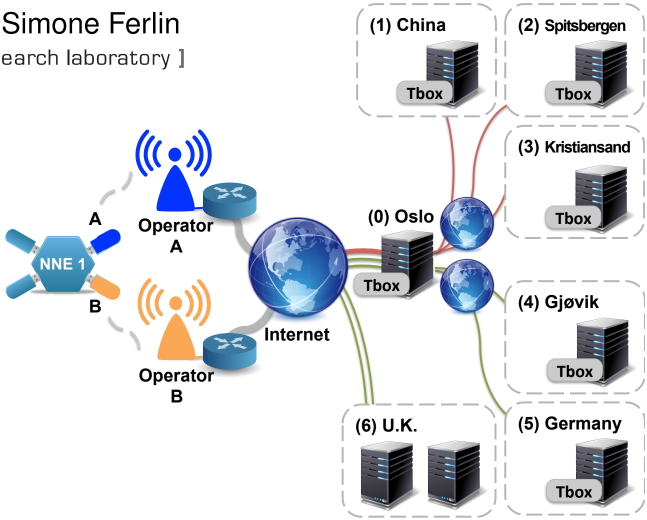

# Working with Coupled Congestion Control

- Summary statistics are gathered at the receivers
- Shared bottlenecks to a receiver
  - Receiver does grouping and sends information to senders
- Shared bottlenecks from a sender:
  - Receivers send summary statistics for grouping at sender.
- Can provide the necessary information for a future multi-sender multi-receiver coupled congestion control.

# Conclusions and further work

## Finalising this stage

- Finish real network experiments
- Paper submission soon (LCN)
- Draft (referring to paper)
- Quantitative results of % correct grouping
    - simulation based where “ground truth” is known
    - bottleneck definition based on queue empty rate or avg. queue size
    - extended version in journal

## Next steps

- Protocol for sender/receiver information exchange
- Integration with coupled congestion control
    - time scales of detection
    - dealing with SBD errors
    - oscillating bottlenecks

UiO : University of Oslo

# Extra slides

# Time domain summary statistics

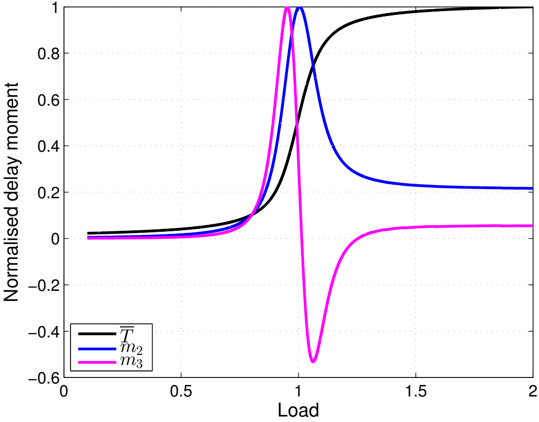


Mean, variance ($m_2$), skewness ($m_3$)

# Practical estimation of skewness

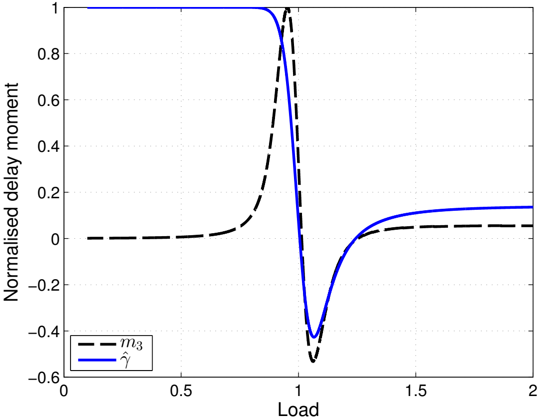

# Practical estimation of variance

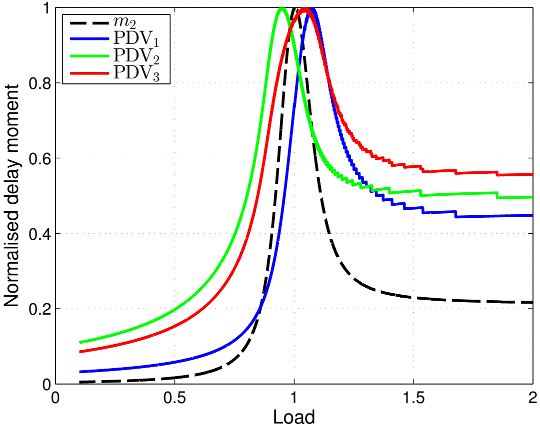

# Practical estimation of key frequency $\hat{f}$