UiO : University of Oslo

Shared Bottleneck Detection for Coupled Congestion Control for RTP Media

(draft-hayes-rmcat-sbd-01)

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

[simula . research laboratory]

REDUCING INTERNET TRANSPORT LATENCY

SEVENTH FRAMEWORK PROGRAMME

3 Host, 3 Links, 3 potential bottlenecks

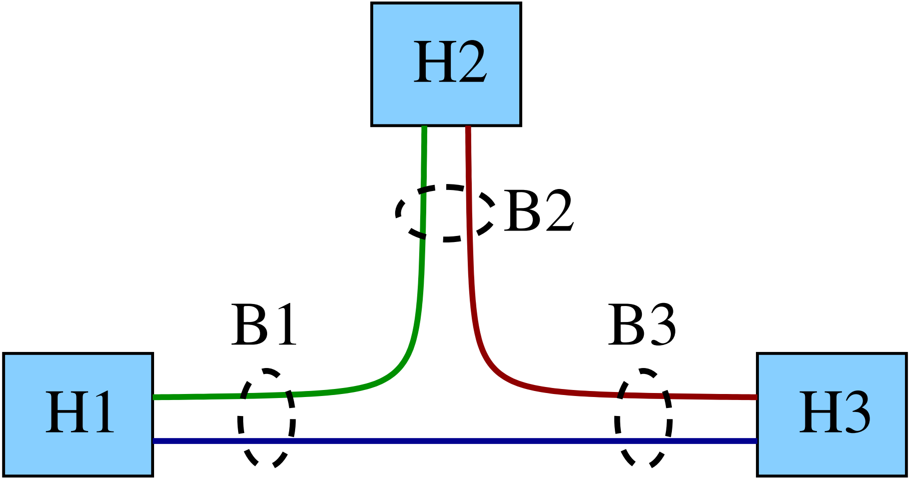

Aims

Provide a Shared Bottleneck Detection (SBD) mechanism that:

- can potentially handle sender and receiver grouped bottlenecks
  - Initially draft will outline it in terms of sender oriented grouped bottlenecks.
- requires minimal feedback
- requires no changes to the network or particular congestion control behaviour
- can handle future advances
  - e.g. changes in networks and new signals

Sender–Receiver Communication

Initialisation

- list of summary statistics to be gathered
- define interval T, and numbers of intervals N & M

Regular Communication

- Summary statistics every T
- eventually also relevant receiver detected SBNs

Summary Statistics

Why summary statistics?

- To limit feedback from receivers
- To deal with noise
- To deal with lag

Statistics Used

- variance – Packet Delay Variation (PDV)
- skewness (skew_est)
- oscillation (freq_est)

Skew_est

$$\text{skew\_est} = \frac{\text{sum\_T(OWD} < \text{E\_N(E\_T(OWD)))} - \text{sum\_T(OWD} > \text{E\_N(E\_T(OWD)))}}{\text{num\_T(OWD)}}$$

where

if (OWD < E_N(E_T(OWD)) 1 else 0

if (OWD > E_N(E_T(OWD)) 1 else 0

for grouping average over N intervals:

E_N(skew_est) = sum_N(skew_est) / N

Skew_est

long term mean over N intervals

skew_est =

$$\frac{\text{sum\_T}(\text{OWD} < \text{E\_N}(\text{E\_T}(\text{OWD}))) - \text{sum\_T}(\text{OWD} > \text{E\_N}(\text{E\_T}(\text{OWD})))}{\text{num\_T}(\text{OWD})}$$

where

if (OWD < E_N(E_T(OWD)) 1 else 0

if (OWD > E_N(E_T(OWD)) 1 else 0

for grouping average over N intervals:

E_N(skew_est) = sum_N(skew_est) / N

Skew_est

or long term mean over M < N intervals

skew_est =

$$\frac{\text{sum\_T}(\text{OWD} < \text{E\_M}(\text{E\_T}(\text{OWD}))) - \text{sum\_T}(\text{OWD} > \text{E\_M}(\text{E\_T}(\text{OWD})))}{\text{num\_T}(\text{OWD})}$$

where

if (OWD < E_N(E_T(OWD)) 1 else 0

if (OWD > E_N(E_T(OWD)) 1 else 0

for grouping average over N intervals:

E_N(skew_est) = sum_N(skew_est) / N

Skew_est

test each OWD sample as it arrives

$$\text{skew\_est} = \frac{\text{sum\_T}(\text{OWD} < \text{E\_N}(\text{E\_T}(\text{OWD}))) - \text{sum\_T}(\text{OWD} > \text{E\_N}(\text{E\_T}(\text{OWD})))}{\text{num\_T}(\text{OWD})}$$

where

if (OWD < E_N(E_T(OWD)) 1 else 0

if (OWD > E_N(E_T(OWD)) 1 else 0

for grouping average over N intervals:

E_N(skew_est) = sum_N(skew_est) / N

Skew_est

keep running sum of test results in T

skew_est =

$$\frac{\text{sum\_T}(\text{OWD} < \text{E\_N}(\text{E\_T}(\text{OWD}))) - \text{sum\_T}(\text{OWD} > \text{E\_N}(\text{E\_T}(\text{OWD})))}{\text{num\_T}(\text{OWD})}$$

where

if (OWD < E_N(E_T(OWD)) 1 else 0

if (OWD > E_N(E_T(OWD)) 1 else 0

for grouping average over N intervals:

E_N(skew_est) = sum_N(skew_est) / N

Skew_est

every T, normalise difference of sums

$$\text{skew\_est} = \frac{\text{sum\_T}(\text{OWD} < \text{E\_N}(\text{E\_T}(\text{OWD}))) - \text{sum\_T}(\text{OWD} > \text{E\_N}(\text{E\_T}(\text{OWD})))}{\text{num\_T}(\text{OWD})}$$

where

if (OWD < E_N(E_T(OWD)) 1 else 0

if (OWD > E_N(E_T(OWD)) 1 else 0

for grouping average over N intervals:

E_N(skew_est) = sum_N(skew_est) / N

PDV

PDV_max = (max_T(OWD) - E_T(OWD))
or
PDV_min = (min_T(OWD) - E_T(OWD))

for grouping average over N intervals:

E_N(PDV) = sum_N(PDV) / N

Freq_est

freq_est = number_of_crossings / N

where
a *crossing* is E_T(OWD) that crosses the E_N(E_T(OWD)) extending a further p_v * E_N(PDV). (incrementally calculated every T)

Freq_est

freq_est = number_of_crossings / N

where
a *crossing* is E_T(OWD) that crosses the E_N(E_T(OWD)) extending a further p_v * E_N(PDV). (incrementally calculated every T)

Freq_est

freq_est = number_of_crossings / N

where
a *crossing* is E_T(OWD) that crosses the E_M(E_T(OWD)) extending a further p_v * E_N(PDV). (incrementally calculated every T)

Freq_est

freq_est = number_of_crossings / N

where
a *crossing* is E_T(OWD) that crosses the E_N(E_T(OWD)) extending a further p_v * E_N(PDV). (incrementally calculated every T)

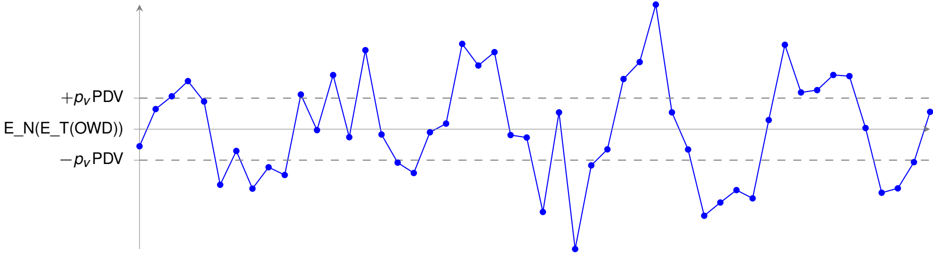

Freq_est

freq_est = number_of_crossings / N

where
a *crossing* is E_T(OWD) that crosses the E_N(E_T(OWD)) extending a further p_v * E_N(PDV). (incrementally calculated every T)

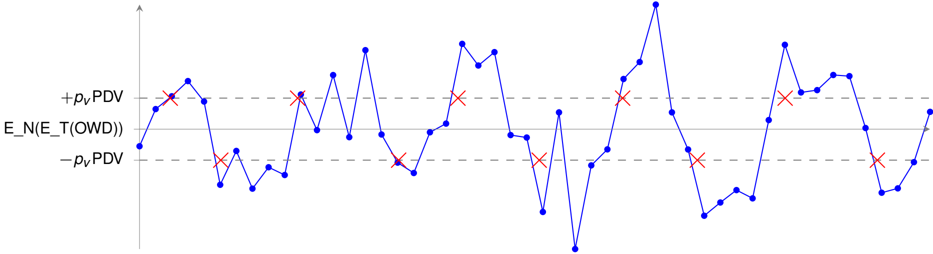

- freq_est$= 10/50 = 0.2$

Flow grouping

- Simple example algorithm given.
  - sufficient for RMCAT
  - not required to be common
- Basic algorithm:
  - Group flows experiencing congestion (based on skew_est)
    - are summary statistics strong enough to be grouped?
    - use of hysteresis
  - divide based on freq_est
  - divide based on PDV
  - divide based on skew_est

Some example results using the algorithm

Simulation results

- Skew_est
  - relation to congestion
  - hysteresis
- Varying path delays

Network results

- Nornet (`https://www.nntb.no/`)
- Same SBD parameter settings as simulations

Simulation setup

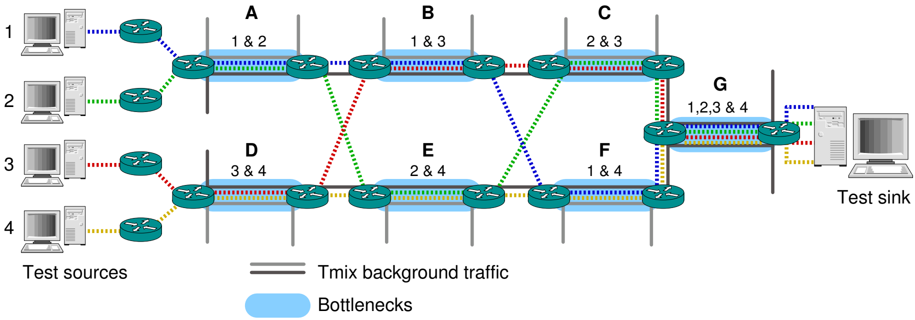

Simulation setup

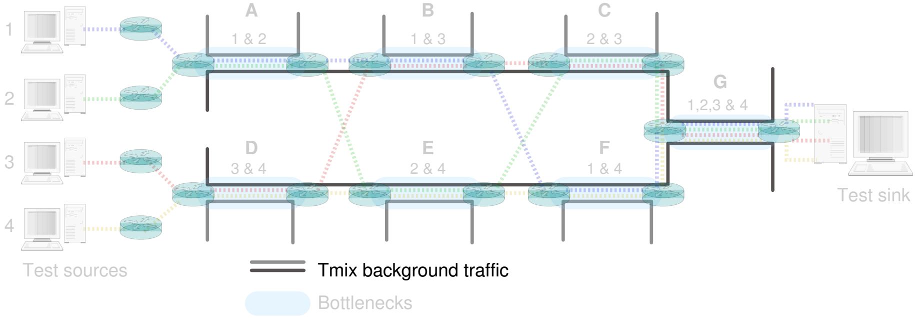

- Background traffic based on real traffic traces
  - > 90%

Simulation setup

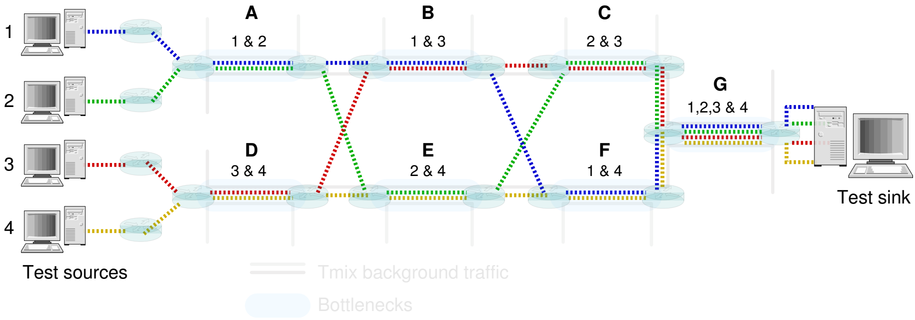

- Background traffic based on real traffic traces
  - > 90%
- Flows 1 & 2 send at twice the rate of 3 & 4

Simulation setup

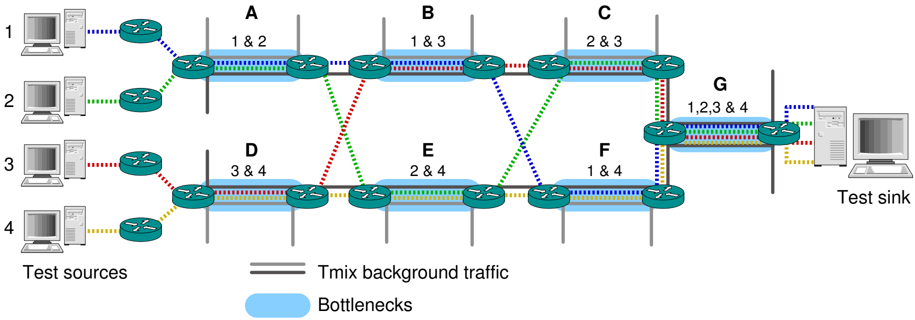

- Background traffic based on real traffic traces
  - > 90%
- Flows 1 & 2 send at twice the rate of 3 & 4
- Various combinations of bottlenecks activated

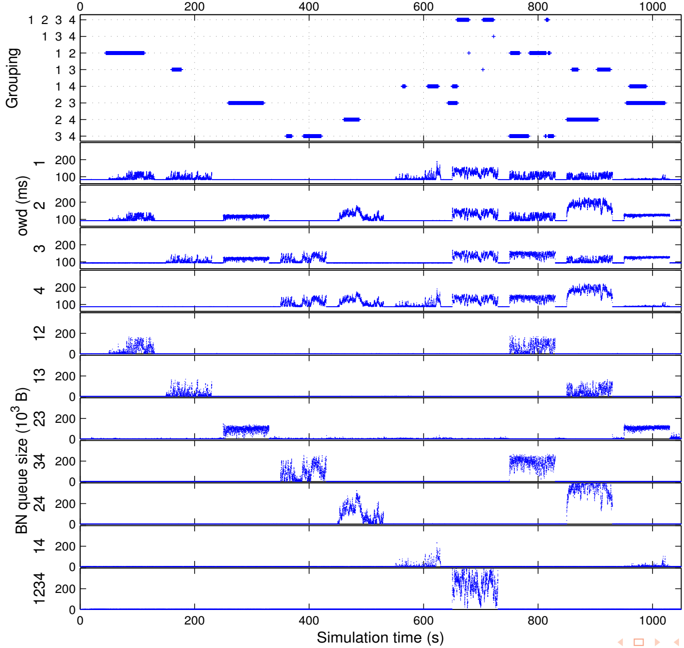

Simulation

(Queue sizes subsampled (1:200) and OWDs (1:10))

Exploring Skew_est and congestion

C_s = 0.0

cong/noncong

skew_est test

Path delays ms

link 10ms std 2.0

86.2, 92.1, 93.6, 87.4

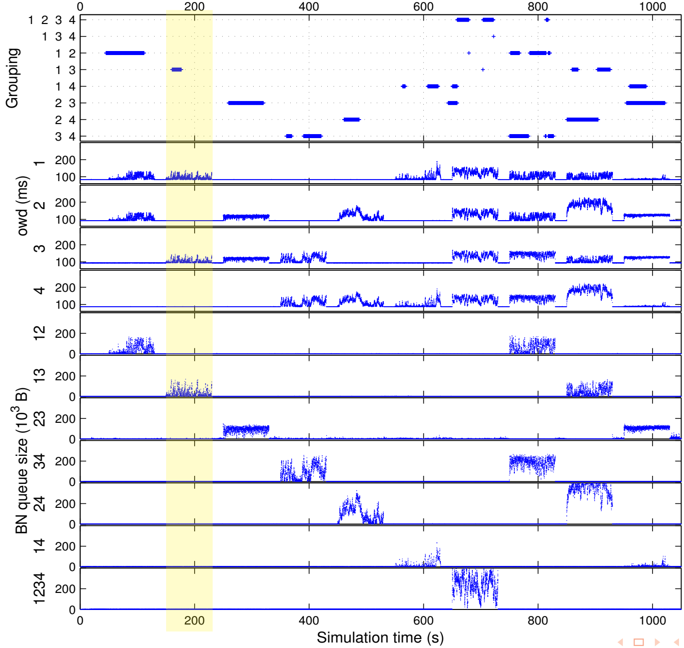

Simulation

(Queue sizes subsampled (1:200) and OWDs (1:10))

Exploring Skew_est and congestion

C_s = 0.0

cong/noncong skew_est test

Path delays ms

link 10ms std 2.0

86.2, 92.1, 93.6, 87.4

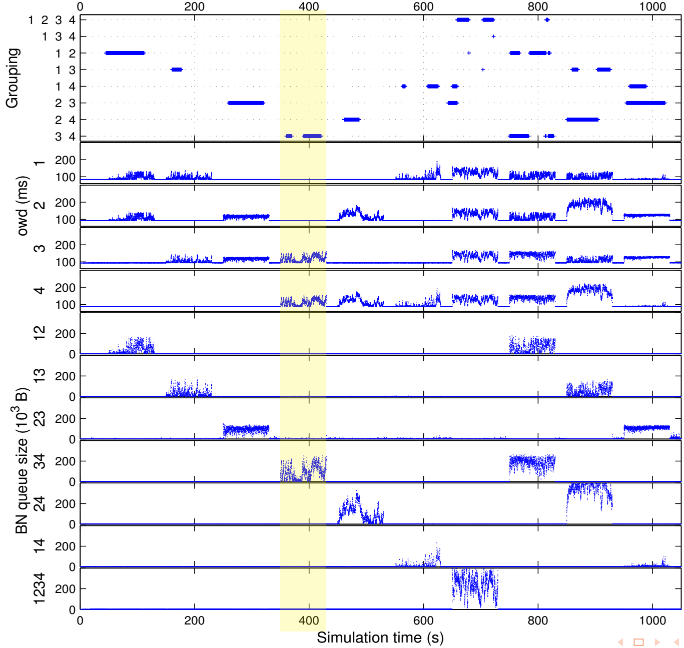

Simulation

(Queue sizes subsampled (1:200) and OWDs (1:10))

Exploring Skew_est and congestion

C_s = 0.0
cong/noncong skew_est test

Path delays ms

link 10ms std 2.0
86.2, 92.1, 93.6, 87.4

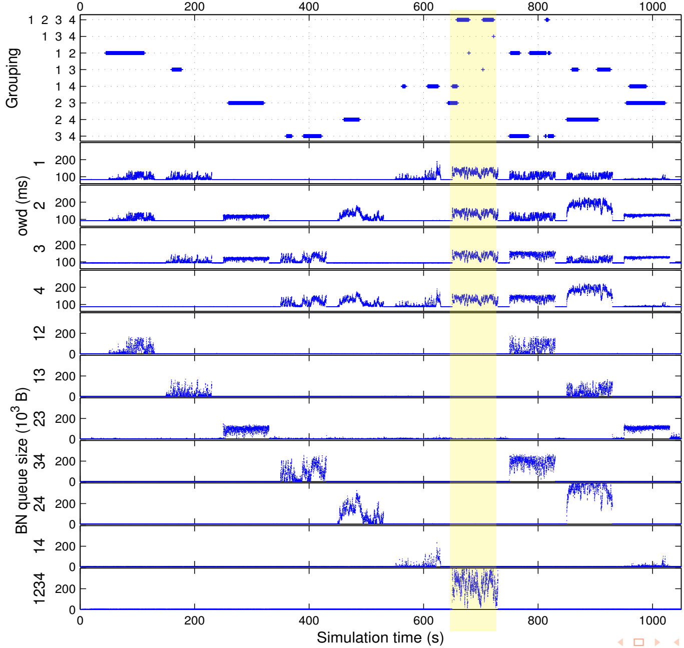

Simulation

(Queue sizes subsampled (1:200) and OWDs (1:10))

Exploring Skew_est and congestion

C_s = 0.0 cong/noncong skew_est test

Path delays ms

link 10ms std 2.0

86.2, 92.1, 93.6, 87.4

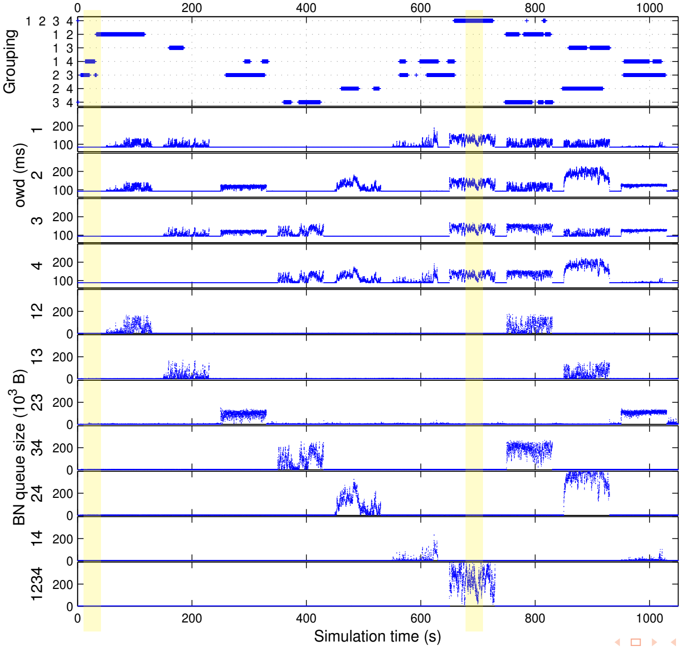

Simulation

(Queue sizes subsampled (1:200) and OWDs (1:10))

Exploring Skew_est and congestion

C_s = 0.3

cong/noncong

skew_est test

Path delays ms

link 10ms std 2.0

86.2, 92.1, 93.6, 87.4

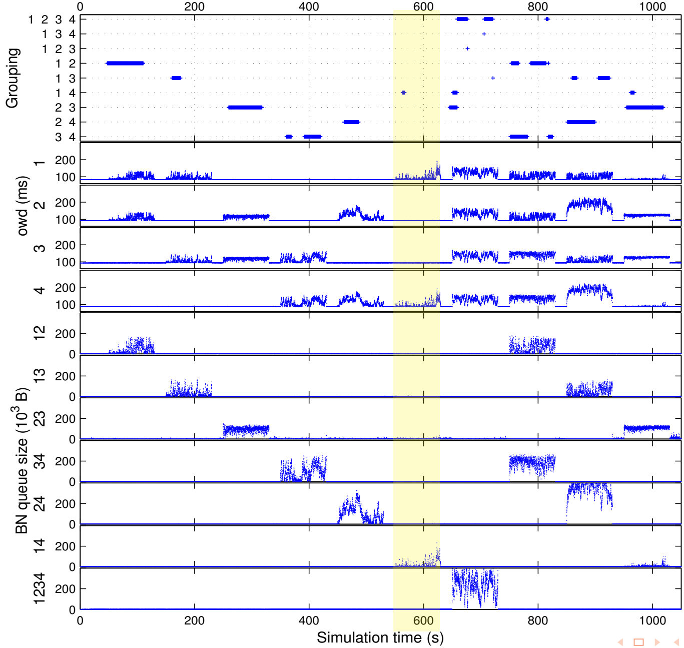

Simulation

(Queue sizes subsampled (1:200) and OWDs (1:10))

Exploring Skew_est and congestion

C_s = -0.1

cong/noncong

skew_est test

Path delays ms

link 10ms std 2.0

86.2, 92.1, 93.6, 87.4

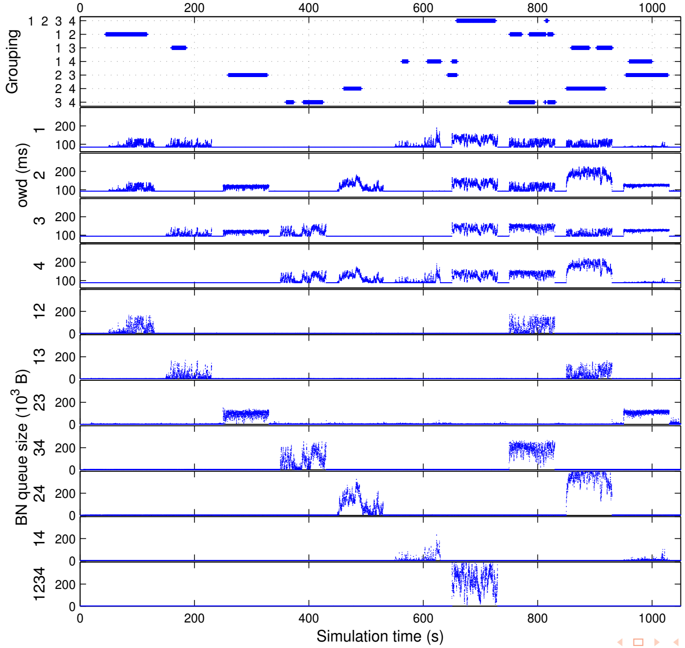

Simulation

(Queue sizes subsampled (1:200) and OWDs (1:10))

Exploring Skew_est and congestion

C_s = 0.0,
C_h = 0.3
cong/noncong skew_est test

Path delays ms

link 10ms std 2.0

86.2, 92.1, 93.6, 87.4

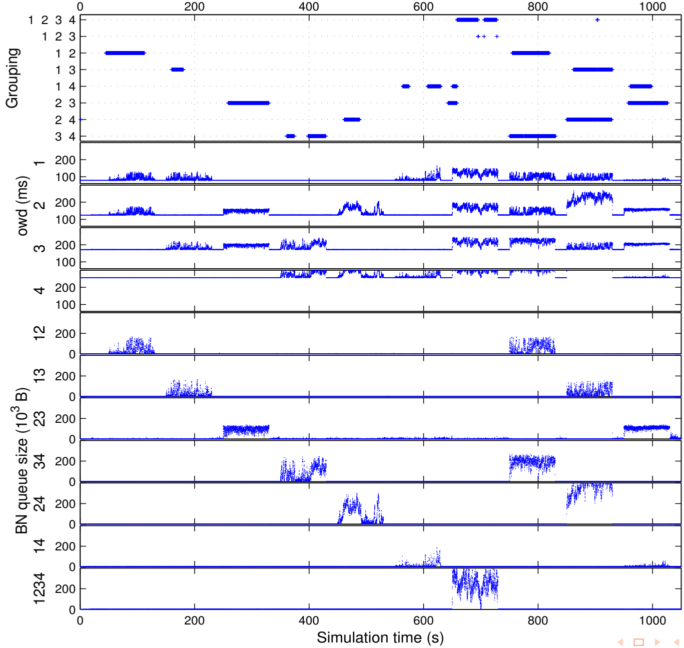

Simulation

(Queue sizes subsampled (1:200) and OWDs (1:10))

Exploring Skew_est and congestion

C_s = 0.0,
C_h = 0.3
cong/noncong
skew_est test

Path delays ms

78.6, 125.1,
171.7, 256.1

Real network experiments

Bottleneck "ground truth"

- cannot be known with 100 % certainty.
- Find thinnest shared link
  - STAB and traceroute
- Load link with distant Internet sources to create bottleneck

Real network experiments

Bottleneck "ground truth"

- cannot be known with 100 % certainty.
- Find thinnest shared link
  - STAB and traceroute
- Load link with distant Internet sources to create bottleneck

What are we testing?

- Robustness in unpredictable "real" environments

Real network experiments with NORNET

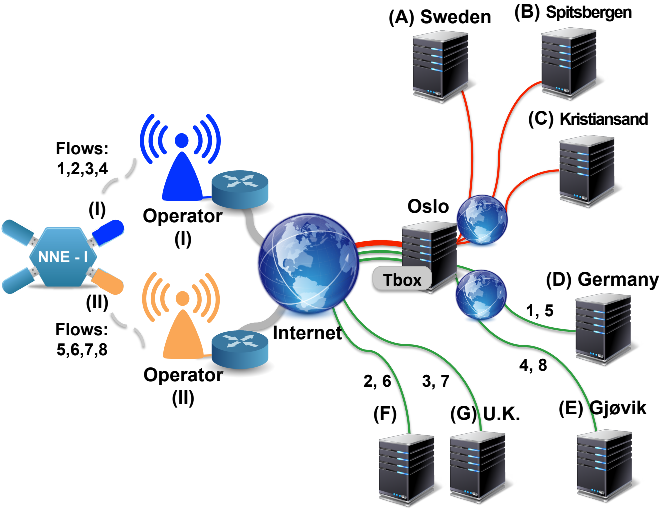

https://www.nntb.no/

Real network experiments with NORNET

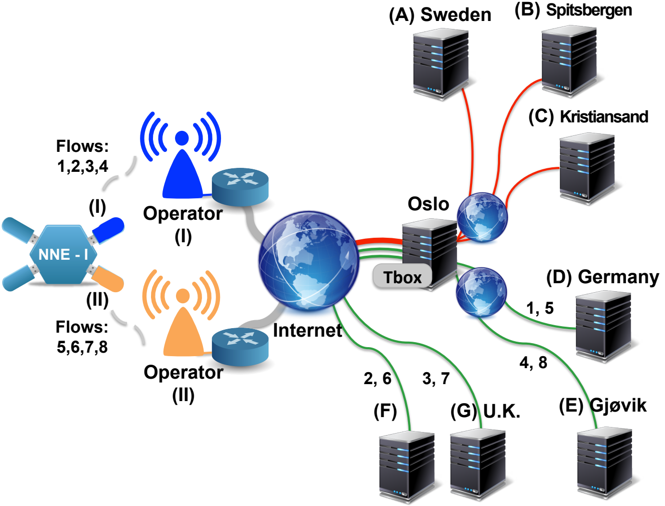

Background Traffic

- A Exponential
- B & C LRD (Hurst=0.8)
- Bottlenecks mostly on 3G link
- majority

https://www.nntb.no/

Real network experiments with NORNET

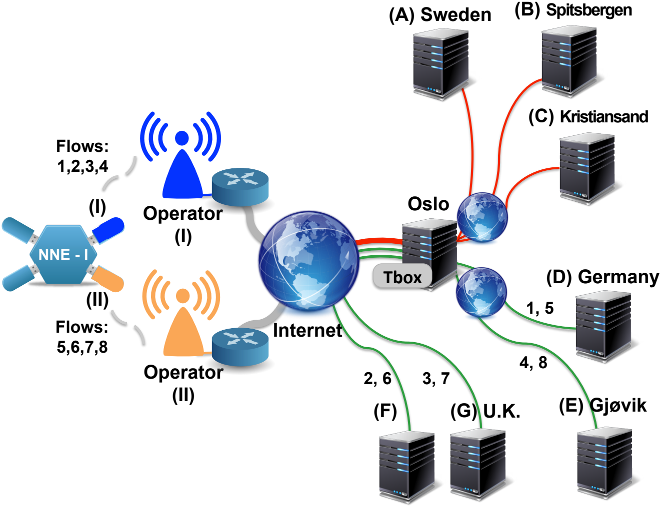

Background Traffic

- A Exponential
- B & C LRD (Hurst=0.8)
- Bottlenecks mostly on 3G link
- majority

Application Traffic

- D,E, & F exp.
- G CBR
- minority (< 10 %)

https://www.nntb.no/

NorNet results

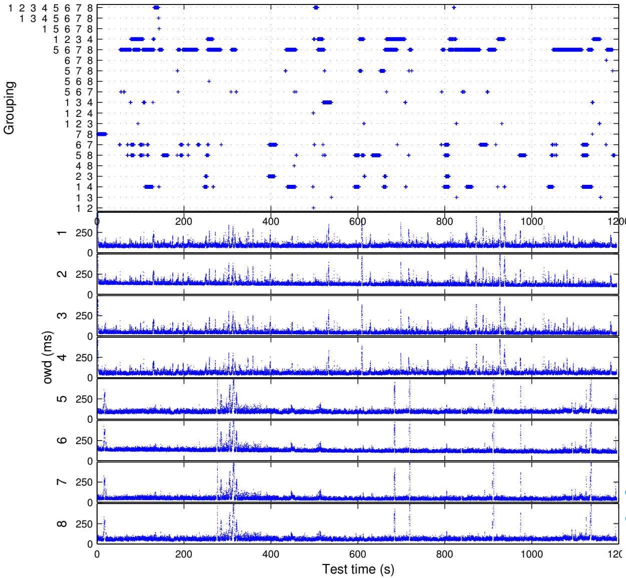

- Very different network
- Same SBD param. settings
- 247 possible non-single grouping arrangements

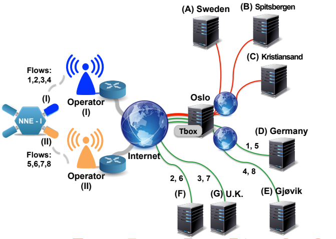

NorNet results

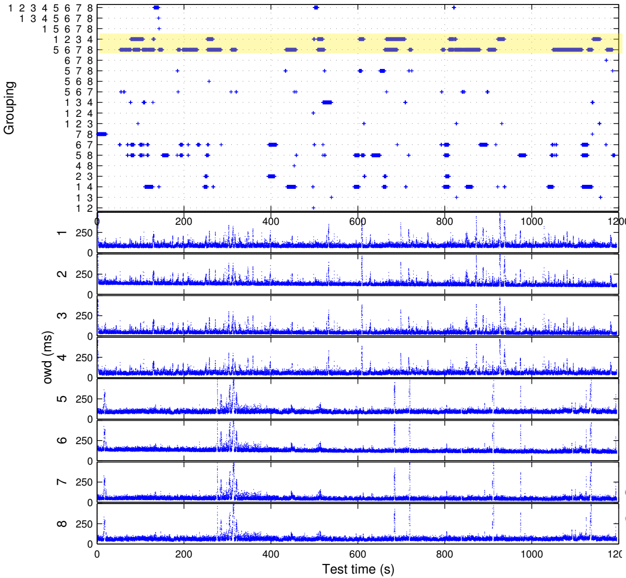

- Very different network
- Same SBD param. settings
- 247 possible non-single grouping arrangements

- Transitory BNs
- Correct during cong.
- Pkts traverse Internet

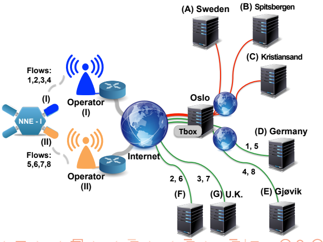

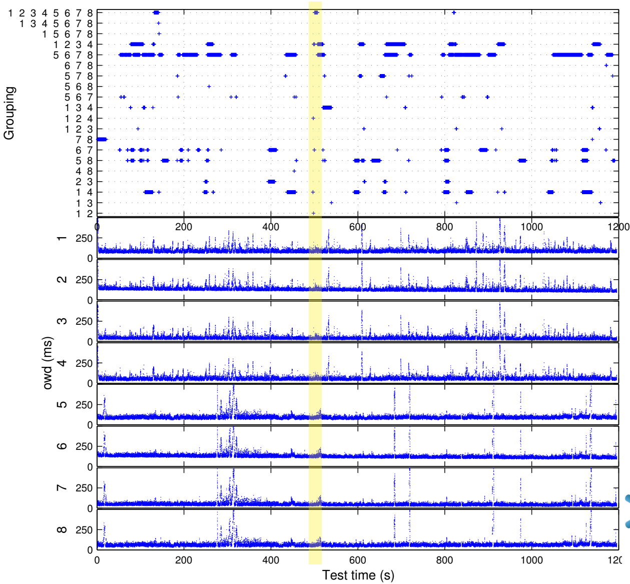

NorNet results

- Very different network
- Same SBD param. settings
- 247 possible non-single grouping arrangements

3G links sharing a common disturbance

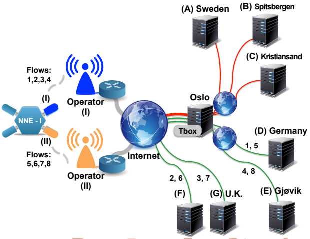

NorNet results

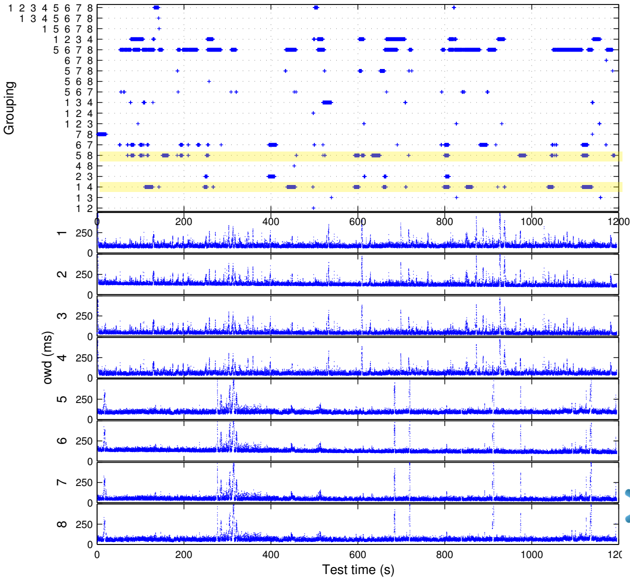

- Very different network
- Same SBD param. settings
- 247 possible non-single grouping arrangements

- 1,4,5,& 8 share Tbox
  - occasionally a BN
- 1 & 4 share 3G link I
- 5 & 8 share 3G link II

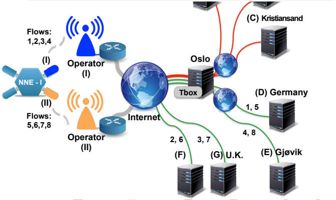

Conclusions and plans

- Define sender receiver interaction
- Evaluate the effect of time resolution
- Extend tests to wifi
- Journal
  - algorithm refinements
  - quantitative tests

Acknowledgements

The authors are funded by the European Community under its Seventh Framework Programme through the Reducing Internet Transport Latency (RITE) project (ICT-317700). The views expressed are solely those of the authors.

UiO : University of Oslo

Extra slides

MM1K PDV

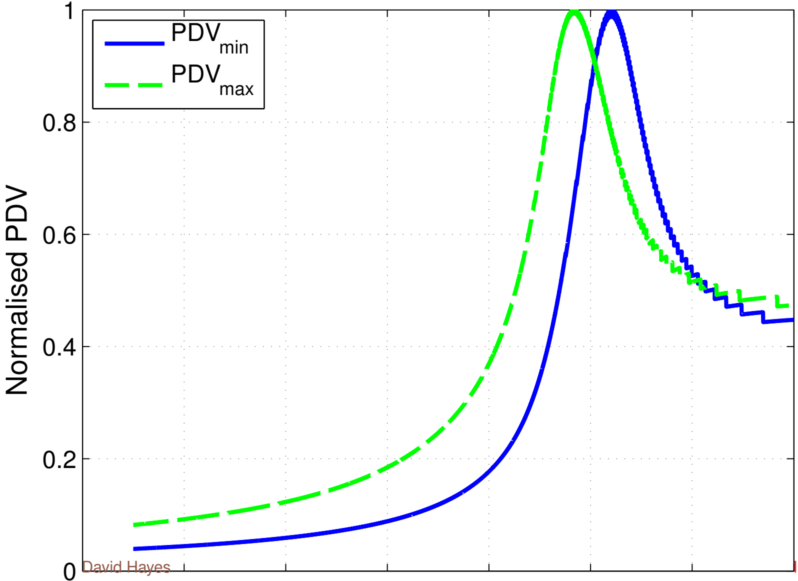

MM1K skew_est