

Update on SBD (RFC 8382)

David Hayes, Michael Welzl, and Simone Ferlin

IETF 112 RMCAT

simulamet

simula

UiO : **University of Oslo**

Some RFC 8382 related work

MPTCP with SBD

S. Ferlin, O. Alay, T. Dreibholz, D. A. Hayes, and M. Welzl, “Revisiting congestion control for multipath TCP with shared bottleneck detection,” in *Proc. of the IEEE International Conference on Computer Communications (INFOCOM)*, Apr. 2016, pp. 1–9. DOI: [10.1109/INFOCOM.2016.7524599](https://doi.org/10.1109/INFOCOM.2016.7524599)

MPQUIC with SBD

T. Paiva, S. Ferlin, and B. Kimura, “Métricas de detecção de gargalos compartilhados em transmissões por múltiplos caminho em mpquic,” in *Proc. 7th Pré-IETF Workshop*, Porto Alegre, RS, Brasil: SBC, 2020, pp. 15–28. DOI: [10.5753/wpietf.2020.13793](https://doi.org/10.5753/wpietf.2020.13793)

Thorough comparative study of RFC 8382

D. Hayes, M. Welzl, S. Ferlin, D. Ros, and S. Islam, “Online identification of groups of flows sharing a network bottleneck,” *IEEE/ACM Trans. Netw.*, vol. 28, no. 5, pp. 2229–2242, 2020. DOI: [10.1109/TNET.2020.3007346](https://doi.org/10.1109/TNET.2020.3007346)

Some RFC 8382 related work

MPTCP with SBD

S. Ferlin, O. Alay, T. Dreibholz, D. A. Hayes, and M. Welzl, “Revisiting congestion control for multipath TCP with shared bottleneck detection,” in *Proc. of the IEEE International Conference on Computer Communications (INFOCOM)*, Apr. 2016, pp. 1–9, DOI:

10.1109/INFCOM.2016.7551411 Thorough comparative study of RFC 8382

D. Hayes, M. Welzl, S. Ferlin, D. Ros, and S. Islam, “Online identification of groups of flows sharing a network bottleneck,” *IEEE/ACM Trans. Netw.*, vol. 28, no. 5, pp. 2229–2242, 2020. DOI: 10.1109/TNET.2020.3007346

Thorough comparative study of RFC 8382

D. Hayes, M. Welzl, S. Ferlin, D. Ros, and S. Islam, “Online identification of groups of flows sharing a network bottleneck,” *IEEE/ACM Trans. Netw.*, vol. 28, no. 5, pp. 2229–2242, 2020. DOI: 10.1109/TNET.2020.3007346

Comparative evaluation of RFC 8382

Algorithms we compare

- Vanilla RFC 8382
- RFC 8382 stats with dynamic online clustering.
- Offline wavelet filtered cross-correlation
- Online adaption of the wavelet filtered cross-correlation

Experiments we run

- Teacup testbed
- Complex simulations
- Internet and Mobile via NorNet

Comparisons presented today

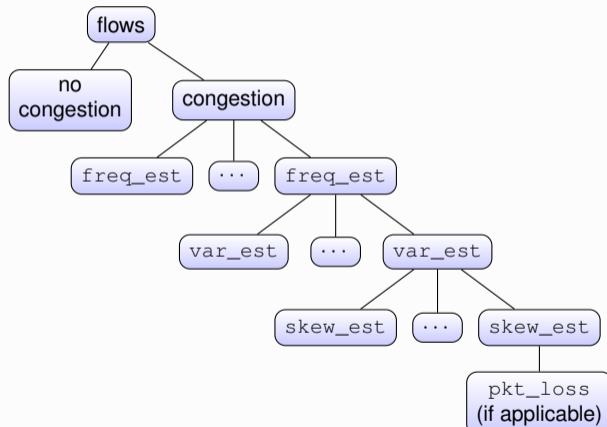
- Time to bottleneck detection
 - start and stop
- Effect of path delays
- Effect of parallel bottlenecks
- ***More in the paper[‡]***

[‡]D. Hayes, M. Welzl, S. Ferlin, D. Ros, and S. Islam, "Online identification of groups of flows sharing a network bottleneck," *IEEE/ACM Trans. Netw.*, vol. 28, no. 5, pp. 2229–2242, 2020. DOI: 10.1109/TNET.2020.3007346

Note: Plots are adapted from the paper and used with permission License Number 5184241200906 © IEEE 2020.

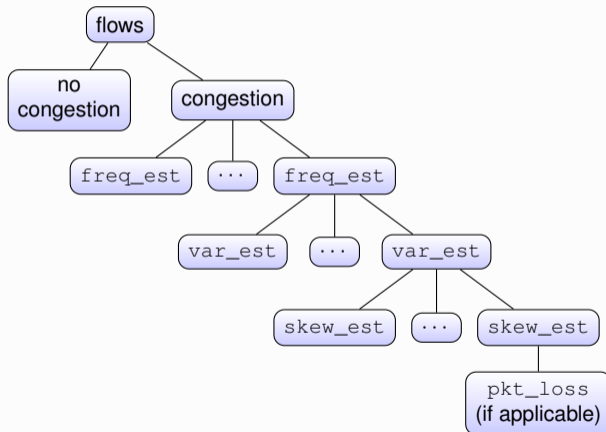
Summary statistic based grouping mechanisms

Divide-and-conquer (rmcatSBD)



Summary statistic based grouping mechanisms

Divide-and-conquer (rmcatSBD)



Dynamic clustering (dcSBD)

Most common question in RMCAT:

Why not use a clustering algorithm?

- Unknown number of groups
- Changing number of groups
- Closest (inverse square law) neighbour in multi-dimensional space
- Iterative based on past grouping



Wavelet filtered correlation methods

Original offline (pDCW)[†]

Matlab calculates optimal filter based on entire trace

Wavelet filter the entire trace

pairwise correlation coefficient calculation at intervals

Group flows where correlation coefficient > 0.5

Online (opDCW)

Calculate filter coefficients whenever `skew_est` becomes < 0

shift OWD measurements through wavelet filters



pairwise correlation coefficient calculation at intervals

Group flows where correlation coefficient > 0.5

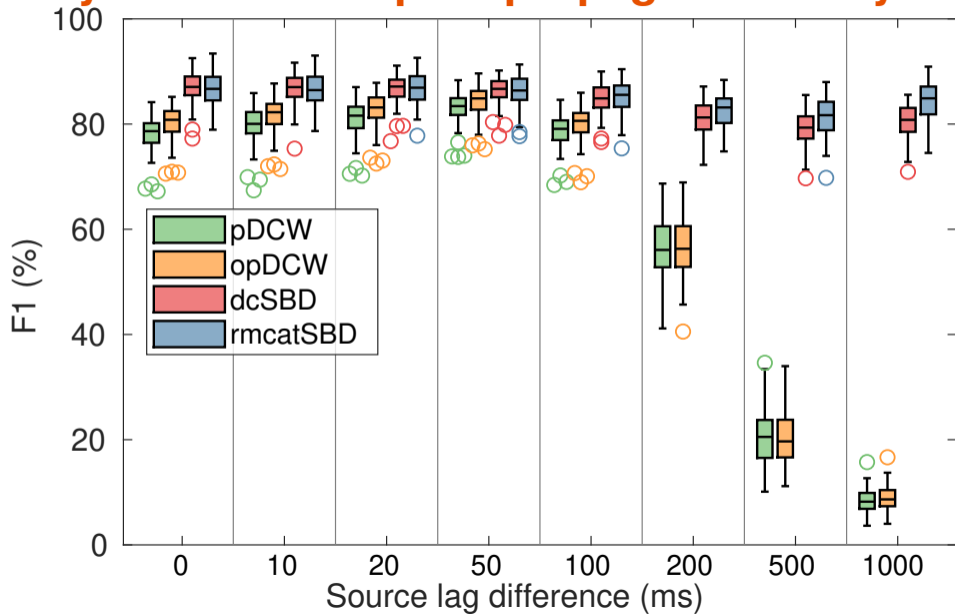
[†]M. S. Kim, T. Kim, Y.-J. Shin, S. S. Lam, and E. J. Powers, "A wavelet-based approach to detect shared congestion," *IEEE/ACM Trans. Netw.*, vol. 16, no. 4, pp. 763–776, Aug. 2008. DOI: 10.1109/TNET.2007.905599

Detection delays


- Start—delay from ground truth indication of a bottleneck starting to correctly grouping flows.
- Stop—delay from ground truth indication of a bottleneck ceasing to correctly detecting that flows cease to share a bottleneck.



Accuracy wrt different path propagation delays



Accuracy wrt the number of parallel bottlenecks



Conclusions

RFC8382 performs comparatively well

- Simple and Light
- Handles path delays well.
- NB summary stats introduce lag
- dcSBD provides improvements for higher numbers of parallel bottlenecks

General limitation

- **All** SBD methods assume BNs have measurable similarities for all flows sharing that BN, ie a *“BN Signature”*
- A counter example:
 - The flow sending pattern completely dominates the BN characteristics.

Conclusions

RFC8382 performs comparatively well

- Simple and Light
- Handles path delays well.
- NB summary stats introduce lag
- dcSBD provides improvements for higher numbers of parallel bottlenecks

General limitation

- **All** SBD methods assume BNs have measurable similarities for all flows sharing that BN, ie a “*BN Signature*”
- A counter example:
 - The flow sending pattern completely dominates the BN characteristics.

Path forward

- We don't think a Proposed Standard RFC is useful unless the limitation is addressed
- How?
 - can be algorithm-dependent; or
 - a generic mechanism may be possible, e.g. measure the sending characteristics, at the sender, and either discount them from the receiver measured characteristics (maybe similar to how a speaker phone cancels the feedback)—requires research; or
 - just flag that it is happening and treat the SBD results accordingly.
- Calling for interested parties to continue this work