

Invited Talk at the New York University, Polytechnic School of Engineering

The NorNet Testbed –
Overview and Selected Results
from Multi-Path Transport Research

Thomas Dreibholz, dreibh@simula.no

Simula Research Laboratory

May 5, 2014



Contents

- Motivation
- The NorNet Testbed
 - NorNet Core
 - NorNet Edge
- Selected Research Topics
- Conclusion

Overview: Motivation

- Motivation
- The NorNet Testbed
 - NorNet Core
 - NorNet Edge
- Selected Research Topics
- Conclusion

Motivation: Robust Networks

- More and more applications rely on ubiquitous Internet access!
- However, our current networks are not as robust as they should be ...




How to make networks more robust?

Resilience by Redundancy

Multi-Homing

- Connections to multiple Internet Service Providers (ISP)
- Idea: if one ISP has problems, another connection still works



Is resilience really improved? What about multi-path transport?

Idea: A Testbed for Multi-Homed Systems

Research in realistic setups is necessary!

- A multi-homed Internet testbed would be useful
 - Something like PlanetLab?
 - Perhaps with better node availability?
 - Support for mobile access (e.g. 3G) as well as wired?
- **NorNet** – A research testbed for multi-homed systems!
 - Lead by the Simula Research Laboratory in Fornebu, Norway
 - Supported by Forskningsrådet

NORNET

<https://www.nntb.no>

Overview: The NorNet Project

- Motivation
- The NorNet Testbed
 - NorNet Core
 - NorNet Edge
- Selected Research Topics
- Conclusion

Goals of the NorNet Project

- Building up a **realistic** multi-homing testbed
- Wired and wireless
 - Wired → “NorNet Core”
 - Wireless → “NorNet Edge”
- **Perform research with the testbed!**



How to get a realistic testbed?

Idea: Distribution of NorNet over whole Norway

- **Challenging topology:**
 - Large distances
 - A few “big” cities, many large rural areas
 - Svalbard:
 - Interesting location
 - Many polar research institutions
- **NorNet Core:**
 - Currently 11+3 sites
- **NorNet Edge:**
 - Currently ca. 400 nodes



Overview: NorNet Core


- Motivation
- The NorNet Testbed
 - NorNet Core
 - NorNet Edge
- Selected Research Topics
- Conclusion

Idea: Tunnelling

- Researchers require control over used ISP interfaces
 - Which outgoing (local site) interface
 - Which incoming (remote site) interface

- Idea: Tunnels among sites

- Router at site A: IPs A_1, A_2, A_3
- Router at site B: IPs B_1, B_2




- IP tunnel for each combination:

$A_1 \leftrightarrow B_1, A_1 \leftrightarrow B_2, A_2 \leftrightarrow B_1, A_2 \leftrightarrow B_2, A_3 \leftrightarrow B_1, A_3 \leftrightarrow B_2$

- Fully-connected tunnel mesh among NorNet Core sites (< 20)
- Each site's router (called **tunnelbox**) maintains the tunnels
 - Static tunnels
 - NorNet-internal addressing and routing over tunnels

A NorNet Core Site Deployment

A usual NorNet Core site:

- 1x switch
- 4x server
 - 1x tunnelbox
 - 3x research systems
- At least two ISP connections
 - Uninett 
 - Other providers
- IPv4 and IPv6 (if available)



UNIS

Longyearbyen 78.2°N,15.6°E



NorNet Core på Svalbard

Site Deployment Status (May 2014)

Site	Location	ISP 1	ISP 2	ISP 3
Simula Research Laboratory	Fornebu, Akershus	UNINETT	Kvantel	Telenor ²
Universitetet i Oslo	Blindern, Oslo	UNINETT	PowerTech	Broadnet ²
Høgskolen i Gjøvik	Gjøvik, Oppland	UNINETT	PowerTech	
Universitetet i Tromsø	Tromsø, Troms	UNINETT ¹	PowerTech	Telenor ²
Universitetet i Stavanger	Stavanger, Rogaland	UNINETT ¹	BKK	
Universitetet i Bergen	Bergen, Hordaland	UNINETT	BKK	
Universitetet i Agder	Kristiansand, Vest-Agder	UNINETT	PowerTech	
Universitetet på Svalbard	Longyearbyen, Svalbard	UNINETT ¹	Telenor ^{2,4}	
NTNU Trondheim	Trondheim, Sør-Trøndelag	UNINETT	PowerTech	
Høgskolen i Narvik	Narvik, Norland	UNINETT	PowerTech	
Høgskolen i Oslo og Akershus	St. Hanshaugen, Oslo	UNINETT		
University of Duisburg-Essen	Essen/Germany	DFN	Versatel ^{2,3}	
Hainan University	Haikou, Hainan/China	CERNET ¹	Unicom ¹	
Karlstads Universitet	Karlstad, Värmland/Sweden	SUNET	– ⁴	

1) IPv6 available from ISP, but not deployed to setup

2) IPv6 not available from ISP ☹

3) Consumer-grade ADSL connection

4) Negotiations in progress

Remote Systems


Our servers may be really remote!

The “road” to Longyearbyen på Svalbard, 78.2°N

Virtualisation

“Anything that can go wrong, will go wrong.”
[Murphy's law]

- Experimentation software is experimental
- How to avoid software issues making a remote machine unusable?
- Idea: virtualisation
 - Lightweight, stable software setup:
Ubuntu Server 12.04 LTS
 - VirtualBox 4.3
 - Other software runs in VirtualBox VMs:
 - Tunnelbox VM on physical server #1
 - 2 LXC-based research node VMs on physical servers #2 to #4
 - In case of problem: manual/automatic restart or reinstall of VM



Idea: *PlanetLab*-based Software for Experiments

- Key idea:
 - Researchers should get virtual machines for their experiments
 - Like *PlanetLab* ...
 - ... but with multi-homing and IPv6, of course
- *PlanetLab* software:
 - Different “stable” distributions: *PlanetLab*, *OneLab*, etc.
 - Current implementation: based on *Linux VServers*
 - Not in mainline kernel
 - Patched kernel, makes upgrades difficult
 - The future: **Linux Containers (LXC)**
 - Active development by *PlanetLab/OneLab*
 - We are maintaining a NorNet-specific branch

Overview: NorNet Edge

- Motivation
- The NorNet Testbed
 - NorNet Core
 - NorNet Edge
- Selected Research Topics
- Conclusion

The NorNet Edge Box: Ready for Deployment (1)

Box contents:

- Ufoboard or Beagle Bone embedded Linux system
- 4x USB UMTS:
 - Telenor, NetCom,
 - Network Norway, Tele2
- 1x ICE CDMA mobile broadband
- 1x Ethernet
- 1x WLAN (optional)
- Power supplies
- Handbook



The NorNet Edge Box: Ready for Deployment (2)

Ufoboard:

- Debian Linux
- Kernel 3.11.x
- **MPTCP (0.88)**



NorNet Edge Visualisation

NorNet Edge Visualization

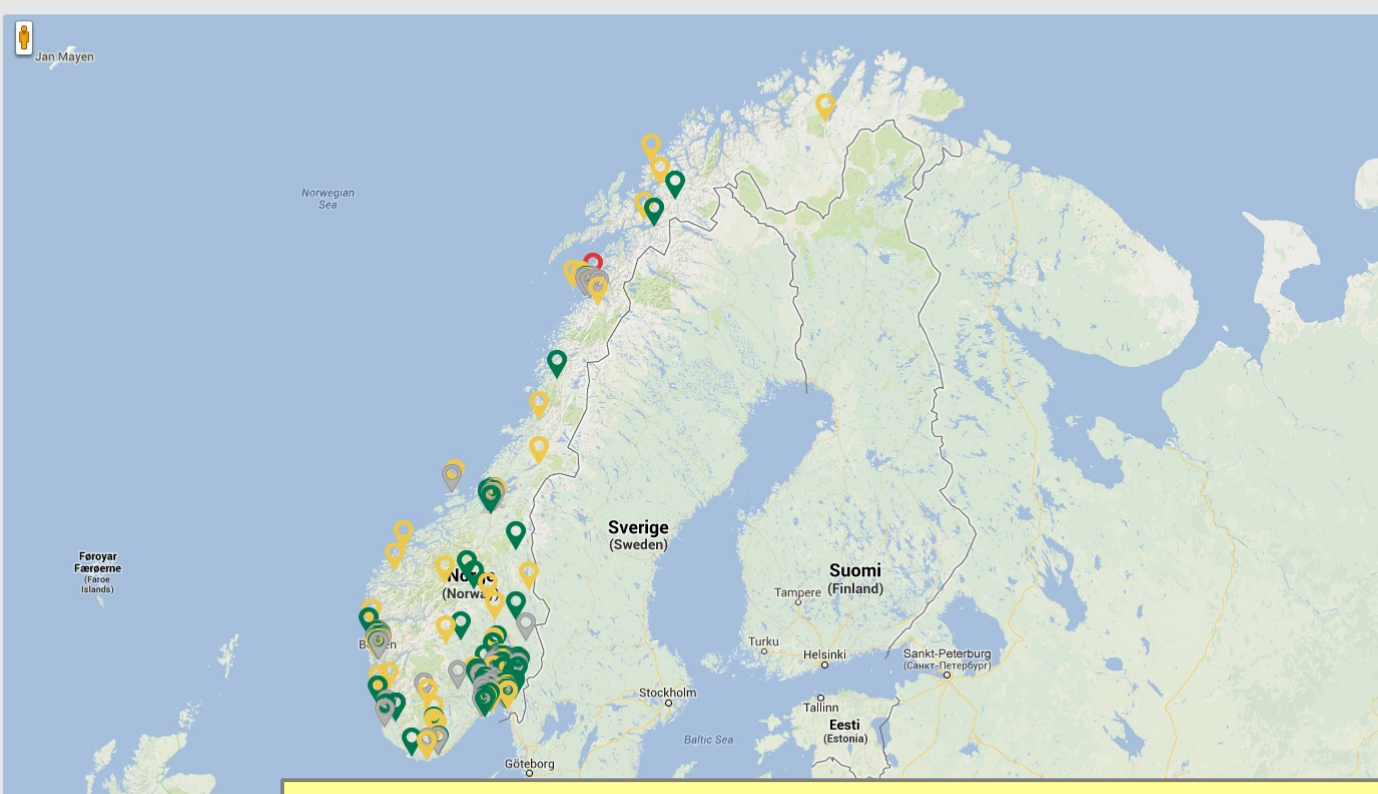
Dashboard Comparative view 2013/7/26 13:54:32

OPERATOR: All

COLOR CODE

- All networks available
- Some networks available
- No networks available
- Node unreachable

UP RATE, %: 83



Search by location or node id

LOCATION	NODE ID	STATUS
Åmot	nne255	●
Asker	nne379	●
Aurskog-Høland	nne365	●
Aurskog-Høland	nne367	●
Aurskog-Høland	nne376	●
Aurskog-Høland	nne377	●
Aurskog-Høland	nne411	●
Austrheim	nne554	●
Bardu	nne272	●
Berg	nne262	●
Bergen	nne202	●
Bergen	nne236	●
Bergen	nne238	●
Bergen	nne283	●
Bergen	nne284	●
Bergen	nne285	●
Bergen	nne286	●
Bergen	nne291	●
Bergen	nne292	●
Bergen	nne293	●

See <http://demo.robustenett.no>

© Simula Research Laboratory. Simula NorNet

Overview:


Selected Research Topics

- Motivation
- The NorNet Testbed
 - NorNet Core
 - NorNet Edge
- Selected Research Topics
- Conclusion


Multi-Path TCP

Multi-Path TCP (MPTCP):

- Multi-path extension for TCP (RFC 6182/RFC 6824)
- Combination of sub-flows (like TCP)
- Idea: improve throughput and resilience




RFC6182



RFC6182

How behaves MPTCP in real networks?

Evaluation in NorNet Edge

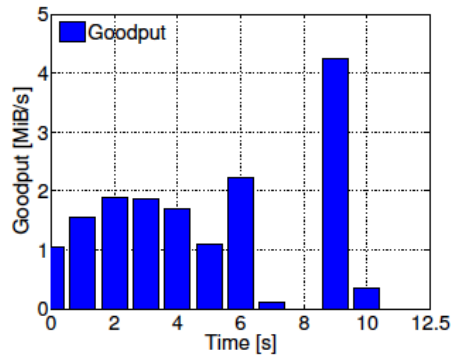


Using NorNet Edge nodes (NNE):

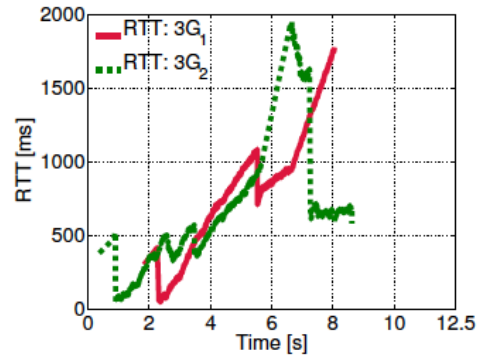
- Two 3G ISPs (2G/3G)
- WLAN at the node location:
real-world public WLAN hotspot (i.e. many users and interferences)

Standard MPTCP does not work very well ...

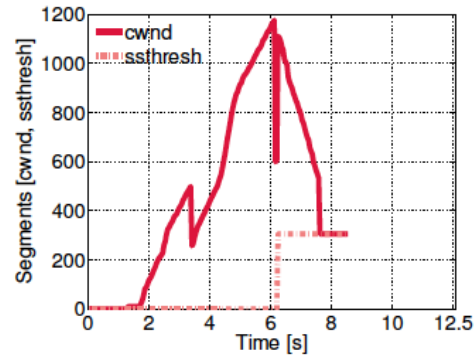
3G + 3G Paths



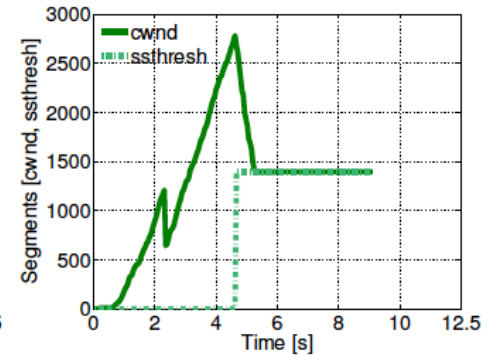
(a) Goodput: 3G₁ and 3G₂



(b) RTT: 3G₁ and 3G₂



(c) Congestion Window: 3G₁



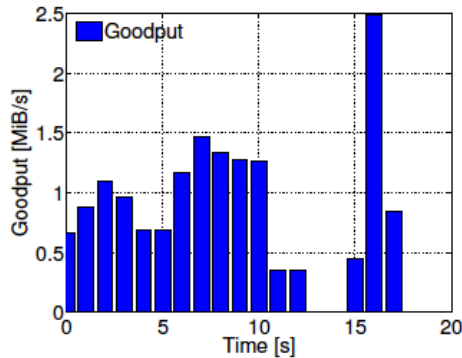
(d) Congestion Window: 3G₂

Paths are heterogeneous

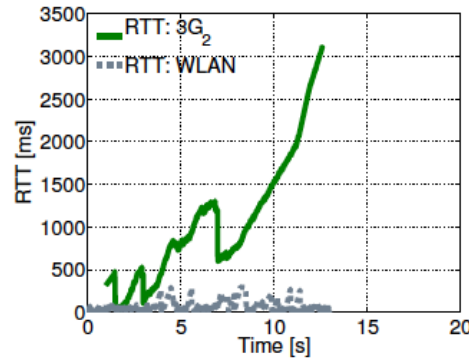
- Varying overall goodput
- High RTTs → bufferbloat!

Performance issues caused by bufferbloat!

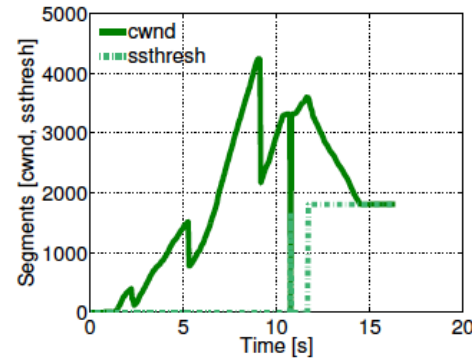
Standard MPTCP does not work very well ... 3G + WLAN Paths



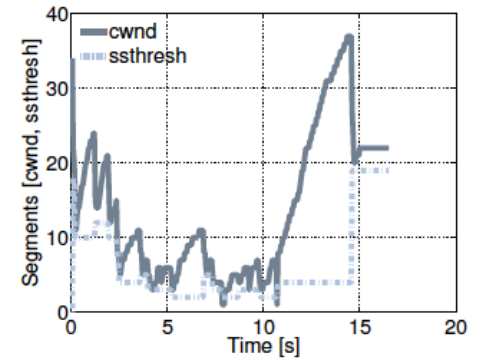
(a) Goodput: 3G₂ and WLAN



(b) RTT: 3G₂ and WLAN



(c) Congestion Window: 3G₂



(d) Congestion Window: WLAN

The same performance issues apply for combining 3G and WLAN

How to avoid the bufferbloat issues?

Multi-Path Transport Bufferbloat Mitigation (MPT-BM)

Algorithm 1 Per-Subflow Bufferbloat Mitigation by MPT-BM

Initialization:

$$sRTT \leftarrow \infty$$

$$sRTT_{\min} \leftarrow \infty$$

RTT estimation:

$$sRTT_{\min} \leftarrow \min(sRTT_{\min}, sRTT)$$

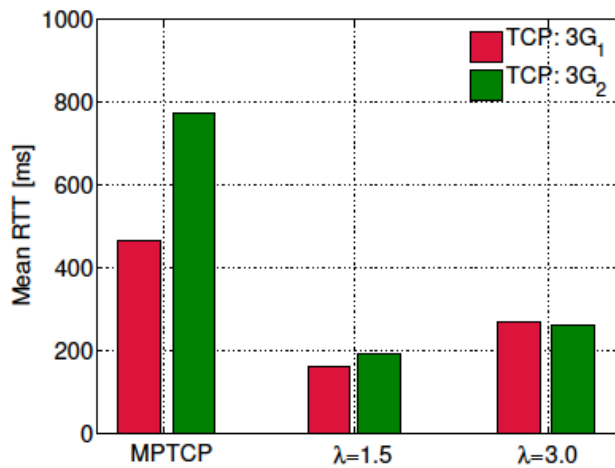
How many segments can be sent?

$$cwnd_{\text{limit}} \leftarrow \lambda * (sRTT_{\min} / sRTT) * cwnd$$


$$\text{send} \leftarrow \begin{cases} \max(0, (\min(cwnd, cwnd_{\text{limit}}) - \text{inflight})) & (RTT_{\min} \geq \Theta) \\ \max(0, cwnd - \text{inflight}) & (RTT_{\min} < \Theta) \end{cases}$$

- Idea:
 - Avoid extreme growth of the congestion window (cwnd)
 - Limitation controllable (parameter: λ)
 - Only necessary for large RTTs (parameter: Θ)


Evaluation: Round-Trip Times



(a) RTT: 3G₁ + 3G₂



(b) RTT: 3G₁ + WLAN




(c) RTT: 3G₂ + WLAN

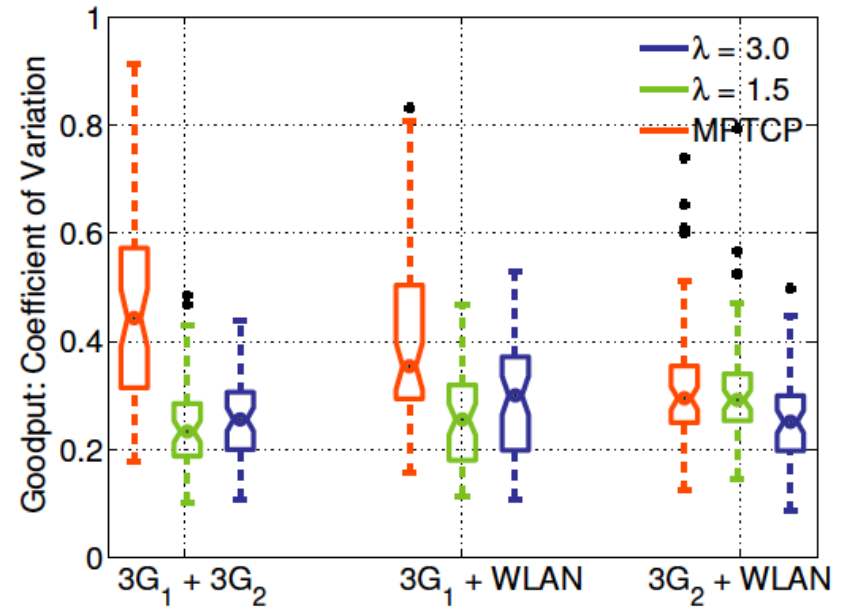
- Significant RTT reduction => bufferbloat is avoided

How does it affect the goodput?

Evaluation: Goodput



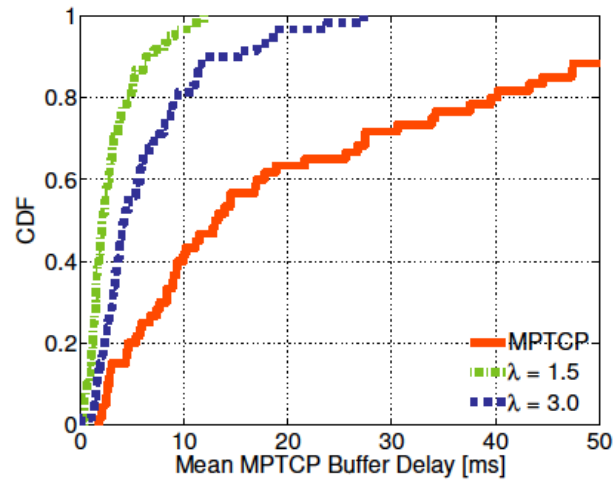
(a) Application Goodput




(b) Coefficient of Variation

- No negative impact, sometimes even slightly better
- Variance is reduced


Evaluation: MPTCP Buffer Delay and Buffer Size



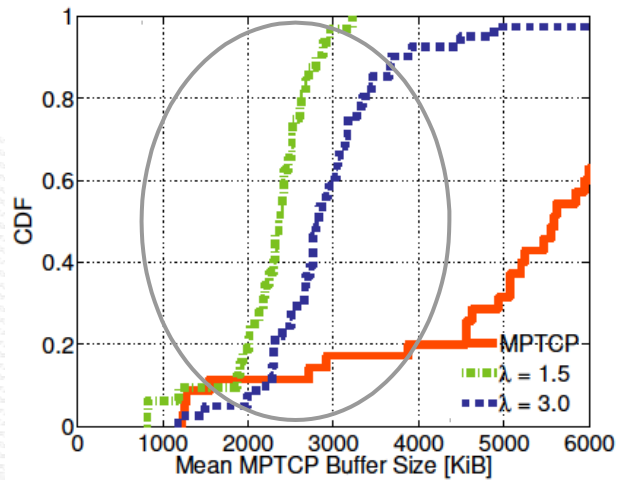
(a) 3G₁ + 3G₂



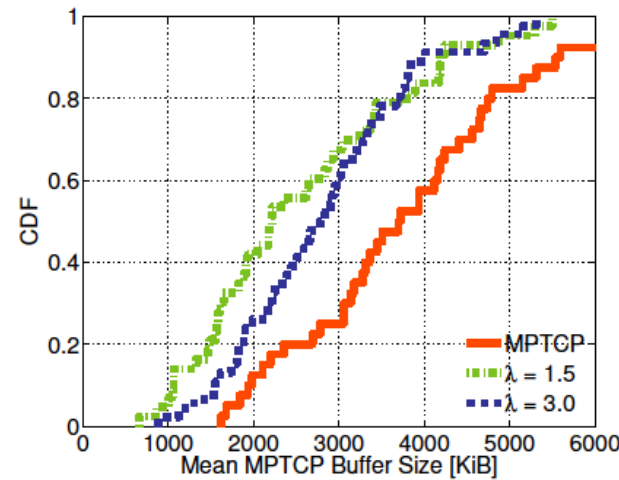
(b) 3G₁ + WLAN



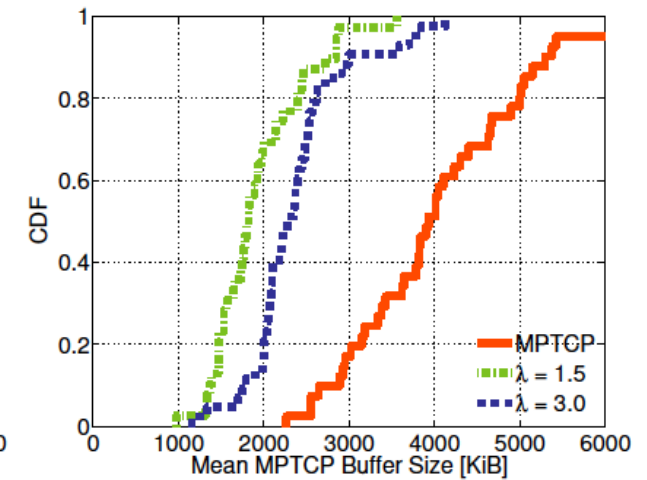
(c) 3G₂ + WLAN



(a) 3G₁ + 3G₂



(b) 3G₁ + WLAN



(c) 3G₂ + WLAN

Further Details

- Ferlin-Oliveira, Simone; Dreibholz, Thomas; Alay, Özgü and Kvalbein, Amund: «*Measuring the QoS Characteristics of Operational 3G Mobile Broadband Networks*» (PDF, 1524 KiB, in English), in Proceedings of the 4th International Workshop on Protocols and Applications with Multi-Homing Support (PAMS), Victoria, British Columbia/Canada, May 16, 2014.
- Ferlin-Oliveira, Simone; Dreibholz, Thomas and Alay, Özgü: «*Tackling the Challenge of Bufferbloat in Multi-Path Transport over Heterogeneous Wireless Networks*», in Proceedings of the IEEE/ACM International Symposium on Quality of Service (IWQoS), ISBN 978-1-4799-4852-9, Hong Kong, May 2014.

<https://www.nntb.no/publications>

Overview: Conclusion

- Motivation
- The NorNet Testbed
 - NorNet Core
 - NorNet Edge
- Selected Research Topics
- Conclusion

Conclusion and Future Work

- The NorNet testbed is progressing!
 - Initial deployment completed
 - Ready for experiments (also for your experiments!)
- Future work:
 - Make more NorNet Core sites multi-homed (additional ISPs, IPv6)
 - Some additional sites
 - Improve and refine management software
 - Get more users (may be you?)

And, of course, do further research!

“NorNet wants to be a building block of the railroad to heaven” ...



... and not be another unused testbed that paves the road to hell!

Coming Soon: The 2nd NorNet Users Workshop (NNUW-2)

See <https://www.nntb.no/>

**Dates will be announced soon!
(Probably: End of August 2014)**



Any Questions?

N  RNET

Visit <https://www.nntb.no> for further information!