

Invited Talk at the Universidad de Castilla-La Mancha

Multi-Path Transport — From Simulations to Real-World Internet Measurements

Thomas Dreibholz

Simula Research Laboratory

16 February 2017



Contents

- About Norway and the Simula Research Laboratory
- From TCP to Multi-Path Transport
- Simulations
- Lab Setups and Internet Setups
- The NorNet Testbed
- Conclusion
- Literature

Overview:

About Norway and the Simula Research Laboratory

- About Norway and the Simula Research Laboratory
- From TCP to Multi-Path Transport
- Simulations
- Lab Setups and Internet Setups
- The NorNet Testbed
- Conclusion
- Literature

Where is Norway?



Oslo 奥斯陆

ca. 2,500 km



Albacete
阿尔瓦塞特

ca. 7,100 km



Beijing 北京

ca. 16,000 km




Sydney
悉尼

Facts about Norway

Capital:	Oslo
Size:	ca. 385,000 km ²
Population:	ca. 5,214,000
Internet TLD:	.no




The Kingdom of Norway (Kongeriket Norge)



King Harald V

Arctic Circle



The Simula Research Laboratory

- Located in Fornebu
 - Just outside of Oslo
 - In the IT Fornebu complex
- Public limited company
 - Non-profit research organisation
 - Ca. 160 people from all over the world
- Research groups
 - Scientific Computing
 - Software Engineering
 - **Resilient Networks and Applications**
- Norway's leading place for computer science research

[**simula** . research laboratory]



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


Overview:

From TCP to Multi-Path Transport

- About Norway and the Simula Research Laboratory
- From TCP to Multi-Path Transport
- Simulations
- Lab Setups and Internet Setups
- The NorNet Testbed
- Conclusion
- Literature

„Classic“ Internet Communication


- Example: World-Wide Web



- Client ↔ Server Communication
 - 1 network interface per device → 1 **IPv4** address
 - Communication with Transmission Control Protocol (**TCP**)

The Current and Future Internet

The Big Picture



- IPv6
 - Devices are frequently IPv4/IPv6 dual stack
 - Usually multiple addresses per interface
- Mobility → address change
- Devices with multiple interfaces
 - Router
 - **Smartphone** (LTE/UMTS, WLAN, Bluetooth?)
 - **Laptop** (Ethernet, WLAN, LTE/UMTS?)



Multi-Homing and Multi-Path Transport

- Multi-Homing
 - Multiple interfaces (addresses)
 - **Redundancy** → Communication even when some paths fail
- Multi-Path Transport
 - Also utilise paths simultaneously → better throughput
 - **MPTCP**: Multi-Path TCP
 - **CMT-SCTP**: Concurrent Multi-Path Transfer for SCTP

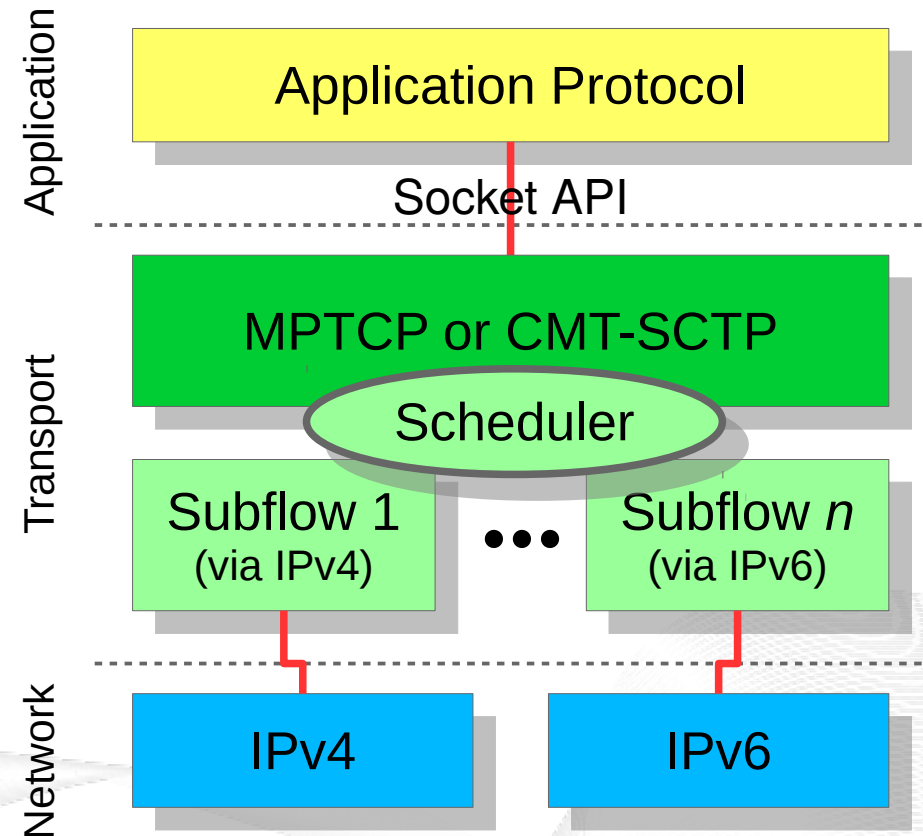


SCTP: Stream Control Transmission Protocol
TCP: Transmission Control Protocol

Hot topic in research and standardisation!

Multi-Path Transport with MPTCP and CMT-SCTP

- Subflow \leftrightarrow path
- Fairness
 - Paths may overlap (fully oder partially)
- Scheduling
 - Different path characteristics
 - Bandwidth
 - Latency and jitter
 - Packet loss



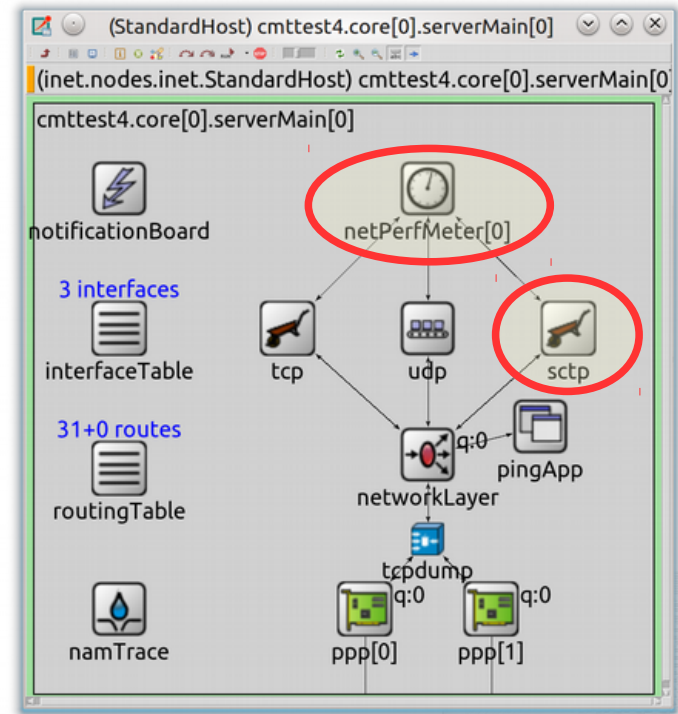
Complex system \rightarrow analyses are necessary!

Overview: Simulations

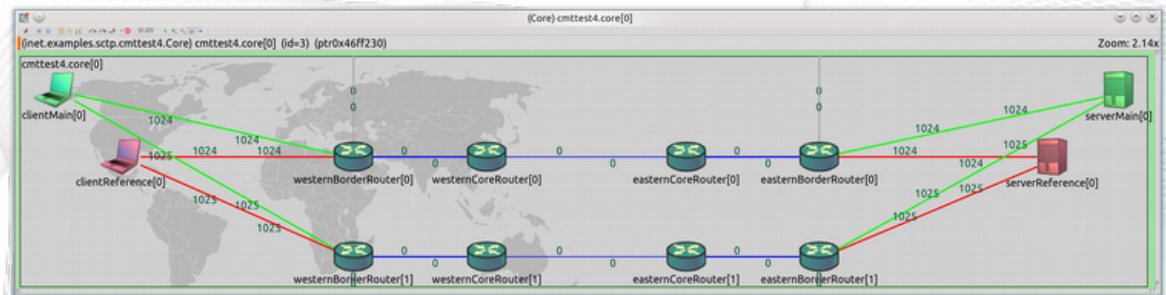
- About Norway and the Simula Research Laboratory
- From TCP to Multi-Path Transport
- Simulations
- Lab Setups and Internet Setups
- The NorNet Testbed
- Conclusion
- Literature

Research (1) – The Beginning: Simulations

- SCTP in OMNeT++
 - SCTP extensions
 - Application model “**NetPerfMeter**”
- Open Source
 - Mostly part of the OMNeT++
INET Framework
 - Some parts still need merging




<https://inet.omnetpp.org>



Challenge: CMT-SCTP over Dissimilar Paths


■ Setup



- 2 paths;
100 Mbit/s, 1 ms, 0% packet loss
- Saturated sender
- Bandwidth variation on path 2

■ Simulation results


- Sctp, primary path via Path 1
- Sctp, primary path via Path 2
- **Expected for CMT-SCTP**
- Original CMT-SCTP



Lesson learned: interaction among mechanisms

Efficient CMT-SCTP over Dissimilar Paths

- **Complexity** due to **interaction** of different mechanisms
 - Congestion control
 - Management of send and receive buffers
 - Handling of retransmissions
 - Acknowledgement mechanism
 - Options for message delivery (in-sequence? lossless?)
- Need for a lot of research



How „good“ are the simulations? Comparison to real system!

Overview:

Lab Setups and Internet Setups

- About Norway and the Simula Research Laboratory
- From TCP to Multi-Path Transport
- Simulations
- Lab Setups and Internet Setups
- The NorNet Testbed
- Conclusion
- Literature

Research (2) – The Next Step: Lab Setup

- Surprisingly big effort:
 - Strange effects of cheap network components:
„It's only cheap on the paper!“
 - Debugging of SCTP in FreeBSD
- But valuable:
 - **The simulations were useful!** 😊
 - Bugfixes for the FreeBSD community
 - Open Source software „**NetPerfMeter**“
 - **Learning effects** and **new ideas!**



Internet protocols → testbed in the Internet!

Research (3) – Real Internet: 3 Cities and 2 Continents



- 3 connected lab setups
 - Establishment of an international cooperation
 - **Essen, Burgsteinfurt** (FH Münster), **Haikou 海口** (Hainan University)
- Very interesting scenario:
 - CMT-SCTP and MPTCP evaluation
 - Very different path characteristics
 - Ideas for further experiments

Many new ideas!

Now really big: NorNet testbed!

The NorNet Testbed

- NorNet Core
 - Cable, up to 4 providers, IPv4+IPv6 (fibre, „consumer-grade” DSL, etc.)
 - Hosts for virtual machines
 - 23 locations (11 in Norway, 12 abroad)

[**simula** . research laboratory]

- NorNet Edge
 - Embedded system „Ufoboard“
 - Up to 4x 2G/3G/4G, 1x CDMA, 1x Ethernet
 - Hundreds of locations (in Norway)



<https://www.nntb.no>



Overview:

The NorNet Testbed

- About Norway and the Simula Research Laboratory
- From TCP to Multi-Path Transport
- Simulations
- Lab Setups and Internet Setups
- The NorNet Testbed
- Conclusion
- Literature

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
- **Deployment:**
 - Core: 11 sites in Norway + CN, DE, SE, US, KR, AU, FR
 - Edge: hundreds of nodes in Norway



Idea for NorNet Core: 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
 - Each site's router (called **tunnelbox**) maintains the tunnels
 - Static tunnels
 - NorNet-internal addressing and routing over tunnels



Address Assignment

- NorNet-internal address spaces:
 - Private NorNet-internal IPv4 “/8” address space (NAT to outside)
 - Public NorNet-internal IPv6 “/48” address space
- Systematic address assignment:
 - IPv4: 10.<Provider ID>.<Site ID>.<Node ID>/24 per site
 - IPv6: 2001:700:4100:<PP><SS>::<NN>/64
(PP=Provider ID; SS=Site ID; NN=Node ID)
- NorNet-internal DNS setup including reverse lookup

Make it as easy as possible to keep the overview!

A NorNet Core Site Deployment

A usual NorNet Core site:



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

Additional researcher-provided sites:

- Varying configurations
- VM setups, powerful servers, “retro-style” PCs ...



Longyearbyen 78.2°N,15.6°E



NorNet Core på Svalbard

NorNet Core Site Deployment Status (December 2016)

No.	Site	ISP 1	ISP 2	ISP 3	ISP 4
1	Simula Research Laboratory	Uninett	Kvantel	Telenor	PowerTech
2	Universitetet i Oslo	Uninett	Broadnet	PowerTech	
3	Høgskolen i Gjøvik	Uninett	PowerTech		
4	Universitetet i Tromsø	Uninett	Telenor	PowerTech	
5	Universitetet i Stavanger	Uninett	Altibox	PowerTech	
6	Universitetet i Bergen	Uninett	BKK		
7	Universitetet i Agder	Uninett	PowerTech		
8	Universitetet på Svalbard	Uninett	Telenor		
9	Universitetet i Trondheim	Uninett	PowerTech		
10	Høgskolen i Narvik	Uninett	Broadnet	PowerTech	
11	Høgskolen i Oslo og Akershus	Uninett	–		
12	Karlstads Universitet	SUNET			
13	Universität Kaiserslautern	DFN			
14	Universität Duisburg-Essen	DFN	Versatel		
15	Hainan University 海南大学	CERNET	China Unicom		
16	The University of Kansas	KanREN			
17	Korea University 고려대학교	KREONET			
18	National ICT Australia (NICTA)	AARNet			
19	HAW Hamburg	DFN			
20	Technische Universität Darmstadt	DFN			
21	Haikou Cg. of Econ. 海口经济学院	China Telecom	CERNET		

IPv4 and IPv6
IPv4 only (ISP without IPv6 support ☹)
IPv4 only (site's network without IPv6 support)
ISP negotiation in progress

<https://www.nntb.no/pub/nornet-configuration/NorNetCore-Sites.html>

Some Site Statistics (February 2017)



Active Sites	23
Distinct ISPs of Active Sites	18
Distinct Countries of Active Sites	8
Total IPv4 Interfaces	42
Total IPv4 Tunnels	861
Total IPv6 Interfaces	26
Total IPv6 Tunnels	325

<https://www.nntb.no/pub/nor-net-configuration/NorNetCore-Sites.html>

Remote Systems

❄ -30°C
Longyearbyen

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 14.04 LTS
 - KVM (Kernel-based Virtual Machine)
 - Other software runs in 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



You may use NorNet Core, too!

**Join the tutorial session today!
Here, at the Universidad de Castilla-La Mancha!**

- Contents:
 - Get access to NorNet Core
 - User and slice management
 - Access to slices
 - Using and configuring slivers with own software
 - How to make use of multi-homing?

Overview: Conclusion

- About Norway and the Simula Research Laboratory
- From TCP to Multi-Path Transport
- Simulations
- Lab Setups and Internet Setups
- The NorNet Testbed
- Conclusion
- Literature

Conclusion and Future Work

“The road to hell is paved with unused testbeds.”
[James P. G. Sterbenz]

- The NorNet Core testbed is ready for experiments!
 - Do you have experiment ideas? → Talk to us!
- Future work:
 - NorNet Core
 - Additional sites, more IPv6 endpoints
 - **OpenStack support**
 - Improve and refine management software
 - Get more users, may be you?



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

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



<https://www.nntb.no>

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

Overview: Literature

- About Norway and the Simula Research Laboratory
- From TCP to Multi-Path Transport
- Simulations
- Lab Setups and Internet Setups
- The NorNet Testbed
- Conclusion
- Literature

Literature (1)

- Wang, K.; Dreibholz, T.; Zhou, X.; Fu, F.; Tan, Y.; Cheng, X.; Tan, Q.: “[On the Path Management of Multi-Path TCP in Internet Scenarios based on the NorNet Testbed](#)” (PDF, 379 KiB), Proceedings of the IEEE International Conference on Advanced Information Networking and Applications (AINA), Taipei, Taiwan/People's Republic of China, March 27, 2017.
- Zhou, F.; Dreibholz, T.; Zhou, X.; Fu, F.; Tan, Y.; Gan, Q.: “[The Performance Impact of Buffer Sizes for Multi-Path TCP in Internet Setups](#)” (PDF, 1006 KiB), Proceedings of the IEEE International Conference on Advanced Information Networking and Applications (AINA), Taipei, Taiwan/People's Republic of China, March 27, 2017.
- Dreibholz, T.: “An Experiment Tutorial for the NorNet Core Testbed at the the Universidad de Castilla-La Mancha”, Tutorial at the Universidad de Castilla-La Mancha, Instituto de Investigación Informática de Albacete, Albacete, Castilla-La Mancha/Spain, February 17, 2017.
- Amer, P. D.; Becke, M.; Dreibholz, T.; Ekiz, N.; Iyengar, J. R.; Natarajan, P.; Stewart, R. R.; Tüxen, M.: “[Load Sharing for the Stream Control Transmission Protocol \(SCTP\)](#)” (TXT, 62 KiB), IETF, Individual Submission, Internet Draft draft-tuexen-tsvwg-sctp-multipath-13, December 22, 2016.
- Dreibholz, T.; Ferlin, S.; Alay, Ö.; Elmokashfi, A. M.; Livadariu, I. A.; Zhou, X.: “[MPTCP Experiences in the NorNet Testbed](#)” (TXT, 21 KiB), IETF, Individual Submission, Internet Draft draft-dreibholz-mptcp-nor-net-experience-02, March 19, 2016.
- Fu, F.; Zhou, X.; Dreibholz, T.; Wang, K.; Zhou, F.; Gan, Q.: “[Performance Comparison of Congestion Control Strategies for Multi-Path TCP in the NorNet Testbed](#)” (PDF, 172 KiB), Proceedings of the 4th IEEE/CIC International Conference on Communications in China (ICCC), pp. 607–612, ISBN 978-1-5090-0243-6, Shenzhen, Guangdong/People's Republic of China, November 3, 2015.
- Dreibholz, T.; Zhou, X.; Fu, F.: “[Multi-Path TCP in Real-World Setups – An Evaluation in the NorNet Core Testbed](#)” (PDF, 211 KiB), 5th International Workshop on Protocols and Applications with Multi-Homing Support (PAMS), pp. 617–622, DOI [10.1109/WAINA.2015.74](#), ISBN 978-1-4799-1775-4, Gwangju/South Korea, March 25, 2015.
- Ferlin, S.; Dreibholz, T.; Alay, Ö.: “[Multi-Path Transport over Heterogeneous Wireless Networks: Does it really pay off?](#)” (PDF, 986 KiB), Proceedings of the IEEE Global Communications Conference (GLOBECOM), pp. 4807–4813, DOI [10.1109/GLOCOM.2014.7037567](#), ISBN 978-1-4799-3512-3, Austin, Texas/U.S.A., December 11, 2014.

Literature (2)

- **Golkar, F.; Dreibholz, T.; Kvalbein, A.:** “**Measuring and Comparing Internet Path Stability in IPv4 and IPv6**” (PDF, 436 KiB), Proceedings of the 5th IEEE International Conference on the Network of the Future (NoF), pp. 1–5, DOI [10.1109/NOF.2014.7119767](https://doi.org/10.1109/NOF.2014.7119767), ISBN 978-1-4799-7531-0, Paris/France, December 4, 2014.
- **Ferlin, S.; Dreibholz, T.; Alay, Ö.:** “**Tackling the Challenge of Bufferbloat in Multi-Path Transport over Heterogeneous Wireless Networks**” (PDF, 820 KiB), Proceedings of the IEEE/ACM International Symposium on Quality of Service (IWQoS), pp. 123–128, DOI [10.1109/IWQoS.2014.6914310](https://doi.org/10.1109/IWQoS.2014.6914310), ISBN 978-1-4799-4852-9, Hong Kong/People's Republic of China, May 26, 2014.
- **Gran, E. G.; Dreibholz, T.; Kvalbein, A.:** “**NorNet Core – A Multi-Homed Research Testbed**” (PDF, 1458 KiB), Computer Networks, Special Issue on Future Internet Testbeds, vol. 61, pp. 75–87, DOI [10.1016/j.bjp.2013.12.035](https://doi.org/10.1016/j.bjp.2013.12.035), ISSN 1389-1286, March 14, 2014.
- **Kvalbein, A.; Baltrūnas, D.; Evensen, K. R.; Xiang, J.; Elmokashfi, A. M.; Ferlin, S.:** “**The NorNet Edge Platform for Mobile Broadband Measurements**” (PDF, 2742 KiB), Computer Networks, Special Issue on Future Internet Testbeds, vol. 61, pp. 88–101, DOI [10.1016/j.bjp.2013.12.036](https://doi.org/10.1016/j.bjp.2013.12.036), ISSN 1389-1286, March 14, 2014.
- **Becke, M.; Adhari, H.; Rathgeb, E. P.; Fu, F.; Yang, X.; Zhou, X.:** “**Comparison of Multipath TCP and CMT-SCTP based on Intercontinental Measurements**” (PDF, 924 KiB), Proceedings of the IEEE Global Communications Conference (GLOBECOM), Atlanta, Georgia/U.S.A., December 10, 2013.
- **Ford, A.; Raiciu, C.; Handley, M.; Bonaventure, O.:** “**TCP Extensions for Multipath Operation with Multiple Addresses**” (TXT, 162 KiB), IETF, RFC 6824, DOI [10.17487/RFC6824](https://doi.org/10.17487/RFC6824), ISSN 2070-1721, January 2013.
- **Dreibholz, T.:** “**Evaluation and Optimisation of Multi-Path Transport using the Stream Control Transmission Protocol**” (PDF, 36779 KiB), University of Duisburg-Essen, Faculty of Economics, Institute for Computer Science and Business Information Systems, URN [urn:nbn:de:hbz:464-20120315-103208-1](https://nbn-resolving.org/urn:nbn:de:hbz:464-20120315-103208-1), March 13, 2012.
- **Dreibholz, T.; Becke, M.; Rathgeb, E. P.; Tüxen, M.:** “**On the Use of Concurrent Multipath Transfer over Asymmetric Paths**” (PDF, 1036 KiB), Proceedings of the IEEE Global Communications Conference (GLOBECOM), DOI [10.1109/GLOCOM.2010.5683579](https://doi.org/10.1109/GLOCOM.2010.5683579), ISBN 978-1-4244-5637-6, Miami, Florida/U.S.A., December 7, 2010.
- **Stewart, R. R.:** “**Stream Control Transmission Protocol**” (TXT, 338 KiB), IETF, RFC 4960, DOI [10.17487/RFC4960](https://doi.org/10.17487/RFC4960), ISSN 2070-1721, September 2007.