Applying TCP-Friendly Congestion Control to Concurrent Multipath Transfer

UNIVERSITÄT DUISBURG ESSEN

Thomas Dreibholz
Institute for Experimental Mathematics
University of Duisburg-Essen, Germany
dreibh@iem.uni-due.de

Table of Contents

UNIVERSITÄT DUISBURG ESSEN

- Basics
  - SCTP and Concurrent Multipath Transfer (CMT)
  - CMT Fairness Problem
  - Resource Pooling (RP)
- Our Approach: CMT/RP-SCTP
  - Congestion Control
  - Simulation Scenario
  - Proof of Concept
- Conclusion and Outlook

Thomas Dreibholz's SCTP Page
http://tdrwww.iem.uni-due.de/dreibholz/sctp/

Stream Control Transmission Protocol (SCTP, RFC 4960)

- SCTP Features
  - Transport-Layer Protocol (like TCP or UDP – but much more powerful!)
  - Reliable, message-oriented, ordered/unordered, multi-streaming
- Multi-Homing
  - Support for multiple addresses per endpoint; may be changed (“Add-IP”)
  - Multiple unidirectional paths in the network (can be disjoint or shared)
  - One path in each direction is chosen for user data (**primary path**)
  - Other paths: backup only (only used for retransmissions)

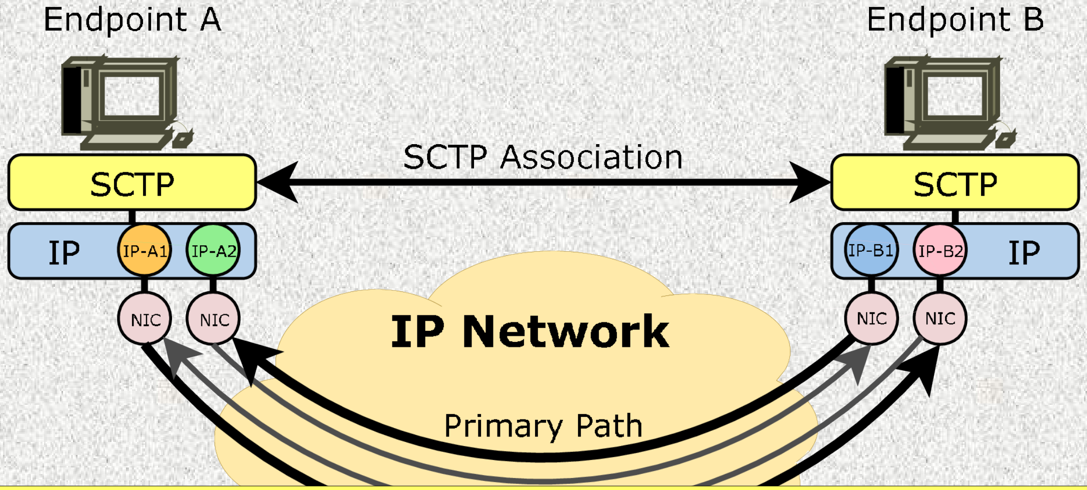

What about utilizing all paths simultaneously?

Concurrent Multipath Transfer (CMT)

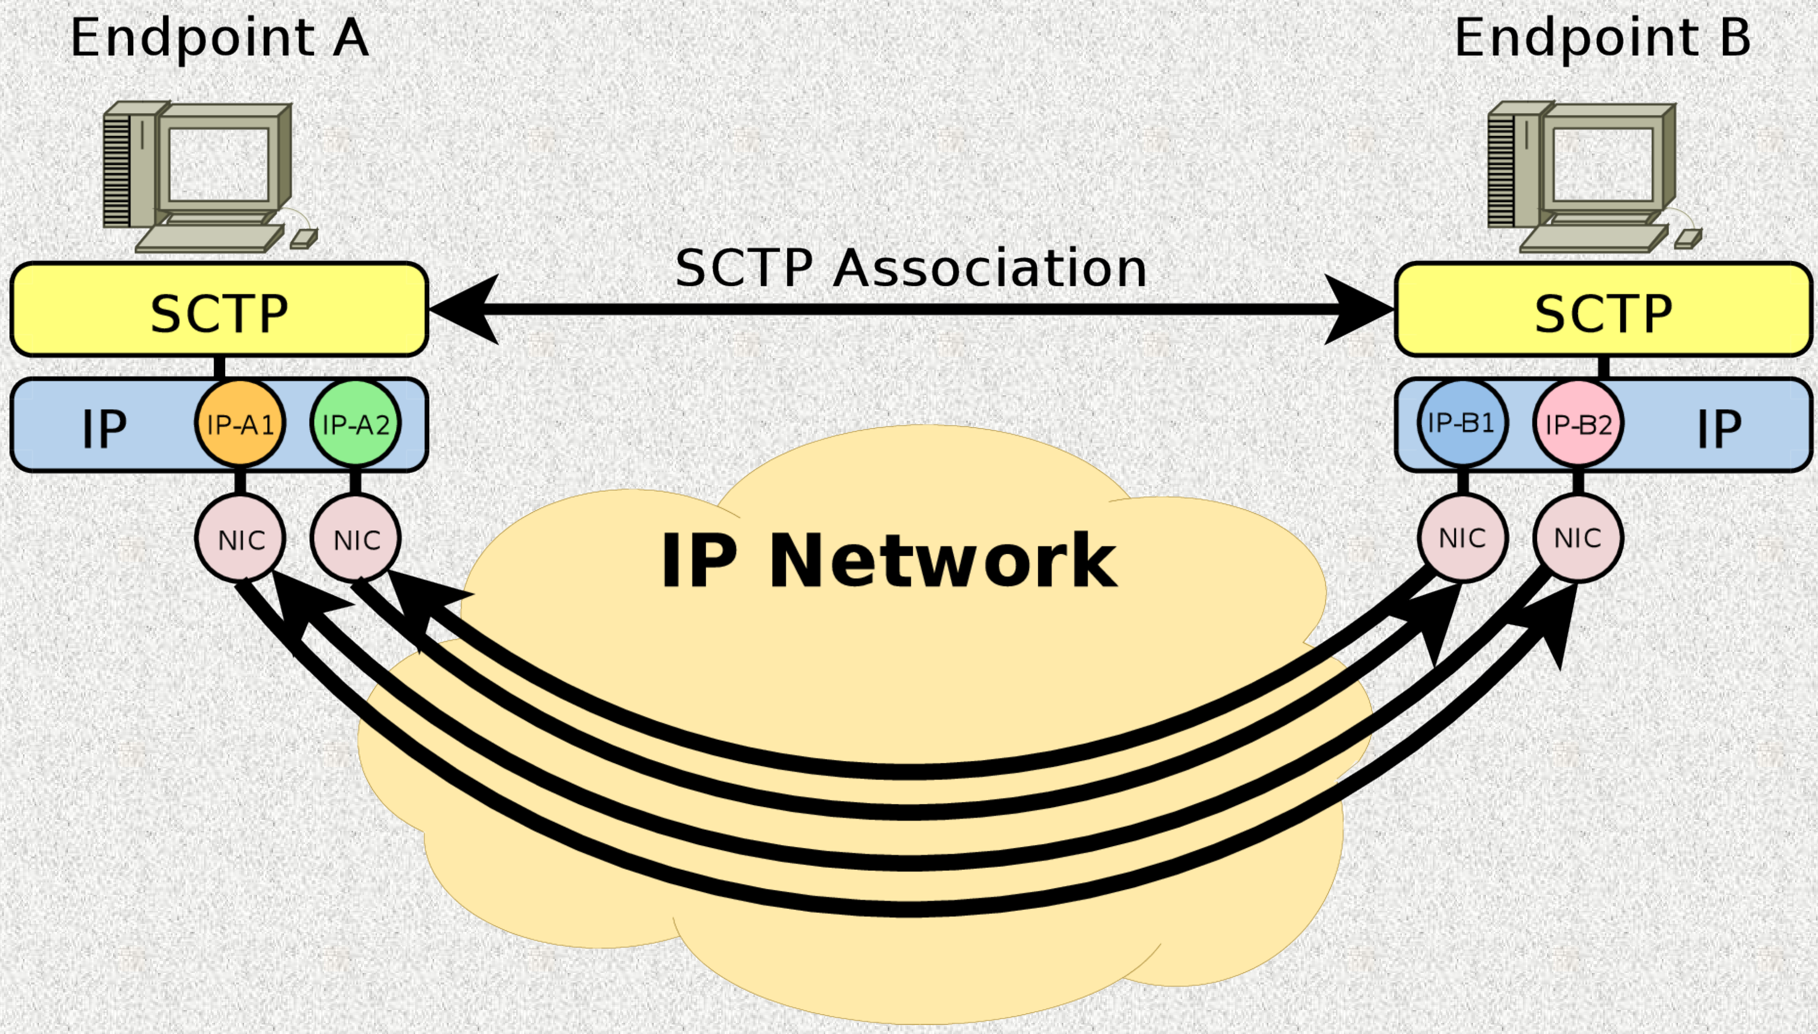

- All paths are used for data transmission
- Assumption of CMT: paths are disjoint → congestion control

CMT Fairness Problem

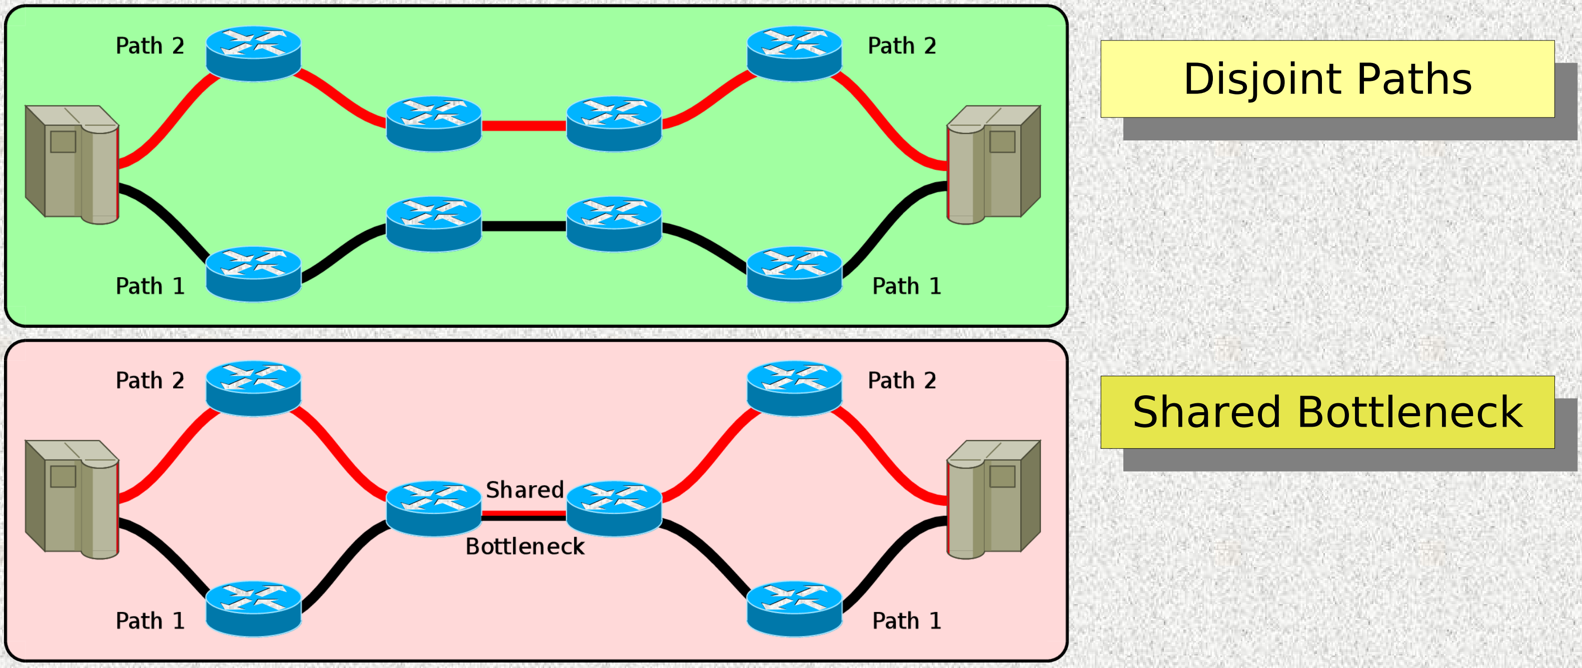

- CMT Congestion Control:
  - TCP-like congestion control *per path*
  - CMT flow behaviour:
    - On path #1: like a single TCP flow
    - On path #2: like a single TCP flow
    - On shared bottleneck: like two TCP flows
      - Doubled bandwidth in comparison to single TCP flow => unfair ☹
      - Cannot be detected reliably in arbitrary networks (e.g. the Internet)

Resource Pooling (RP)

UNIVERSITÄT DUISBURG ESSEN

- Definition of Resource Pooling (RP) from [WHB09]:
  - *"Making a collection of resources behave like a single pooled resource"*
  - Principle can be applied in many cases:
    - Statistical multiplexing
    - Failure resilience
    - ...
    - Load balancing
- Applying RP for multipath transfer:
  - Do not handle each path independently
  - Instead, let the paths of a multi-homed flow behave like one big path ...
  - ... by using a congestion control which is aware of path interaction
  - Idea: combining CMT-SCTP with RP to solve the fairness problem!

How to realize a RP-enabled congestion control for CMT-SCTP?

CMT/RP-SCTP – Combining CMT-SCTP with RP

UNIVERSITÄT DUISBURG ESSEN

- Definitions:
  - $S=\sum s_i$ – the sum of slow start thresholds s_i on path i
  - $C=\sum c_i$ – the sum of congestion windows c_i on path i
- Idea: slow start threshold ratio $\frac{S_P}{S}$ as capacity measure for path P
- Congestion window growth on path P
  - Slow Start: on α acknowledged bytes in fully-utilized congestion window

$$c_P = c_P + \lceil min(MTU_P, \alpha) * \frac{S_P}{S} \rceil$$

  - Congestion Avoidance: on fully-acknowledged congestion window

$$c_P = c_P + \lceil MTU_P * \frac{S_P}{S} \rceil$$

- Congestion window decrease on path P
  - On Fast Retransmission:

$$s_P = max(c_P - \frac{C}{2}, 4 * MTU_P * \frac{s_P}{S}, MTU_P); c_P = s_P$$

  - On Timer-Based Retransmission:

$$s_P = max(c_P - \frac{C}{2}, 4 * MTU_P * \frac{s_P}{S}, MTU_P); c_P = MTU_P$$

Simulation Setup

- **Scenario:**

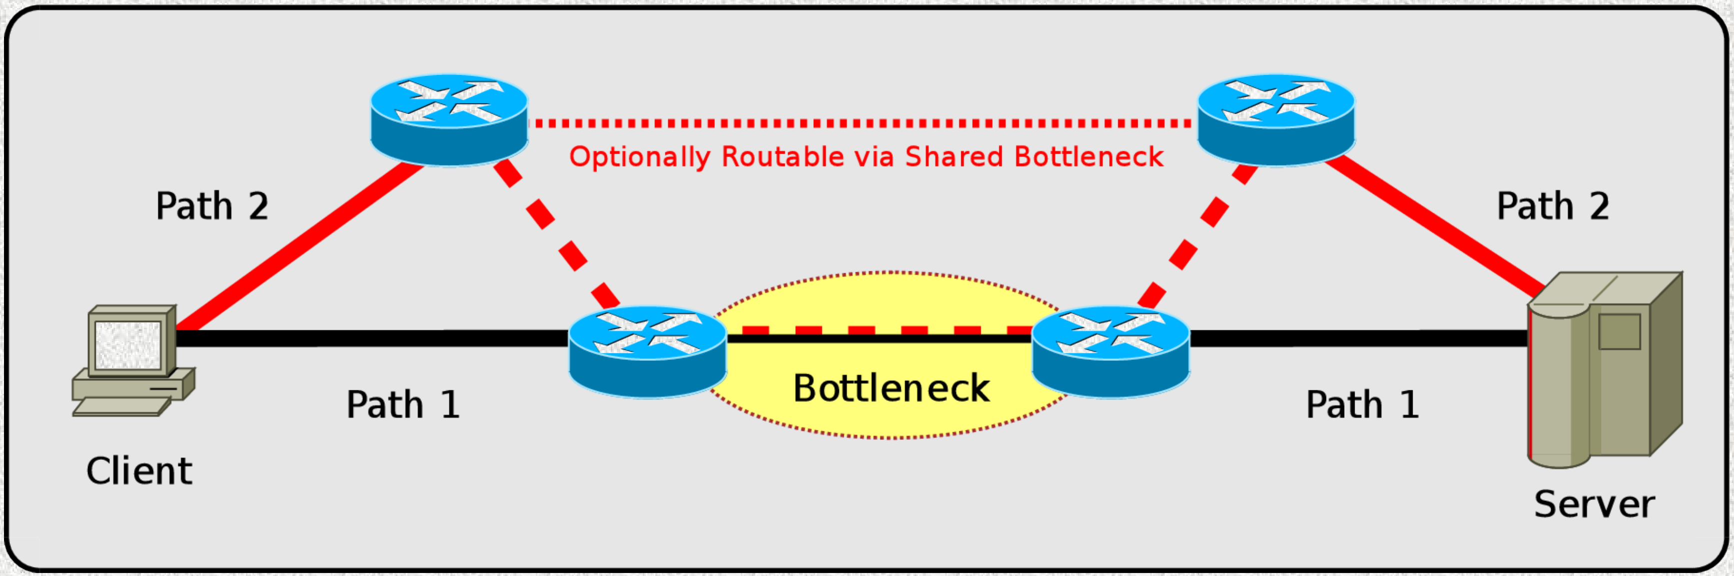

- **Parameters:**
  - Routes switchable: disjoint paths *or* shared bottleneck
  - Saturated senders
  - Background flow(s): CMT-SCTP or CMT/RP-SCTP
  - Reference flow: non-CMT SCTP flow for comparison
    - Goal: this flow should get a fair bandwidth share, of course

Throughput for Exclusive Usage of Two Disjoint 100 Mbit/s Paths

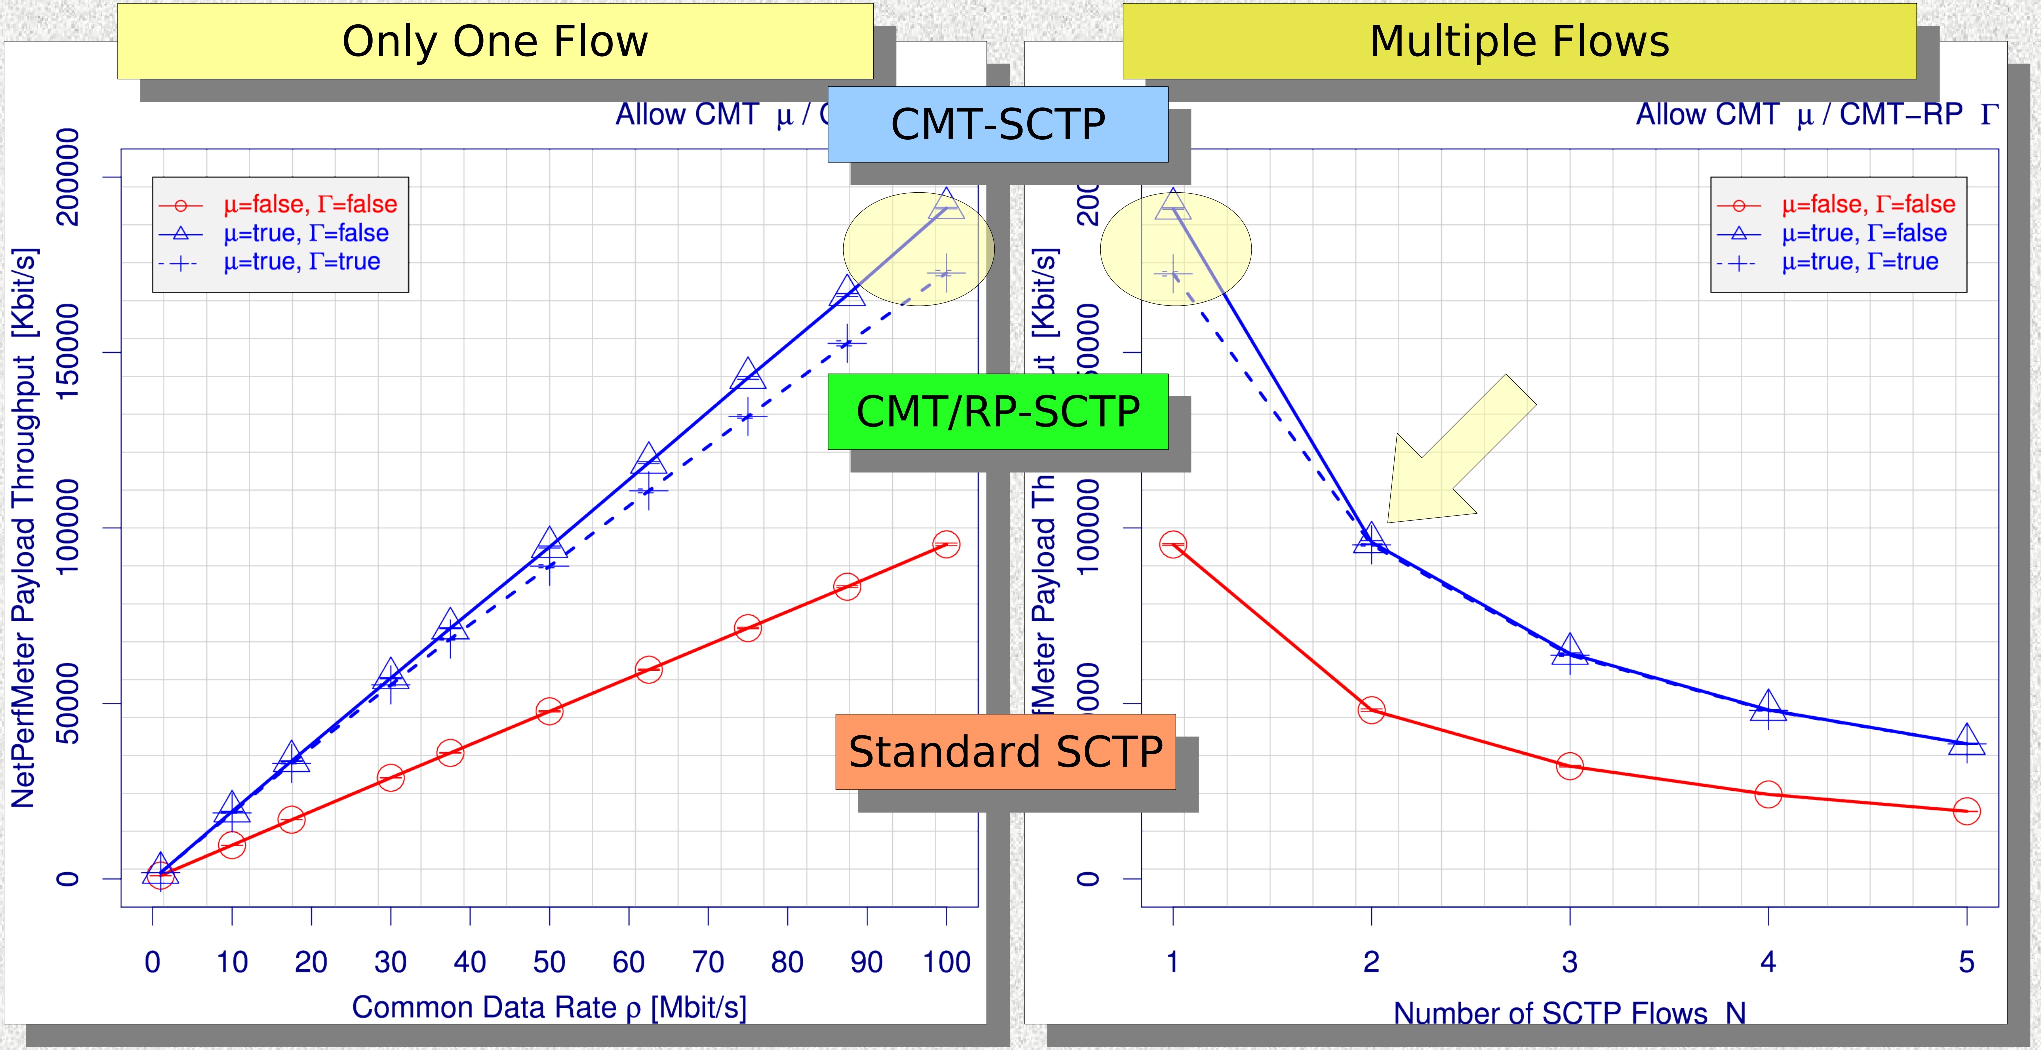

- Thoughput of <u>single</u> CMT/RP-SCTP flow is a little lower than CMT-SCTP
- ... but no significant difference for multiple flows ...

Congestion Window Examples

UNIVERSITÄT DUISBURG ESSEN

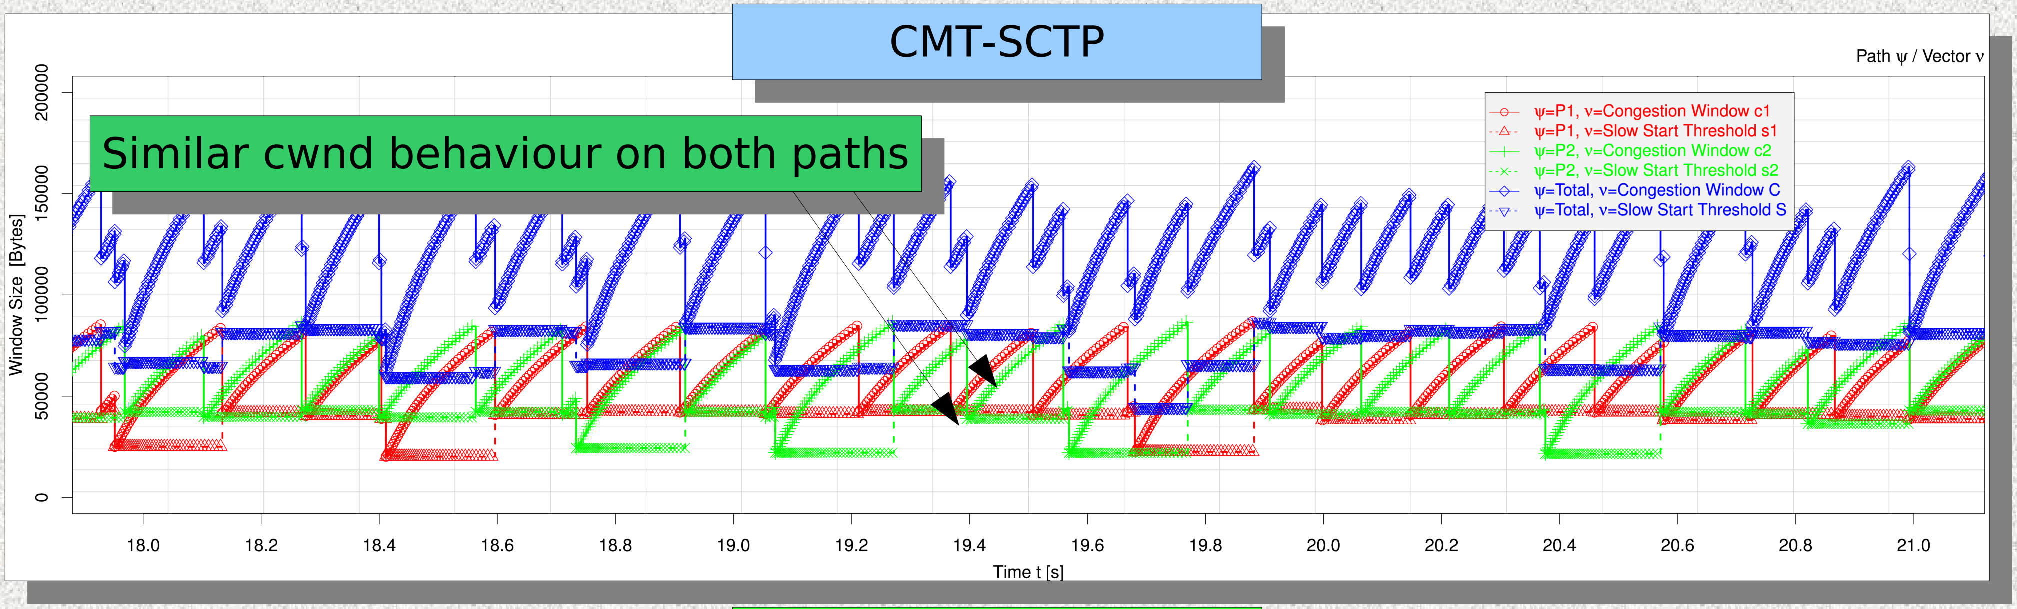

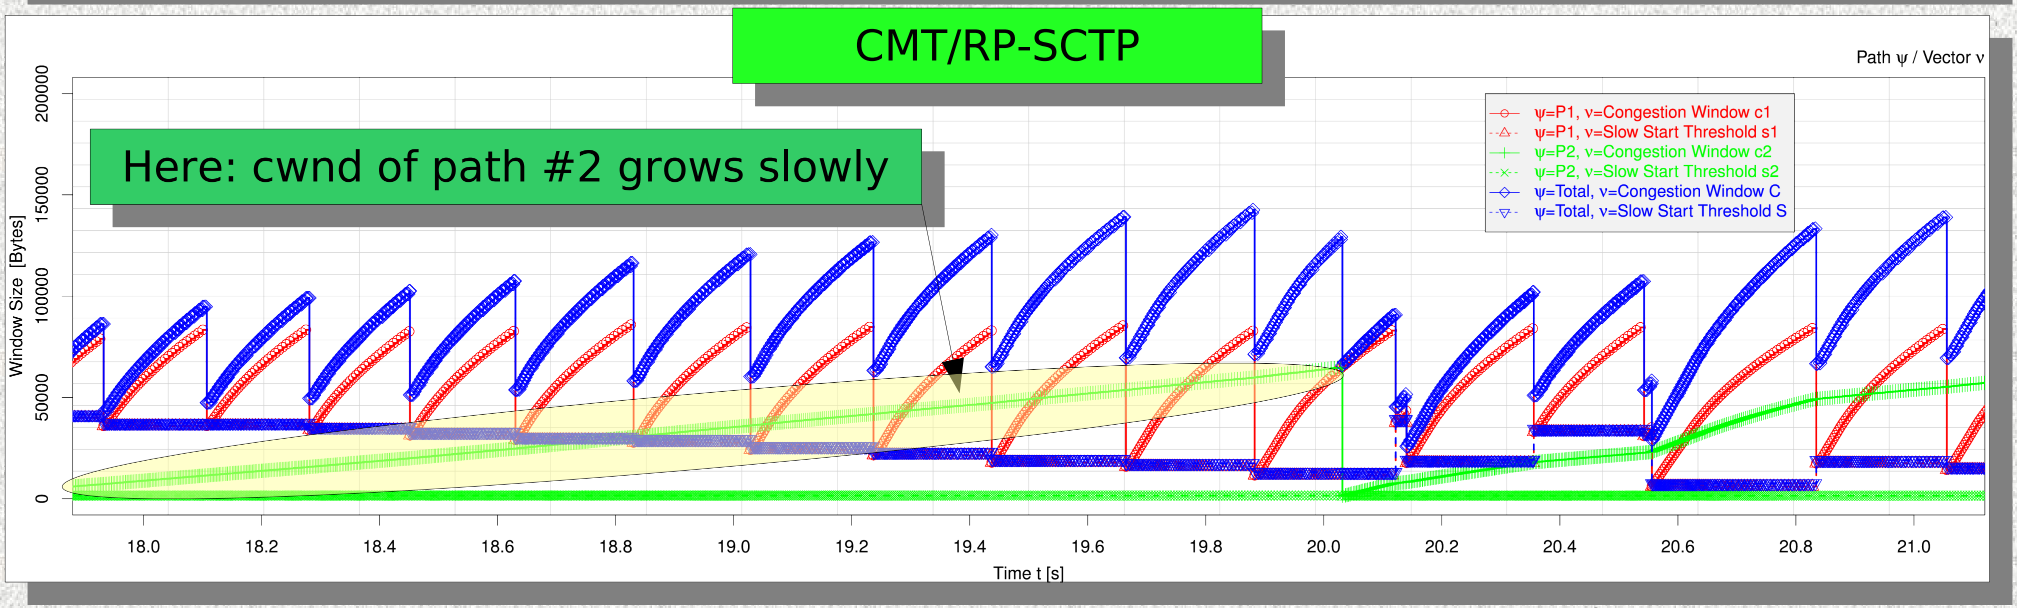

Flow Concurrency in 100/100 Mbit/s Setup: Throughput of the Reference Flow

UNIVERSITÄT DUISBURG ESSEN

Concurrency on Disjoint Paths

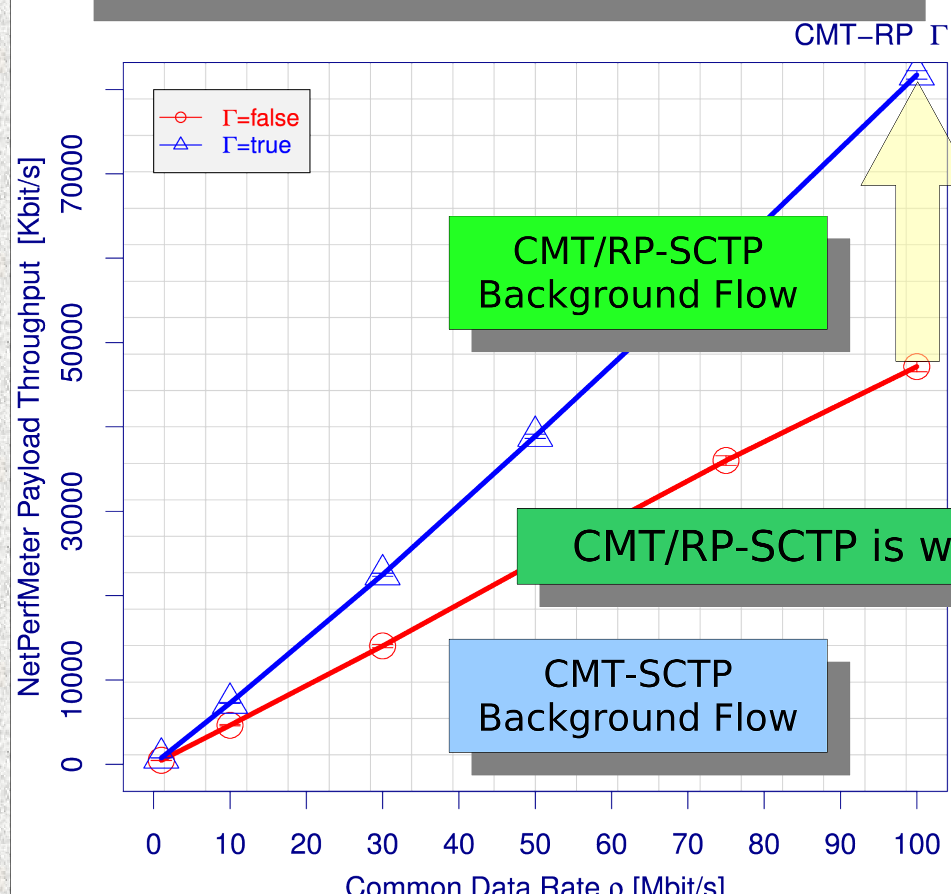

Concurrency on Shared Bottleneck

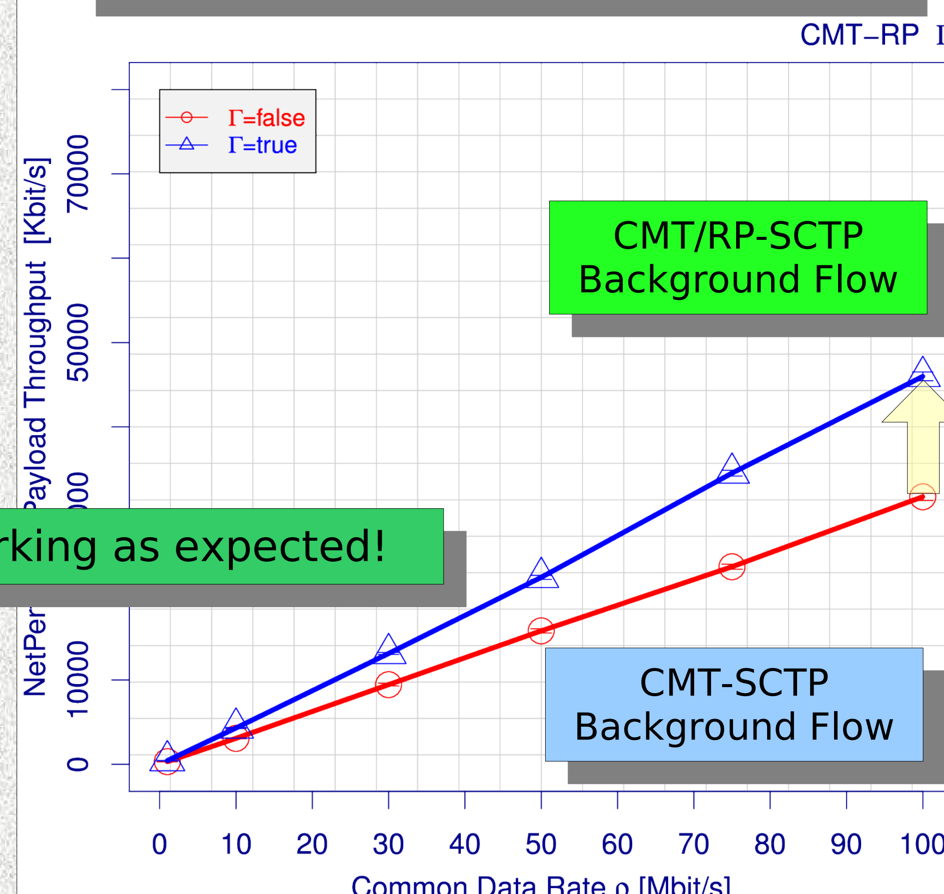

CMT/RP-SCTP is working as expected!

- Disjoint paths: CMT/RP flow takes less bandwidth on shared path
  - Since it already gets 100% of the bandwidth on the non-shared path
- Shared bottleneck: non-CMT flow gets 50% of the bandwidth

Conclusion and Outlook

UNIVERSITÄT DUISBURG ESSEN

- Conclusion
  - CMT-SCTP – Concurrent Multipath Transfer with SCTP
    - Unfair to concurrent non-CMT flows on shared bottlenecks
  - Resource Pooling
    - Take care of congestion interaction among paths
  - Our approach: CMT/RP-SCTP – CMT-SCTP with Resource Pooling
  - Proof of concept by simulations

- Future Work
  - Prototype implementation into FreeBSD networking stack
  - Performance for asymmetric paths
  - Contributions to IETF standardization process

Thank You for Your Attention!
Any Questions?

To be continued ...

Visit Our Project Homepage:

http://tdrwww.iem.uni-due.de/dreibholz/sctp

Thomas Dreibholz, dreibh@iem.uni-due.de