Extracting Performance Metrics from NetFlow in Enterprise Networks

2nd EMANICS Workshop on NetFlow/IPFIX Usage

Jochen Kögel
jochen.koegel@ikr.uni-stuttgart.de

8. October 2009

Universität Stuttgart
Institute of Communication Networks
and Computer Engineering (IKR)
Prof. Dr.-Ing. Andreas Kirstädter

Overview

- Motivation and scenario
  - Monitoring in enterprise networks
  - What could be extracted from NetFlow?
- Metrics extraction process
- Evaluation and results
- Conclusion and outlook

Motivation and scenario

Monitoring in global enterprise networks

- Global MPLS-Cloud connects several locations
- Network metrics (RTT, delay, loss,...) monitored by active probes (partly, no full mesh)
- Unsampled NetFlow (v5) from many routers (own routers + customer edge)
- Application level: Response time measured by active probing (E2E-probes)

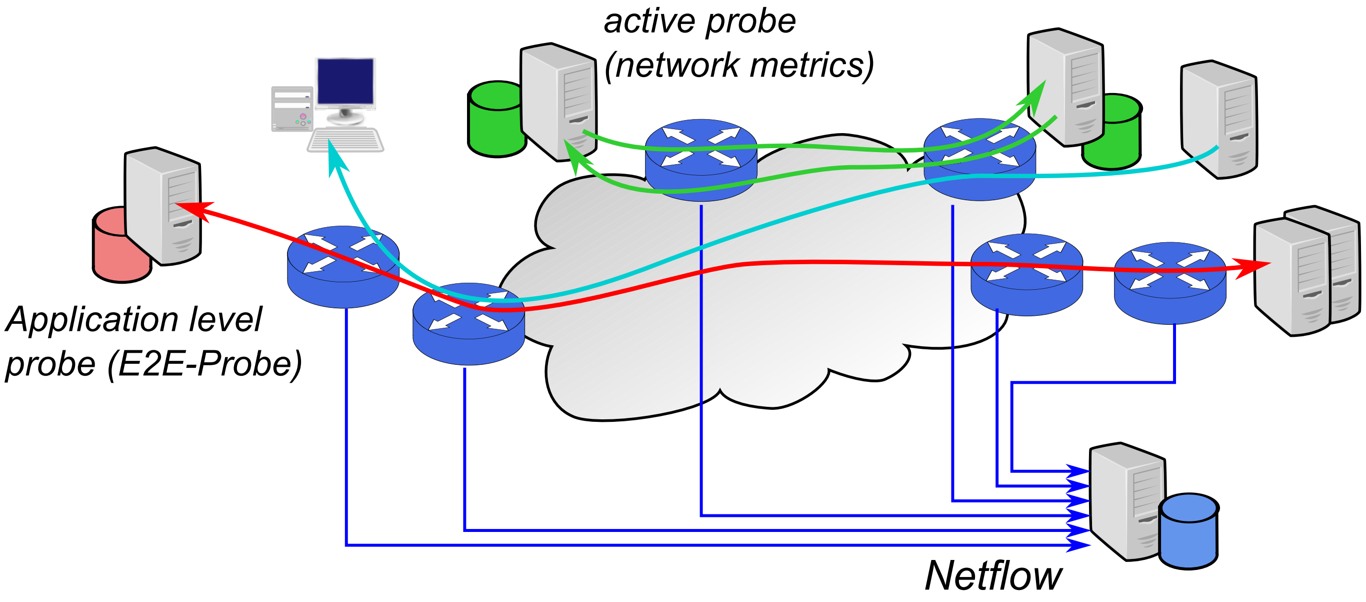

→ Correlating network metrics with application response times: NetFlow-based?

→ Extract metrics from NetFlow-Data to enrich/validate/replace active measurements?

Motivation and scenario

What could be extracted from NetFlow Records?

- Round-Trip-Time (RTT) – data from one router (depends on routing)
- One way delay – data from several routers (+ synchronized clocks)
- Packet loss
- Flow contention (roughly)

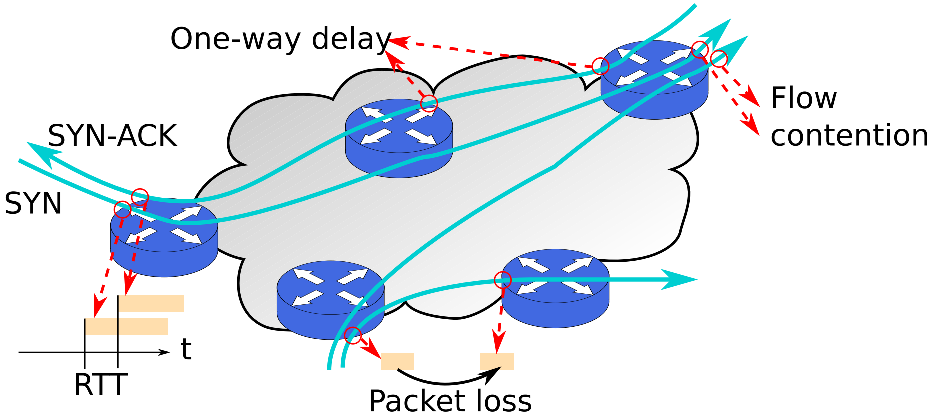

Partly covered in QoS-Monitoring-Section of RFC 5472 (IPFIX Applicability), but no investigation on flow-level so far (?)

Constraints

- Incomplete NetFlow data (table contention, packet loss)
- Rerouting, ECMP, disjoint paths

Metrics extraction process

Preprocessing Steps

1. Join flow records of same forward flow based on 5-Tuple (`JoinedFlow`)
2. Build `FlowAcrossExporters`: associate `JoinedFlows` of all exporters
3. Associate forward and reverse flows (`BiFlow`)

RTT Extraction

1. Take complete `BiFlows`
2. Calculate start-flow record offset (mid-flow offset also contains server response time)

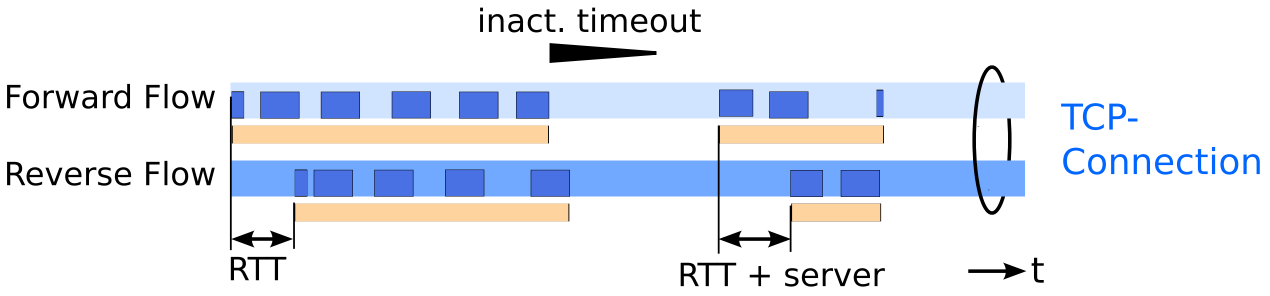

Delay Extraction

1. Step through `JoinedFlows` of `FlowAcrossExporters`
2. Calculate delay between every exporter pair based on start time

Evaluation and results

Evaluation scenario

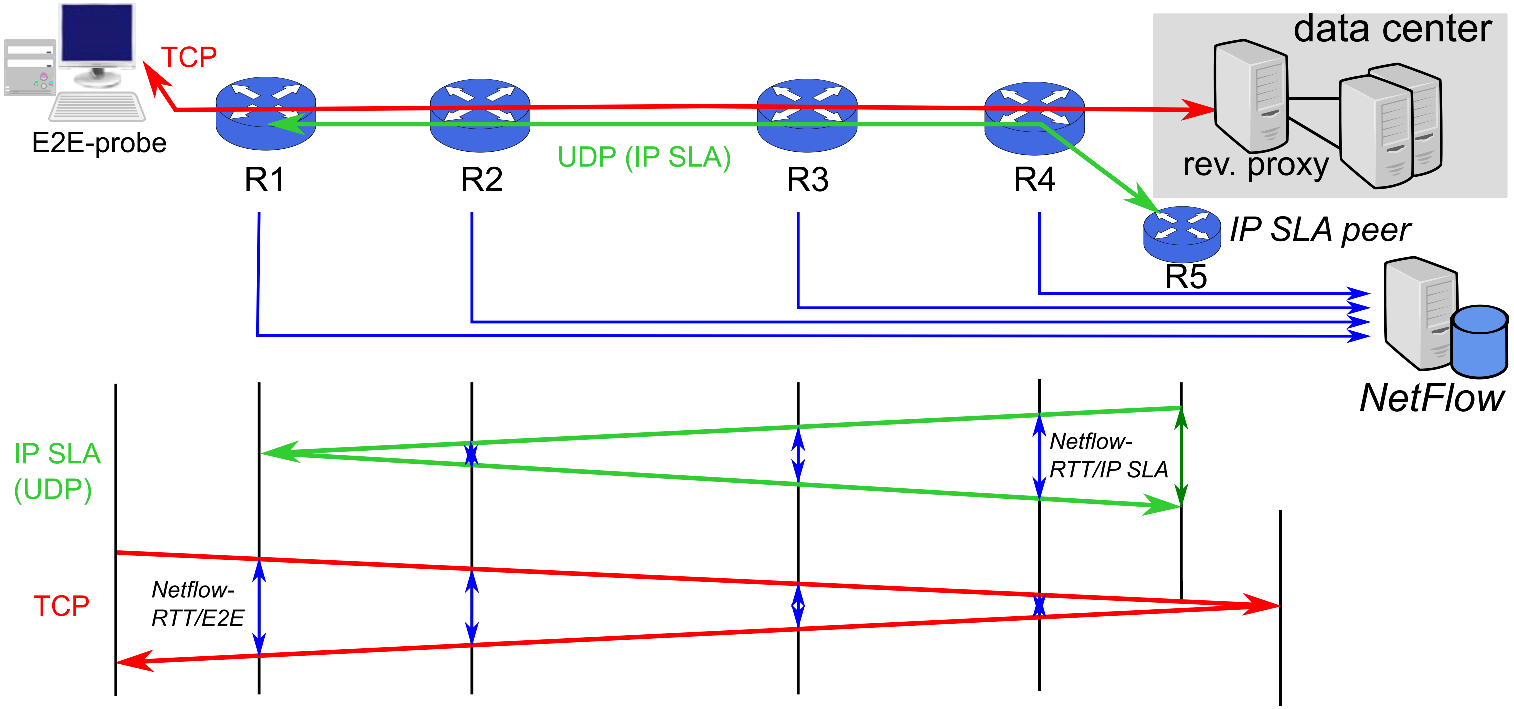

NetFlow data

- NetFlow data of 3 days taken for this evaluation
- Filter: TCP-Connections of E2E-Probe and IP SLA flows

IP SLA (active measurement between routers)

- Three UDP measurement flows per minute
- Reports on average RTT every five minutes

Evaluation and results

Evaluation 1: NetFlow-RTT/E2E vs. IP SLA RTT

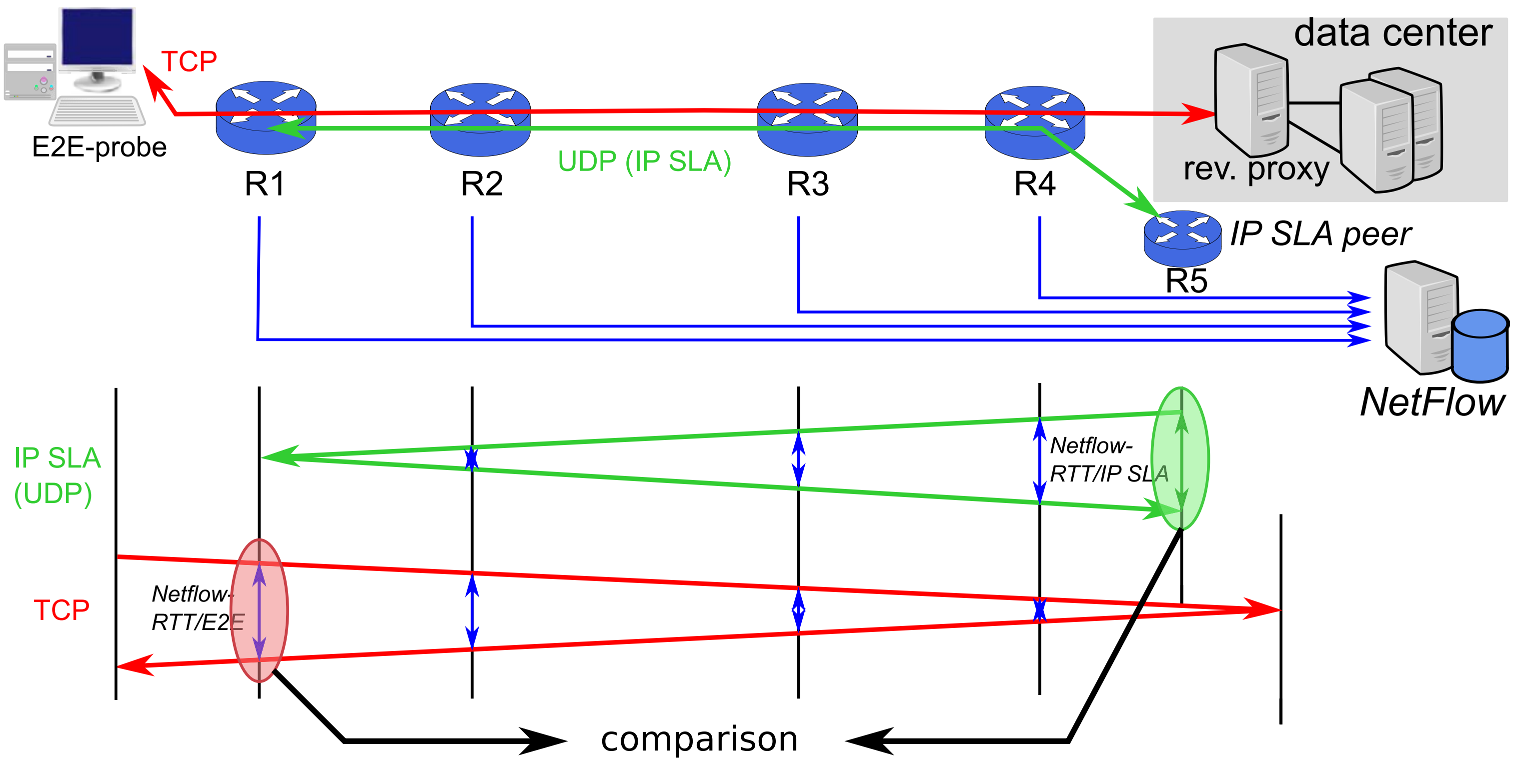

- Comparison of active measurement results and NetFlow-RTT
- RTT measurements in different directions, but same path

Evaluation and results

Evaluation 1: NetFlow-RTT/E2E vs. IP SLA RTT

IPSLA vs. TCP-Flows of E2E-Probes

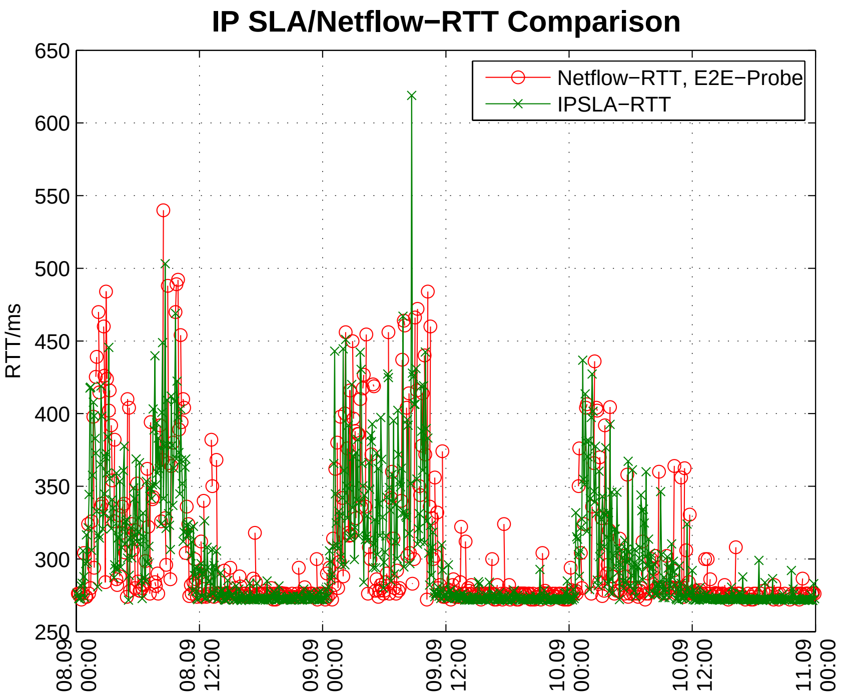

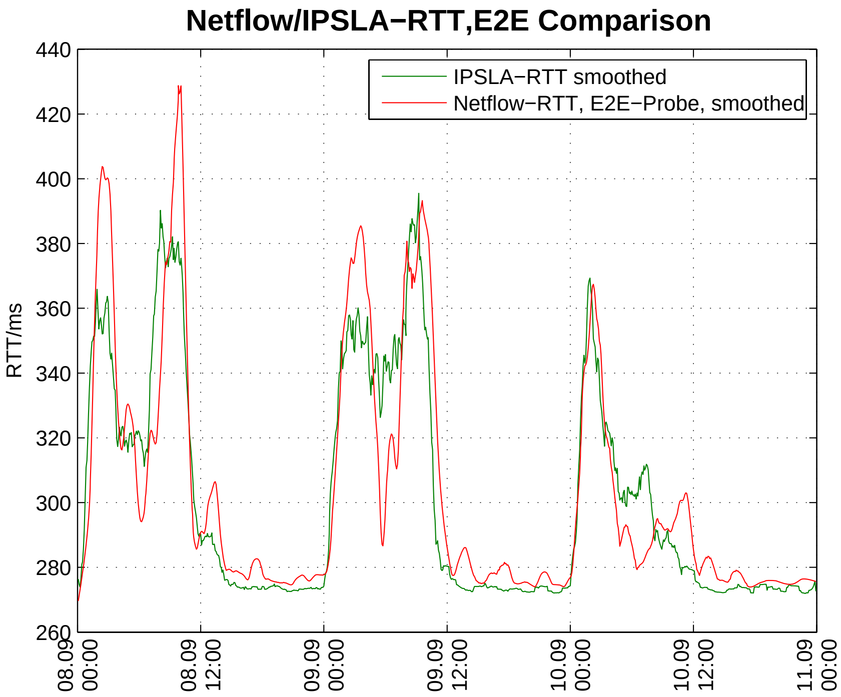

- Left: without further processing, right: smoothed (window 15)
- 1-4 NetFlow-RTT/E2E samples per IPSLA-sample

→ NetFlow-RTT timestamp may differ several seconds from IPSLA-Measurement time

→ Closer look on flows from IPSLA-Measurement

Evaluation and results

Evaluation 2: NetFlow-RTT/IP SLA vs. IP SLA RTT

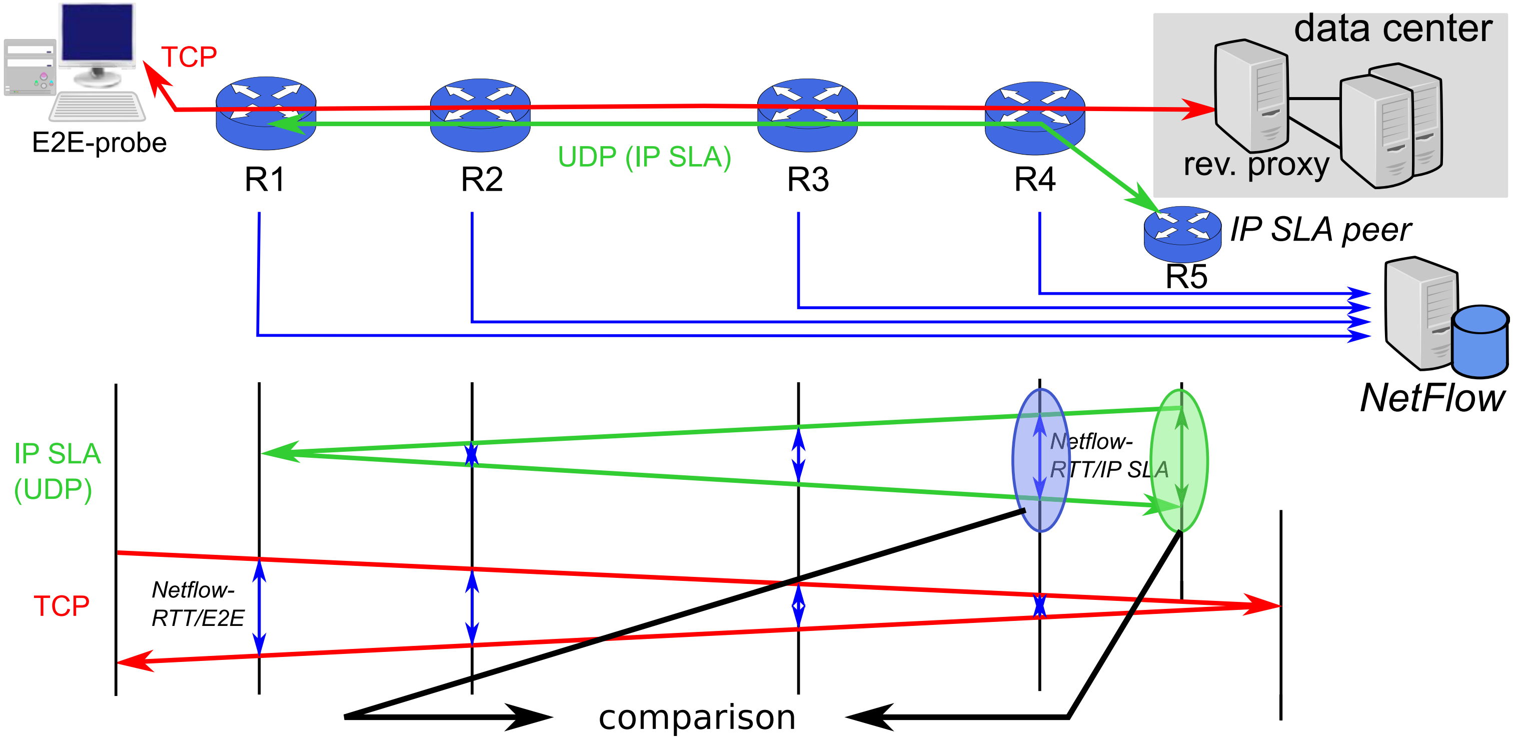

Compare RTT gained from IP SLA with RTT calculated from measurement flows

Evaluation and results

Evaluation 2: NetFlow-RTT/IP SLA vs. IP SLA RTT

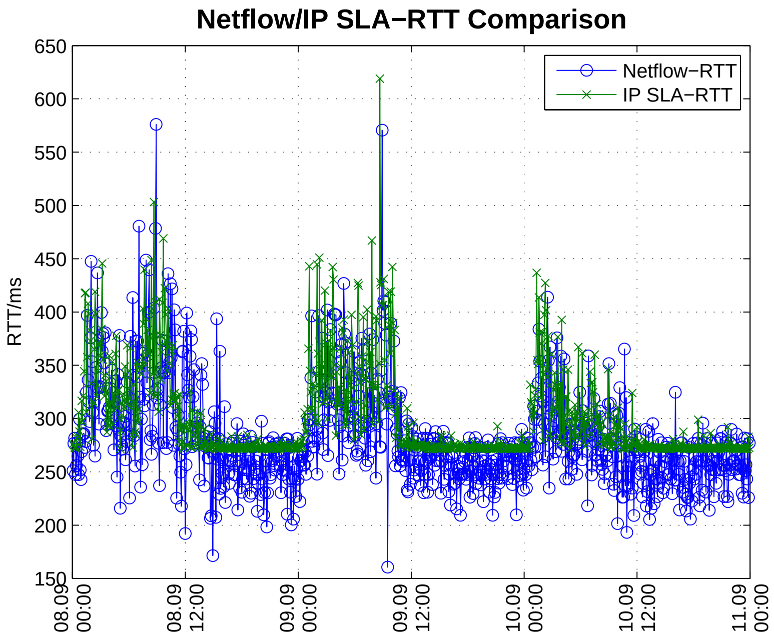

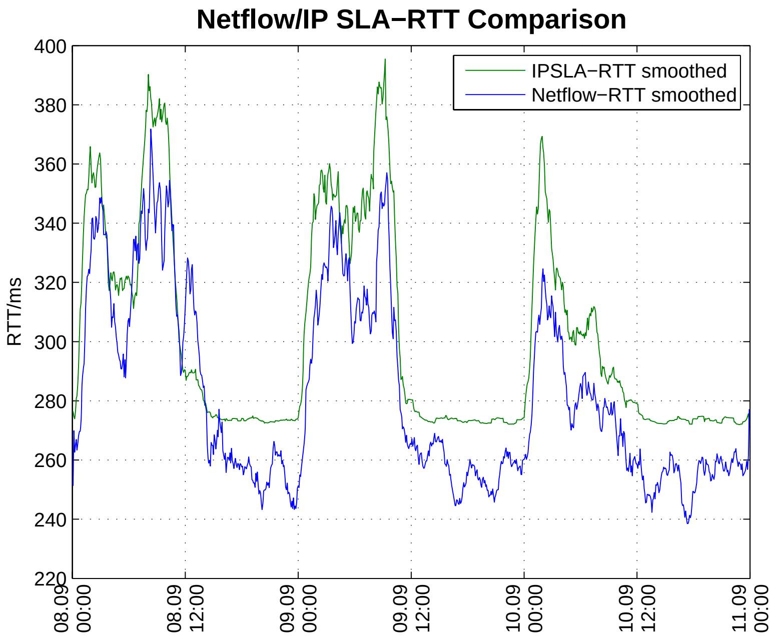

- Left: without further processing, right: smoothed (window 20)
- ~15 NetFlow-RTT samples per IP SLA sample
- IP SLA flows reuse 5-Tuple → matching problem → reason for difference?

Evaluation and results

Evaluation 3: NetFlow-RTT vs. NetFlow-delay

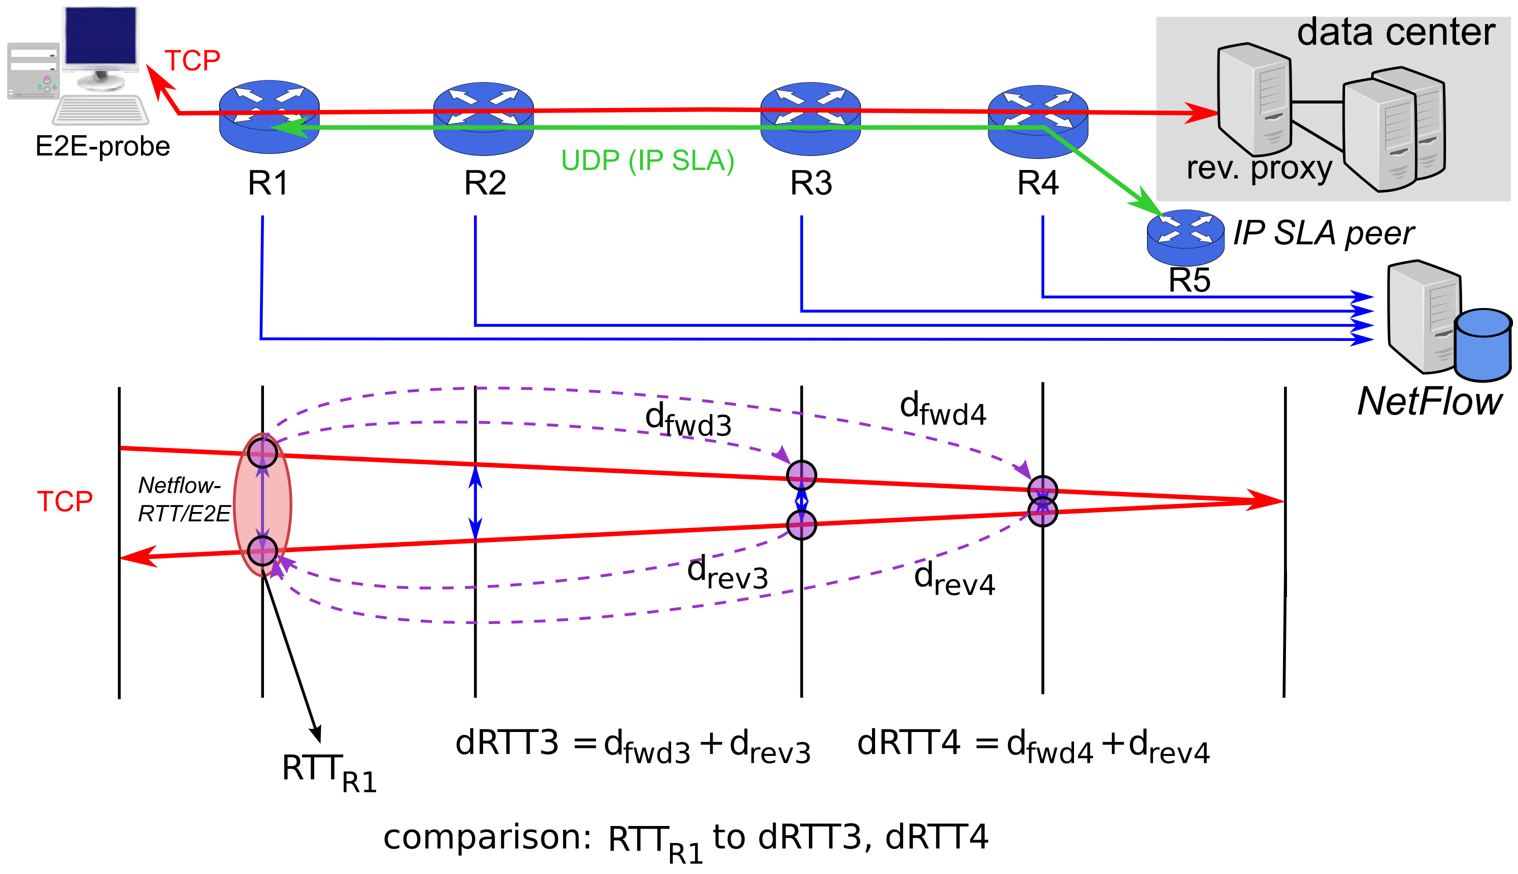

comparison: RTT_{R1} to dRTT3, dRTT4

1. Calculate NetFlow-delays independently of NetFlow-RTT
2. Match Netflow delay samples based on timestamp and sum up

→ Direct comparison possible (if R3 and R4 were inaccurate, we would notice)

Evaluation and results

Evaluation 3: NetFlow-RTT vs. NetFlow-delay

Delay of complete path: R1 - R4

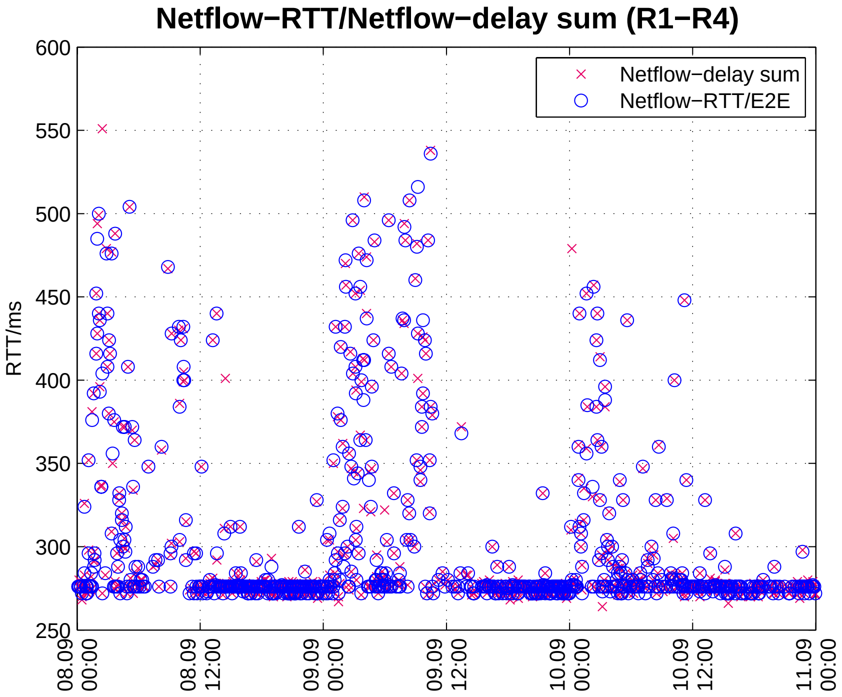
Netflow−RTT/Netflow−delay sum (R1−R4)

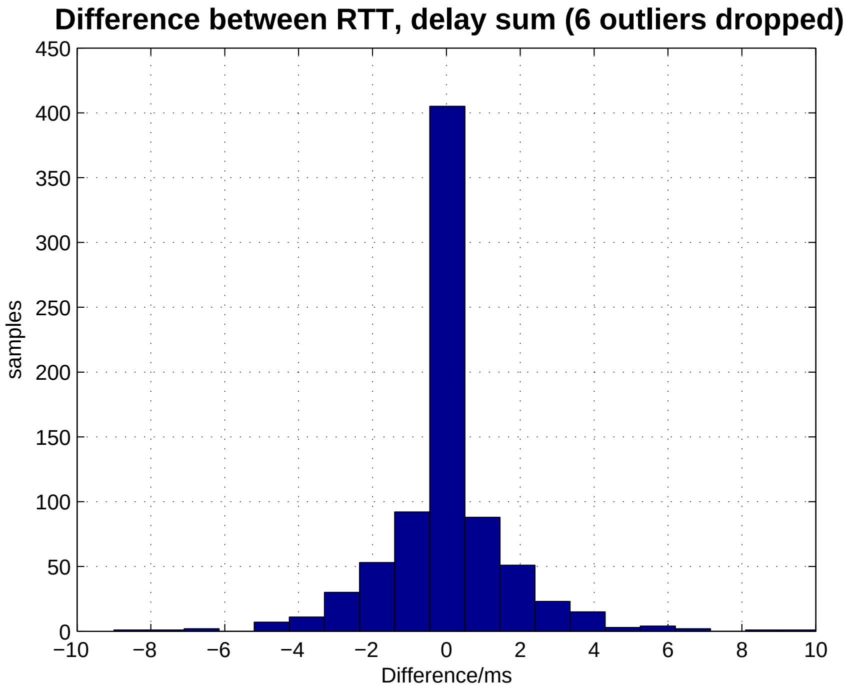
Difference between RTT, delay sum (6 outliers dropped)

- Sum of forward and reverse delay almost equal to NetFlow RTT
- Delay contribution of segment from R4 to reverse proxy negligible

Evaluation and results

Evaluation 3: NetFlow-RTT vs. NetFlow-delay

Delay of partial path: R1 - R3

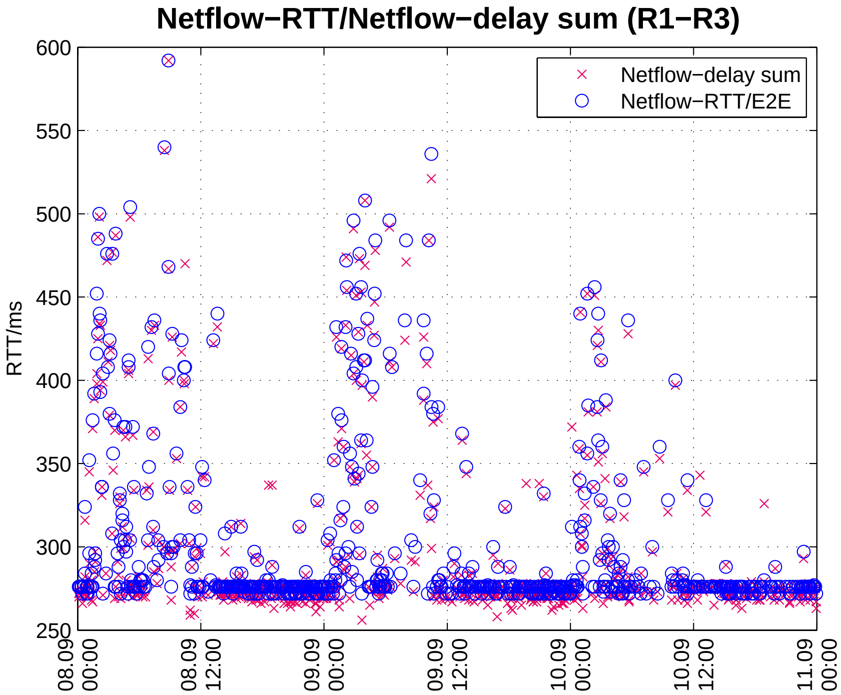

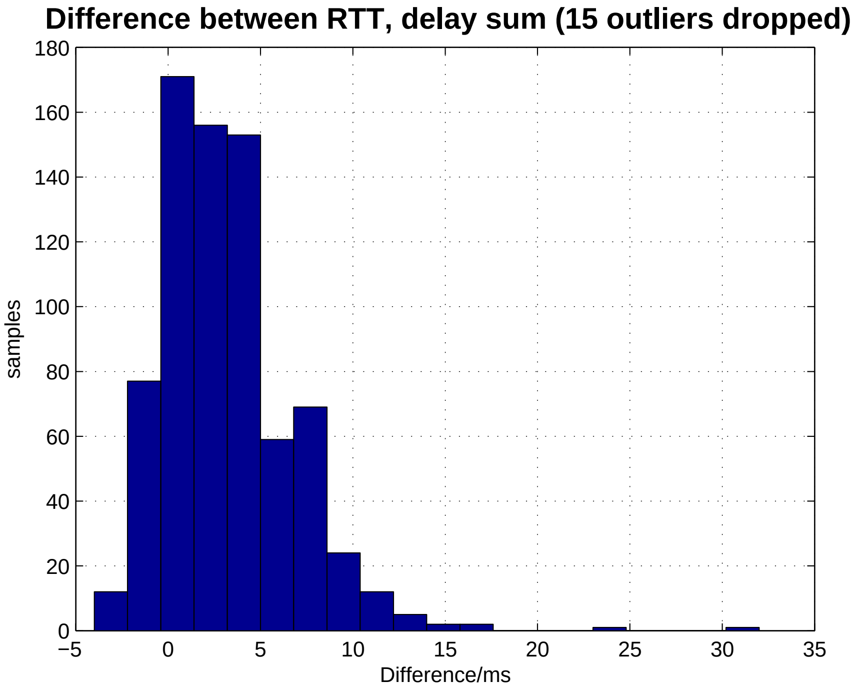

- Delay contribution of segment between R3 and R4 measurable

Conclusion and Outlook

Conclusion

- NetFlow as basis for performance metrics: "QoS Monitoring" and server response
- Comparison to RTT of IP SLA data: differences, but trends are the same
- Comparison of NetFlow-delay and NetFlow-RTT
  delay contribution of network segments measurable

Outlook

- Comparison on a large scale
- Improved algorithms to deal with missing or inaccurate records
- Compensation of record loss by combining information from several routers
- Take knowledge about record loss into account (IPFIX Reliability Statistics?)

- Is active per-packet measurement required at all?