Transparent Transmission Segmentation for Software-Defined Networks

NetSoft 2017

Andreas Schmidt, Thorsten Herfet
Telecommunications Lab
Saarland Informatics Campus - Saarbrücken

July 04, 2017

SAARLAND UNIVERSITY
SAARBRÜCKEN GRADUATE SCHOOL OF COMPUTER SCIENCE

Motivation

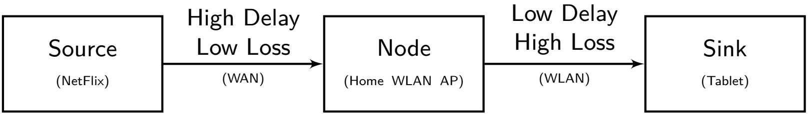

Source
(NetFlix)
High Delay
Low Loss
(WAN)
Node
(Home WLAN AP)
Low Delay
High Loss
(WLAN)
Sink
(Tablet)

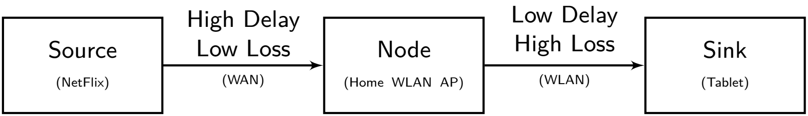

The End-To-End (E2E) Principle [Saltzer1984]

- "The **end-to-end argument** suggests that functions placed at low levels of a system may be **redundant** or of **little value** when compared with the **cost** of providing them at that low level."
- "A **great deal of information** about system implementation is needed to make this **choice [of end-to-end or local implementation] intelligently**."

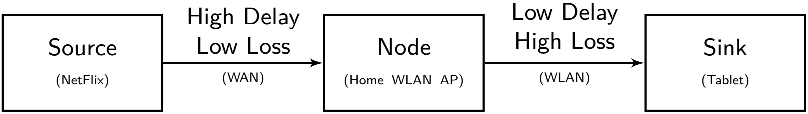

The End-To-End (E2E) Principle [Saltzer1984]

- "The **end-to-end argument** suggests that functions placed at low levels of a system may be **redundant** or of **little value** when compared with the **cost** of providing them at that low level."
- "A **great deal of information** about system implementation is needed to make this **choice [of end-to-end or local implementation] intelligently**."

Operating **transport protocols** (e.g. TCP) purely E2E is **not ideal**.

Transparent Transmission Segmentation (TTS)

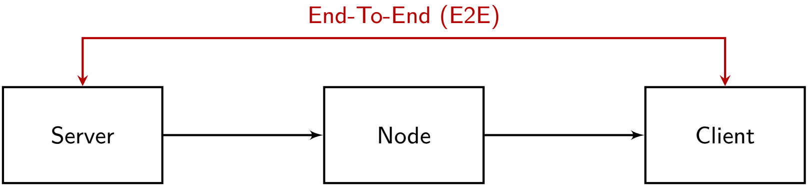

Transparent Transmission Segmentation (TTS)

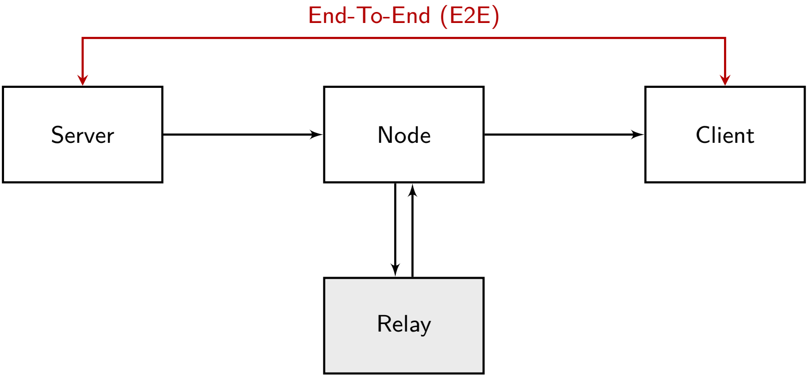

Transparent Transmission Segmentation (TTS)

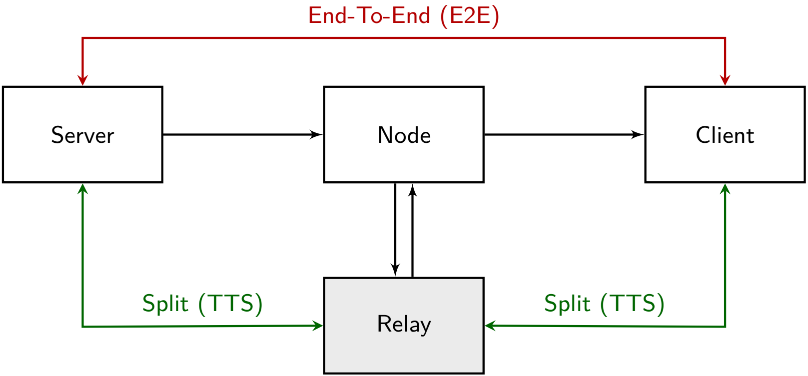

TTS: Requirements and Domains

Requirements

- **Transparency**: The application should only notice a performance improvement.
- **Ease of Deployment**: Use minimal number of components and hardware.
- **Evolutionary Approach**: Leave protocols (e.g. TCP/IP) as is.

Requirements

- **Transparency**: The application should only notice a performance improvement.
- **Ease of Deployment**: Use minimal number of components and hardware.
- **Evolutionary Approach**: Leave protocols (e.g. TCP/IP) as is.

Segmentation Domains

- **Latency**: Propagation, processing and queueing delays.
- **Contention**: Many flows contend, even though they might only share one link.
- **Buffers**: Sizes, fill states and protocol layers.
- **Capacities**: Bottleneck throughput, utilizations.

Error Control

- **Retransmit locally**, reducing how long the resend takes (TCP).
- Set **retransmit timer** to lower values (TCP).
- Use an **ideal forward-error code** (real-time protocols).

SAARLAND UNIVERSITY
SAARBRÜCKEN GRADUATE SCHOOL OF COMPUTER SCIENCE

Error Control

- **Retransmit locally**, reducing how long the resend takes (TCP).
- Set **retransmit timer** to lower values (TCP).
- Use an **ideal forward-error code** (real-time protocols).

Congestion Control

- “Negotiate” sending rate for a **single conflicting link** (TCP).
- **Adapt faster** for RTT-dependent CCAs (TCP Tahoe, Reno).

TTS: Affected Network Functions

Error Control

- **Retransmit locally**, reducing how long the resend takes (TCP).
- Set **retransmit timer** to lower values (TCP).
- Use an **ideal forward-error code** (real-time protocols).

Congestion Control

- "Negotiate" sending rate for a **single conflicting link** (TCP).
- **Adapt faster** for RTT-dependent CCAs (TCP Tahoe, Reno).

Flow Control

- Adapt to **local BDP** with reduced RTT and local bandwidth.
- Achieve **higher utilization** using intermediate buffers.

A Critique of "End-To-End Arguments" [Moors2002]

- Criticizes [Saltzer1984] pointing out **new perspectives**.
- "The decision to implement reliable transfer in the transport layer is not justified on the basis of end-to-end arguments, but rather on the basis of **trust**."

A Critique of "End-To-End Arguments" [Moors2002]

- Criticizes [Saltzer1984] pointing out **new perspectives**.
- "The decision to implement reliable transfer in the transport layer is not justified on the basis of end-to-end arguments, but rather on the basis of **trust**."

Performance-Enhancing Proxies [RFC3135]

- Introduces the approach of **terminating transmissions**.
- Describes **transparency** considerations (user, app, transport, network).

A Critique of "End-To-End Arguments" [Moors2002]

- Criticizes [Saltzer1984] pointing out **new perspectives**.
- "The decision to implement reliable transfer in the transport layer is not justified on the basis of end-to-end arguments, but rather on the basis of **trust**."

Performance-Enhancing Proxies [RFC3135]

- Introduces the approach of **terminating transmissions**.
- Describes **transparency** considerations (user, app, transport, network).

TCP

- Split-TCP providing a similar service but **no transparency**.
- Certain congestion control algorithms are **RTT-independent** (e.g. CUBIC).
- Recent advances: **congestion-based** congestion-control (BBR [Google2016]).

TTS in Softwarized Networks

SAARLAND UNIVERSITY
SAARBRÜCKEN GRADUATE SCHOOL OF COMPUTER SCIENCE

Operating System: **Linux**

- Ubuntu 16.04.
- General purpose solution.

Switching: **OpenvSwitch**

- OpenFlow 1.3 support.

Virtualization: **Docker**

- Lightweight containers.
- Quick startup.
- Low footprint.

Off-the-shelf and **open source**.

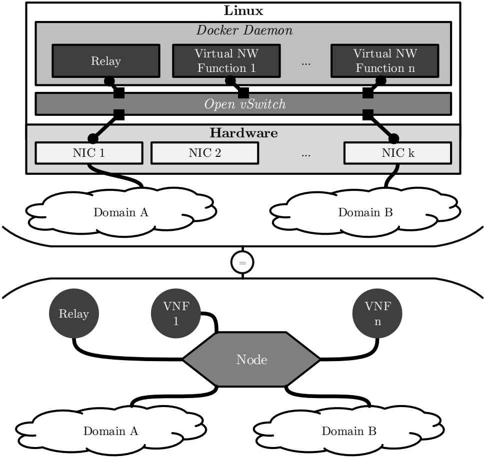

Relay Implementation (NFV)

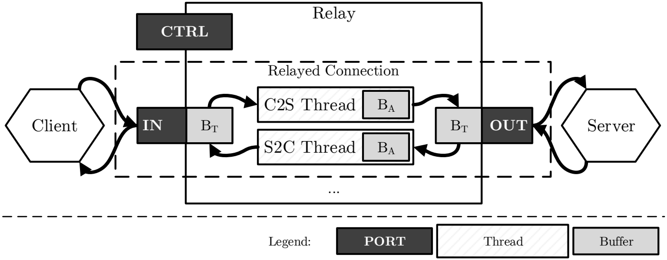

Relay Implementation (NFV)

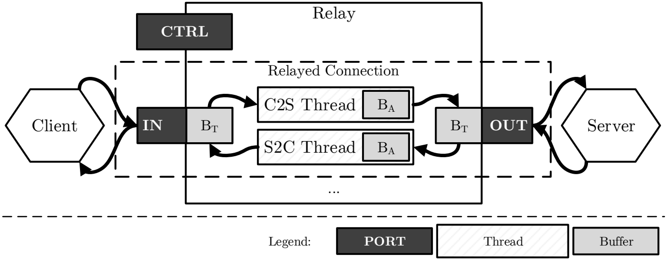

Details

- Pure software written in plain C (glibc, pthreads).
- Configuration includes...
  - addresses (IP, ports) and
  - buffer sizes for app-layer and TCP send/recv buffer.

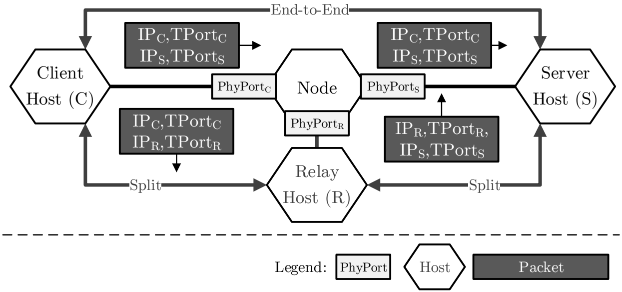

Packet Headers

- Packets on the path from C to S look as normal.
- Packets between N and R look as if R is communicating with both hosts directly.

Relaying Process (SDN)

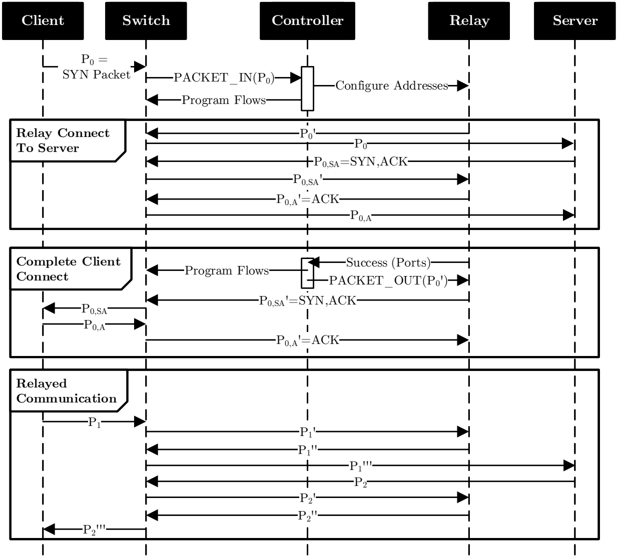

Evaluation

Setup and Methodology

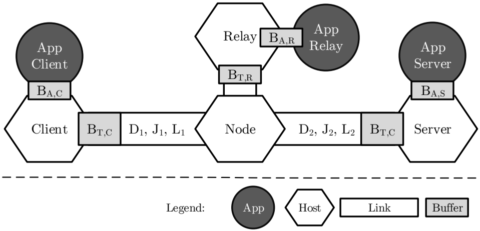

Test System

- Hardware: 8 Cores, 8GB RAM.
- Mininet + netem (network and link simulation).
- TCP Cubic (congestion control algorithm).

Compare E2E (C-S) with TTS (C-R-S).
Metric: Stream a large file using TCP, measure time.

TTS with Lossy Last-Mile

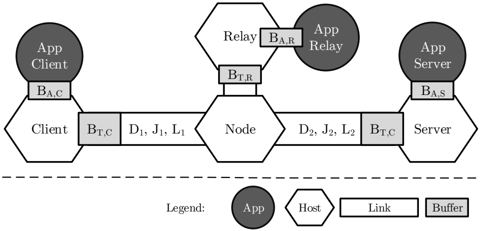

Parameters:

$D_1 = J_1 = D_2 = J_2 = 1ms,\ R = 100MBps,\ B_{A,Relay} = 3000Bytes$

Results (200 trials):

$L_1[\%]$	$L_2[\%]$	μ_{E2E}	μ_{TTS}	σ_{E2E}	σ_{TTS}	A_{12}
10^{-6}	1	6.534	5.923	0.782	0.073	0.884
10^{-6}	10^{-2}	5.864	5.672	0.261	0.071	0.755
10^{-6}	10^{-6}	5.924	5.722	0.295	0.078	0.748

TTS with Lossy Last-Mile

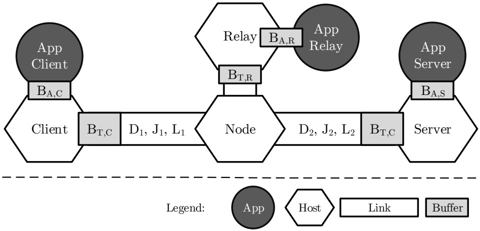

Parameters:

$D_1 = J_1 = D_2 = J_2 = 1ms,\ R = 100MBps,\ B_{A,Relay} = 3000Bytes$

Results (200 trials):

$L_1[\%]$	$L_2[\%]$	μ_{E2E}	μ_{TTS}	σ_{E2E}	σ_{TTS}	A_{12}
10^{-6}	1	6.534	5.923	0.782	0.073	0.884
10^{-6}	10^{-2}	5.864	5.672	0.261	0.071	0.755
10^{-6}	10^{-6}	5.924	5.722	0.295	0.078	0.748

Reduced mean and standard deviation. Higher loss = higher gain through TTS.

TTS with High Jitter

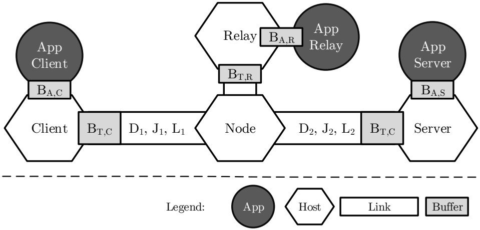

Parameters:

$L_1 = L_2 = 10^{-6}\%,\ D_1 = D_2 = 100ms,\ R = 100MBps,\ B_{A,Relay} = 3000Bytes$

Results (200 trials):

J_1	J_2	μ_{E2E}	μ_{TTS}	σ_{E2E}	σ_{TTS}
1.0	1.0	11.697	6.685	0.010	0.116
10.0	10.0	11.746	7.174	0.039	0.139
100.0	100.0	35.134	8.085	4.652	1.148

TTS with High Jitter

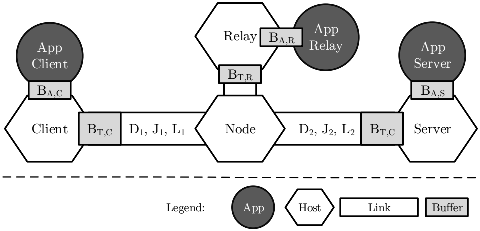

Parameters:

$L_1 = L_2 = 10^{-6}\%,\ D_1 = D_2 = 100ms,\ R = 100MBps,\ B_{A,Relay} = 3000Bytes$

Results (200 trials):

J_1	J_2	μ_{E2E}	μ_{TTS}	σ_{E2E}	σ_{TTS}
1.0	1.0	11.697	6.685	0.010	0.116
10.0	10.0	11.746	7.174	0.039	0.139
100.0	100.0	35.134	8.085	4.652	1.148

Mean and standard deviation are smaller (and grow slower) for TTS than for E2E.

Practical

- Apply concepts to **other protocols** than TCP (RTP, PRRT).
- **Comparative study** of our approach to **Xen or in-router implementations**.
- Automate and streamline the **relay deployment and operation** process.

Theoretical

- Find a heuristic for chosing appropriate **location and number of segmentation points**, given link parameters and application constraints.
- Formalize the **gains** of TTS for TCP with respect to the different network functions.

SAARLAND UNIVERSITY
SAARBRÜCKEN GRADUATE SCHOOL OF COMPUTER SCIENCE

Transparent Transmission Segmentation for Software-Defined Networks

- The decisions based on the **E2E principle** have to be **reconsidered** with SDN.
- Segmenting **connections** can provide significant performance enhancements.
- SDN/NFV can be used for TTS with little deployment overhead and effort.
- Segmentation can be done **without**...
  - **... changing the router**
  - **... changing the protocol**
  - **... special hardware**
  - **... machine-level virtualization**
- More details can be found at `https://www.on.uni-saarland.de`.

Conclusion

Transparent Transmission Segmentation for Software-Defined Networks

- The decisions based on the **E2E principle** have to be **reconsidered** with SDN.
- Segmenting **connections** can provide significant performance enhancements.
- SDN/NFV can be used for TTS with little deployment overhead and effort.
- Segmentation can be done **without**...
  - **... changing the router**
  - **... changing the protocol**
  - **... special hardware**
  - **... machine-level virtualization**
- More details can be found at `https://www.on.uni-saarland.de`.

Thank you for your attention. Questions?

References

[Google2016] N. Cardwell, Y. Cheng, C. S. Gunn, S.H. Yeganeh and V. Jacobson "BBR: Congestion-Based Congestion Control," ACM Queue, vol. 14, no. 5, pp. 50:20-50:53, 2016.

[Mann1947] H. B. Mann and D. R. Whitney, "On a test of whether one of two random variables is stochastically larger than the other," The annals of mathematical statistics, pp. 50–60, 1947.

[Moors2002] T. Moors, "A critical review of 'end-to-end arguments in system design'," in IEEE International Conference on Communications (ICC), vol. 2, pp. 1214–1219, IEEE, 2002.

[RFC3135] J. Border, M. Kojo, J. Griner, G. Montenegro, and Z. Shelby, "RFC 3135 - Performance Enhancing Proxies Intended to Mitigate Link-Related Degradations," tech. rep., 2001.

[Saltzer1984] J. H. Saltzer, D. P. Reed, and D. D. Clark., "End-To-End Arguments in System Design," ACM Transactions on Computer Systems, vol. 2, no. 4, pp. 277–288, 1984.

[Vargha2000] A. Vargha and H. D. Delaney, "A Critique and Improvement of the CL Common Language Effect Size Statistics of McGraw and Wong," Journal of Educational and Behavioral Statistics, vol. 25, no. 2, pp. 101–132, 2000.

Backup

Gaussian

- $D_{total} = \sum_i D_i$ with D_i independent.
- $f(D_{total} = x) = \frac{1}{\sqrt{2\pi\sigma^2}} \cdot e^{-\frac{(x-\mu)^2}{2\sigma^2}}$
- Compare E2E and TTS via μ, σ.

Significance

- For Gaussian, we relied on running enough evaluations.
- For Non-Parametric we used p values from a Mann-Whitney U-test by [Mann1947].

Non-Parametric

- No assumptions about X_1 and X_2.
- A_{12} metric by [Vargha2000]:

$$A_{12} \begin{cases} < 0.5 & : X_1 \text{ is smaller} \\ = 0.5 & : \text{same} \\ > 0.5 & : X_2 \text{ is smaller} \end{cases}$$

- Evaluate N times 1 and M times 2.
- Build cross-product of samples.
- Replace tuple by a 0 (if $1 < 2$), 0.5 (if $1 = 2$) and 1 (if $2 > 1$).
- Average over the replaced values.
- Compare E2E (X_1) and TTS (X_2) via A_{12}.