Endpoint-Transparent Multipath in Software Defined Networks

Master Thesis

Dario Banfi

February 3, 2016

Chair for Network Architectures and Services
Department of Informatics
Technische Universität München

Introduction & Motivation

Goals

Methodology
Scheduling
Reordering
Monitoring

Architecture

Evaluation
MPTCP Comparison
Demo

Contribution

Conclusion

Introduction & Motivation

Multipath Forwarding is a L2/L3 strategy where next-hop packet forwarding to a destination can occur over multiple paths.

Advantages:

- Aggregated Path Throughput
- Load Sharing
- Connection Resiliency

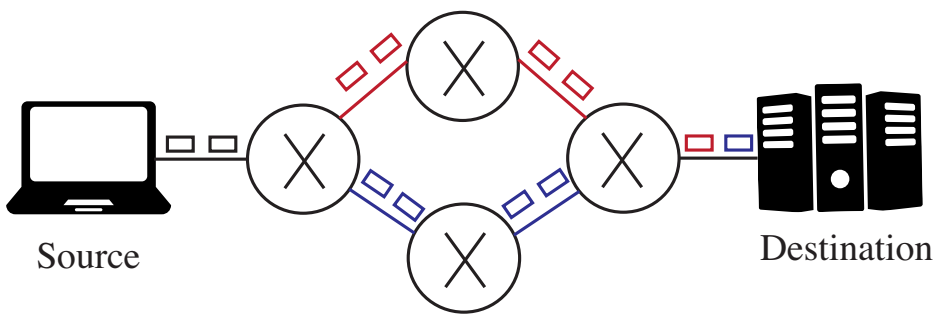

Introduction & Motivation

The multipath approach is used in:

- Transport-Layer Multipath Protocols
  - Multipath TCP (MPTCP) and CMT-SCTP
  - Endpoint-support is necessary
  - Requires multihomed hosts
- In-network Multipath:
  - ECMP, Load-Balancing Techniques
  - Multipath done per-flow
  - No bandwidth aggregation and static configuration

Goals

The goals of the research were the following:

- Use multipath forwarding to provide aggregated path bandwidth
- Endpoint-transparent solution
- Easily deployable into current networks
  → Software Defined Networking
  → Open vSwitch

Methodology

Scheduling

- A Multipath Scheduling Algorithm decides how to forward packets over a set of paths.
- Weighted Round Robin
  - It allocates traffic proportionally to the paths' capacities thanks to weight parameters
  - Sends packet bursts of variable size over the available paths
  - Bursts of packets reduce the probability of reordering
- Implemented in Open vSwitch and configurable by flow rules issued by the controller

```
 match: protocol, source=src, destination=dst
action: weight:10,output:port1 weight:10,output:port2
```

Methodology
Scheduling

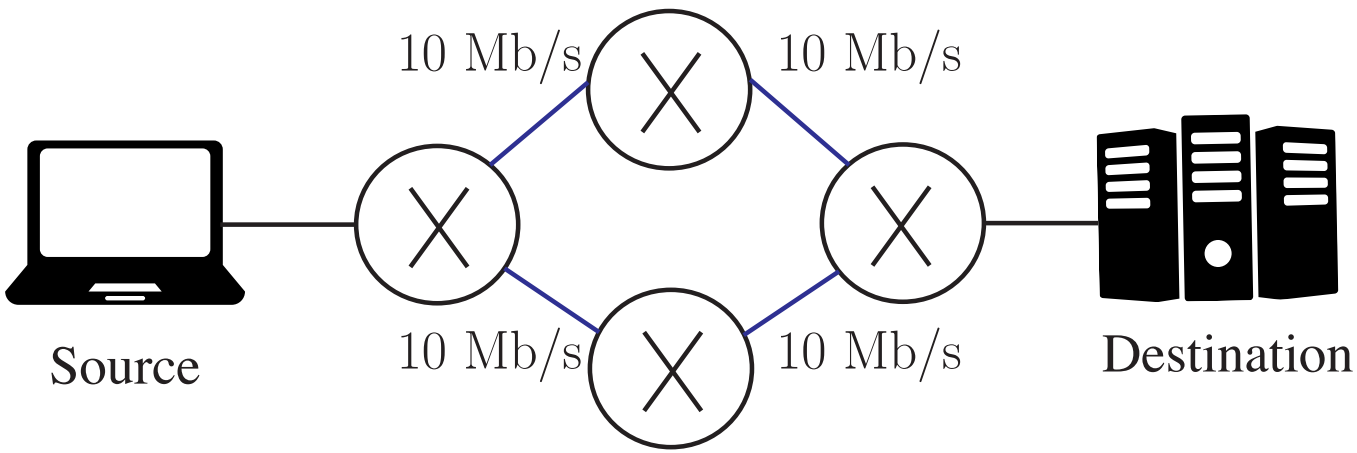

Methodology

Scheduling

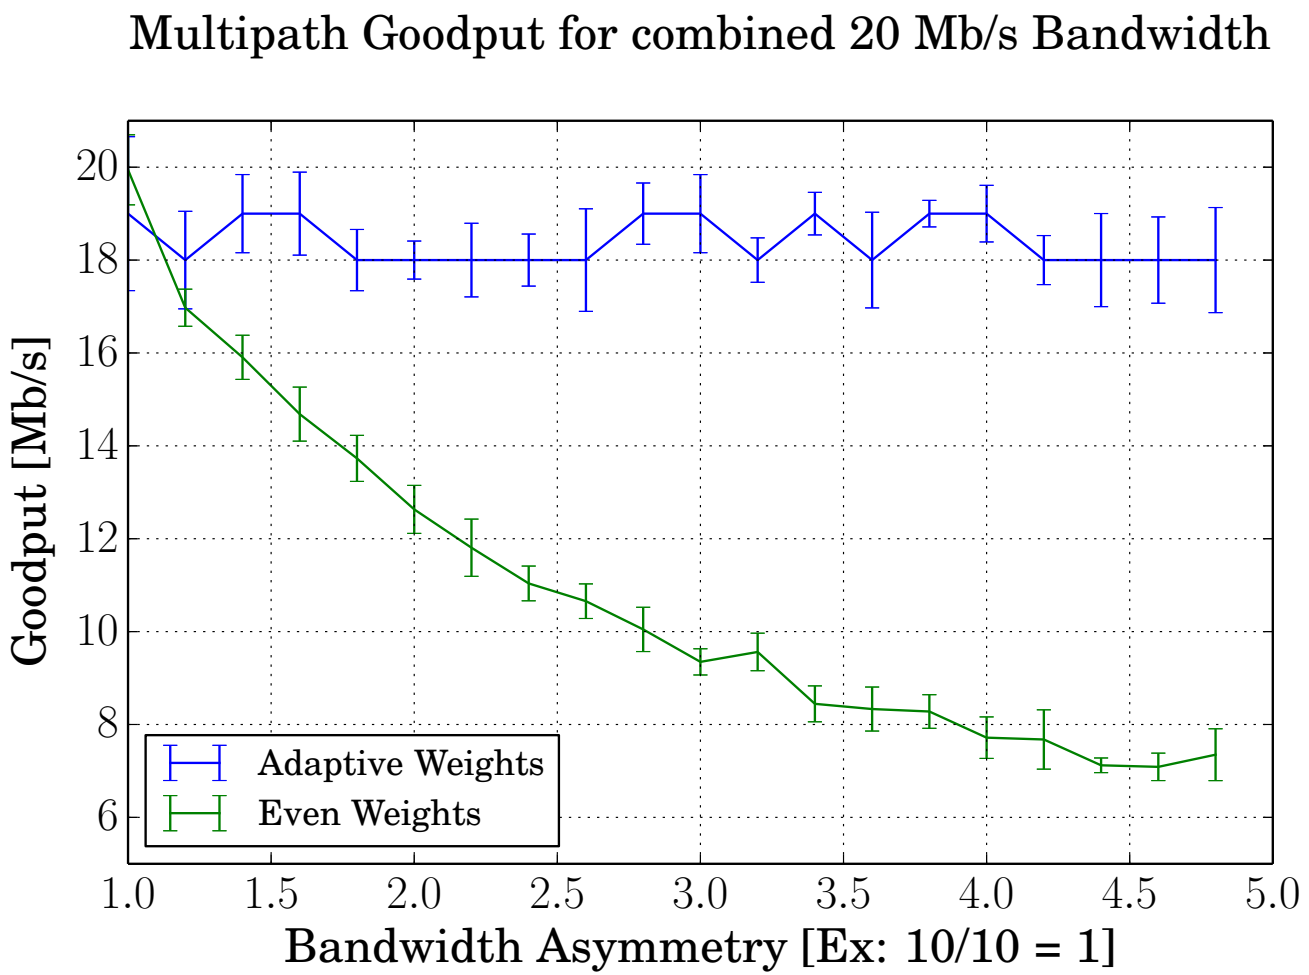

Methodology
Reordering

- Packet reordering causes critical performance degradation, especially on TCP connections
- Solution: Reordering Buffer in the network hardware
  - Reorders packets based on their TCP Sequence Number
  - Reduces unnecessary retransmission and sending rate limitations
  - Best-effort approach
- Implemented in Open vSwitch and configurable by the controller for specific flows

```
match: protocol, source=src, destination=dst
action: reorder,output:port
```

Methodology

Reordering

Multipath performance with two paths of 10 Mb/s

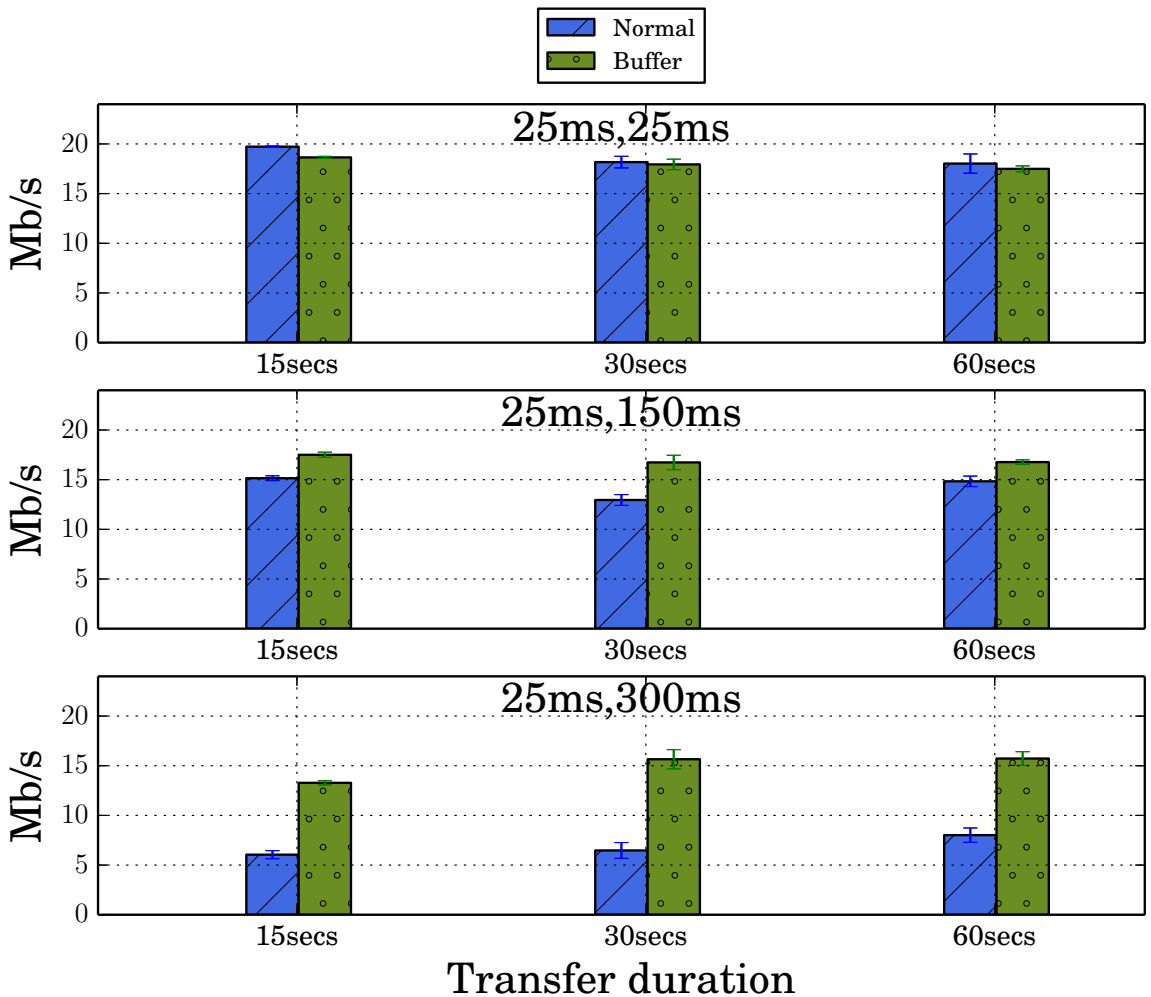

Methodology
Monitoring

- Scheduler's *path weights* and *reordering rules* must be properly configured
- Weights that do not reflect the path capacities will cause worse performance than single-path!
- Necessity of accurate Network Monitoring
  - Performed by the Controller with active and passive Measurements
  - Latency and Bandwidth Measurements between SDN Switches
  - Path characteristics are used to configure multipath forwarding tables

Architecture

- Topology Discovery Component inserts the switches into a graph structure
- Edges of the graph have latency and bandwidth costs
- A Max-Flow algorithm is used to set up forwarding tables
- Adaptive reconfiguration to failures/congestion

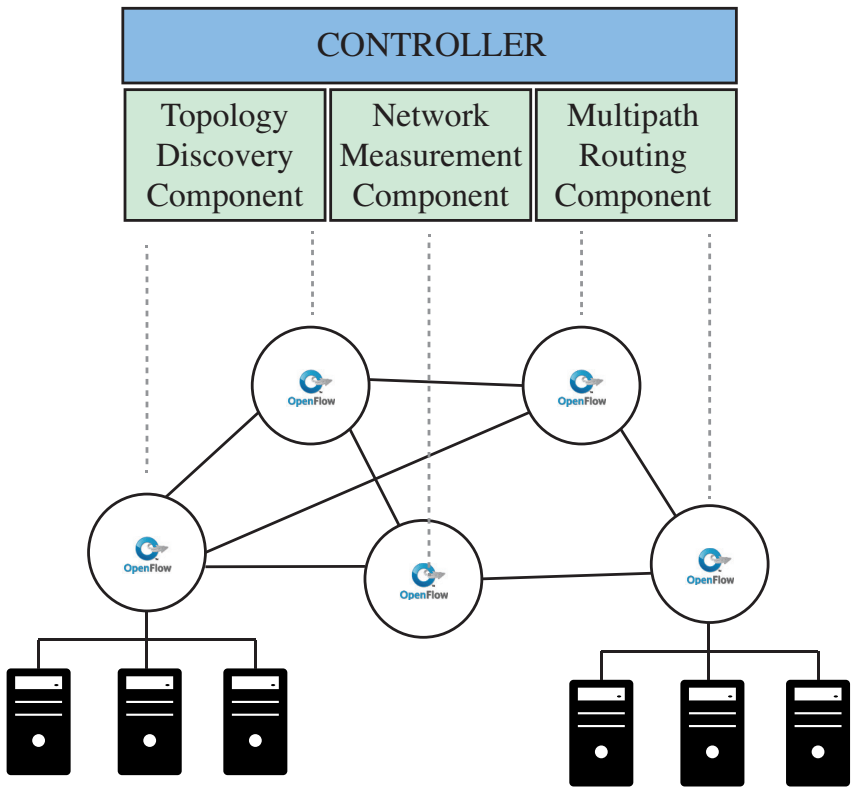

Evaluation

MPTCP Comparison

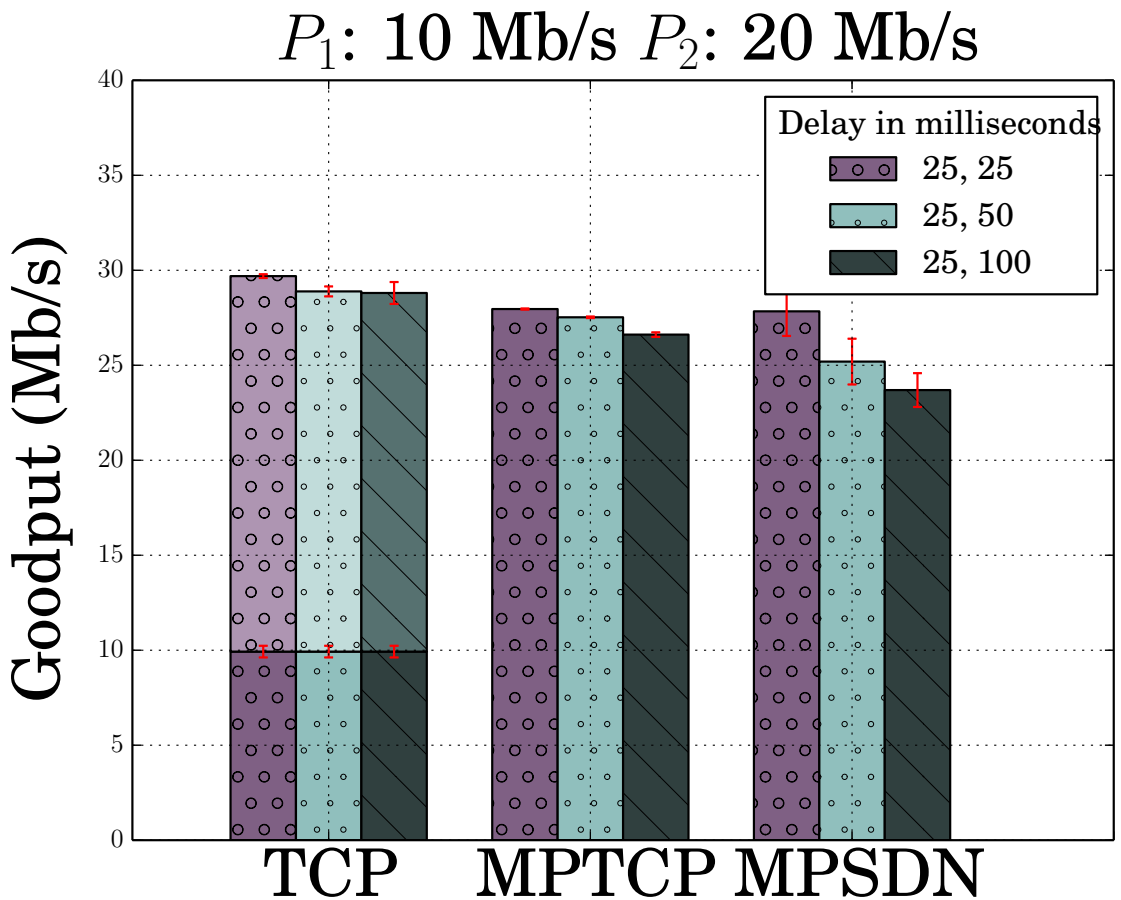

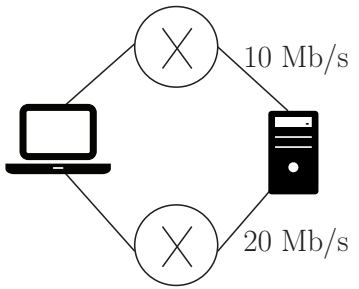

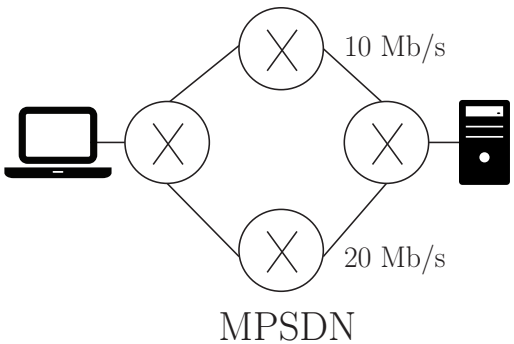

Evaluation

Live Demo (Approximately 120 seconds)

This demo will show:

- Increased Multipath Throughput
- Controller's path setup and dynamic reconfiguration
- Drawbacks of congestion in a sub-path

The steps will be the following:

1. Stream Video Single-Path
2. Controller sets up Multipath Forwarding
3. Congestion in a sub-path
4. Controller readapts forwarding automatically

Contribution

- Analysis on optimal multipath parameters to maximize throughput and reduce out-of-order delivery
- Open vSwitch modification for packet-granularity scheduling
- Open vSwitch modification for TCP Reordering
- SDN Controller performing network monitoring and multipath forwarding rules computation
- Evaluation on Emulated Networks and Large-Scale Multihomed Testbed (**NorNet**)

Conclusion

+ Increased throughput and resource utilization
+ Adaptive reconfiguration
+ L2/L3 Capabilities
+ Endpoint-Transparent

But

– Network monitoring must be accurate to configure correctly multipath parameters
– Reordering Buffer is a performance bottleneck
– Increased complexity

Conclusion

Q & A

Thanks for the attention!

Dario Banfi
dario.banfi@tum.de