Institut für Kommunikations-Technik

TNOVA

Towards a Scalable SDN Virtualization Platform

IFIP/IEEE SDNMO 2014

Zdravko Bozakov, Panagiotis Papadimitriou

Leibniz Universität Hannover, Germany

Leibniz Universität Hannover

Introduction

- Network virtualization in multi-tenant data-centers:
  - ✓ Elastic provisioning
  - ✓ Robustness to failures
  - ✓ Network abstraction
  - ✓ Reduction in OPEX and CAPEX for enterprise networks
  - ✘ Limited control and access on virtualized network devices

SDN Virtualization

- Virtual networks programmable as SDNs (vSDNs)
- Benefits:
  - Tenants:
    - Advanced control and access on virtualized network devices
  - Providers:
    - Less configuration overhead for vSDN operation
    - New cloud service model:
      - SDN as a Service (SDNaaS)

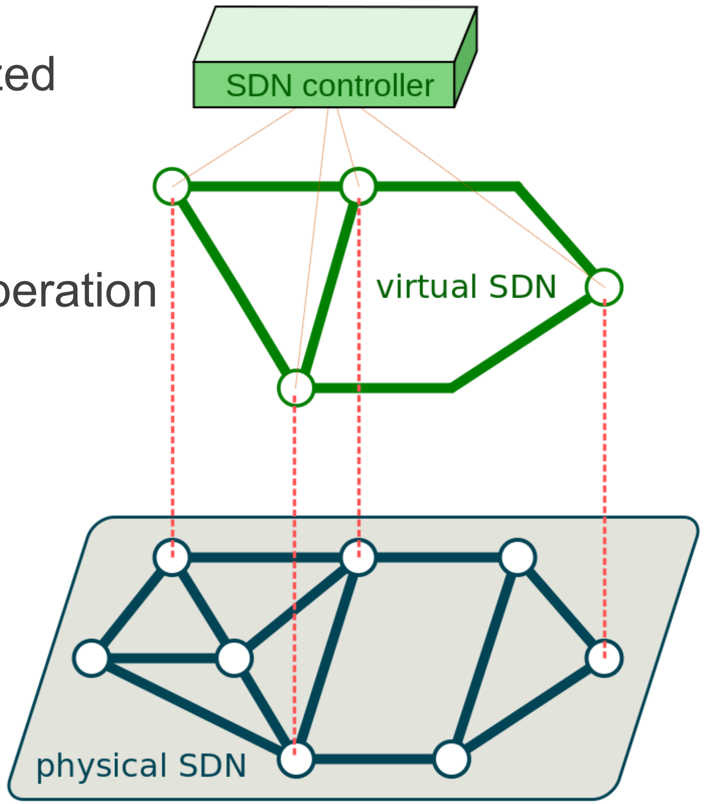

SDN Virtualization Challenges

- Automation of vSDN setup
  - vSDN mapping
  - Transparent allocation of isolated flowspaces
    - Selection of identifiers
  - Generation and installation of flow entries for packet forwarding and encapsulation
    - Binding traffic to logical context using tagging
- Deployment of arbitrary vSDN topologies
  - Mapping multiple virtual switches onto the same switch

Existing Solutions

- FlowVisor
  - ✓ Flow table isolation
  - ✓ Transparent control message filtering and rewriting
  - ✘ Virtual address space allocation not supported
  - ✘ Configuration overhead for SDN slice deployment
- OpenVirtex
  - ✓ Virtual SDN deployment
  - ✓ Transparent control message rewriting
  - ✘ Scalability

Outline

- SDN hypervisor
  - SDN data plane virtualization
  - vSDN setup automation
  - Control message translation
- Implementation
- Conclusions

SDN Hypervisor Design

SDN Hypervisor Overview

- Distributed SDN hypervisor
  - Multiple autonomous controller proxies (CPX)
    - Coordination by a management module (MM)
- Dataplane segmentation
  - Multiple SDN domains
  - Switches within a domain controlled by the same controller proxy

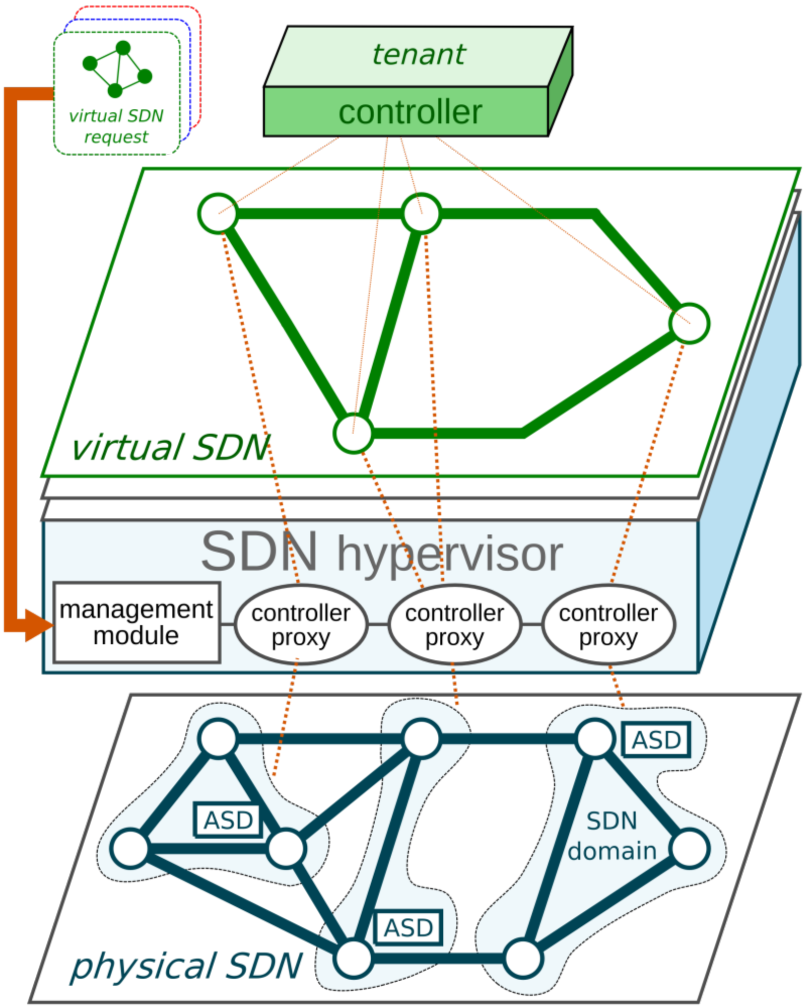

Data Plane Scalability

- Switch flow tables can't hold the flow entries of many tenants
  - Solution: SW datapaths in commodity servers
- Leverage on Internet flow distribution:
  - A small subset of flows carries most of Internet traffic
    - Traffic statistics from an access router at a large European ISP
      - 100 prefixes → 50% of total traffic
      - 1000 prefixes → 80% of total traffic
- Dual-datapath approach:
  - Datapath in OpenFlow switch handling elephant flows
  - Auxiliary SW datapath (ASD) in commodity server handling mice traffic

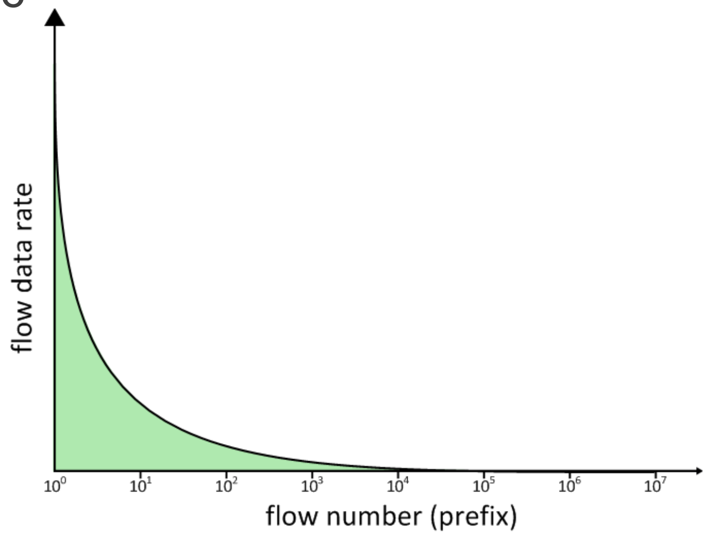

Control Plane Functions

- Management module
  - Topology embedding
  - SDN domain segmentation
  - CPX coordination for network-wide resource management
- Controller proxy
  - Infrastructure flow entry installation
  - Message translation
  - Flow cache management
  - SDN domain optimizations

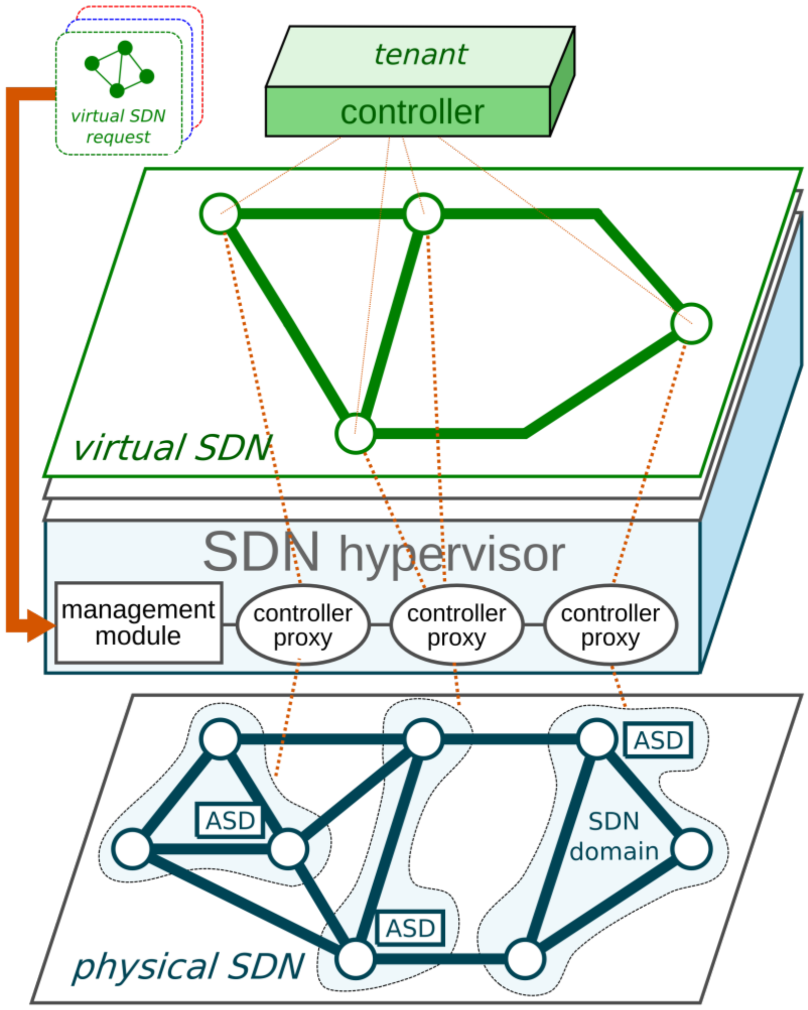

SDN Data Plane Virtualization

- Flow table segmentation

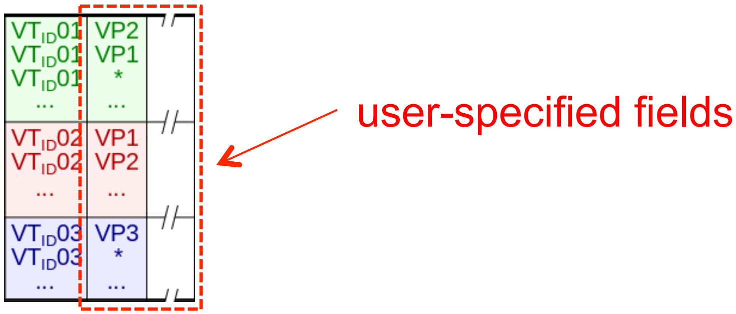

- Link identifier (LID) assignment and resolution
  - Resolver table stored in the switch flow table
    - Simplified LID remapping for vSDN resource migrations

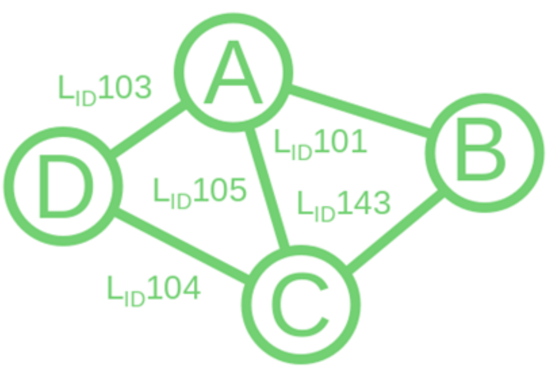

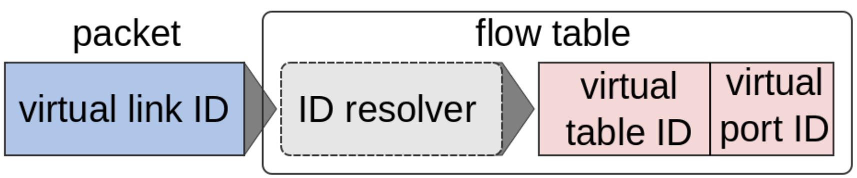

Packet Forwarding

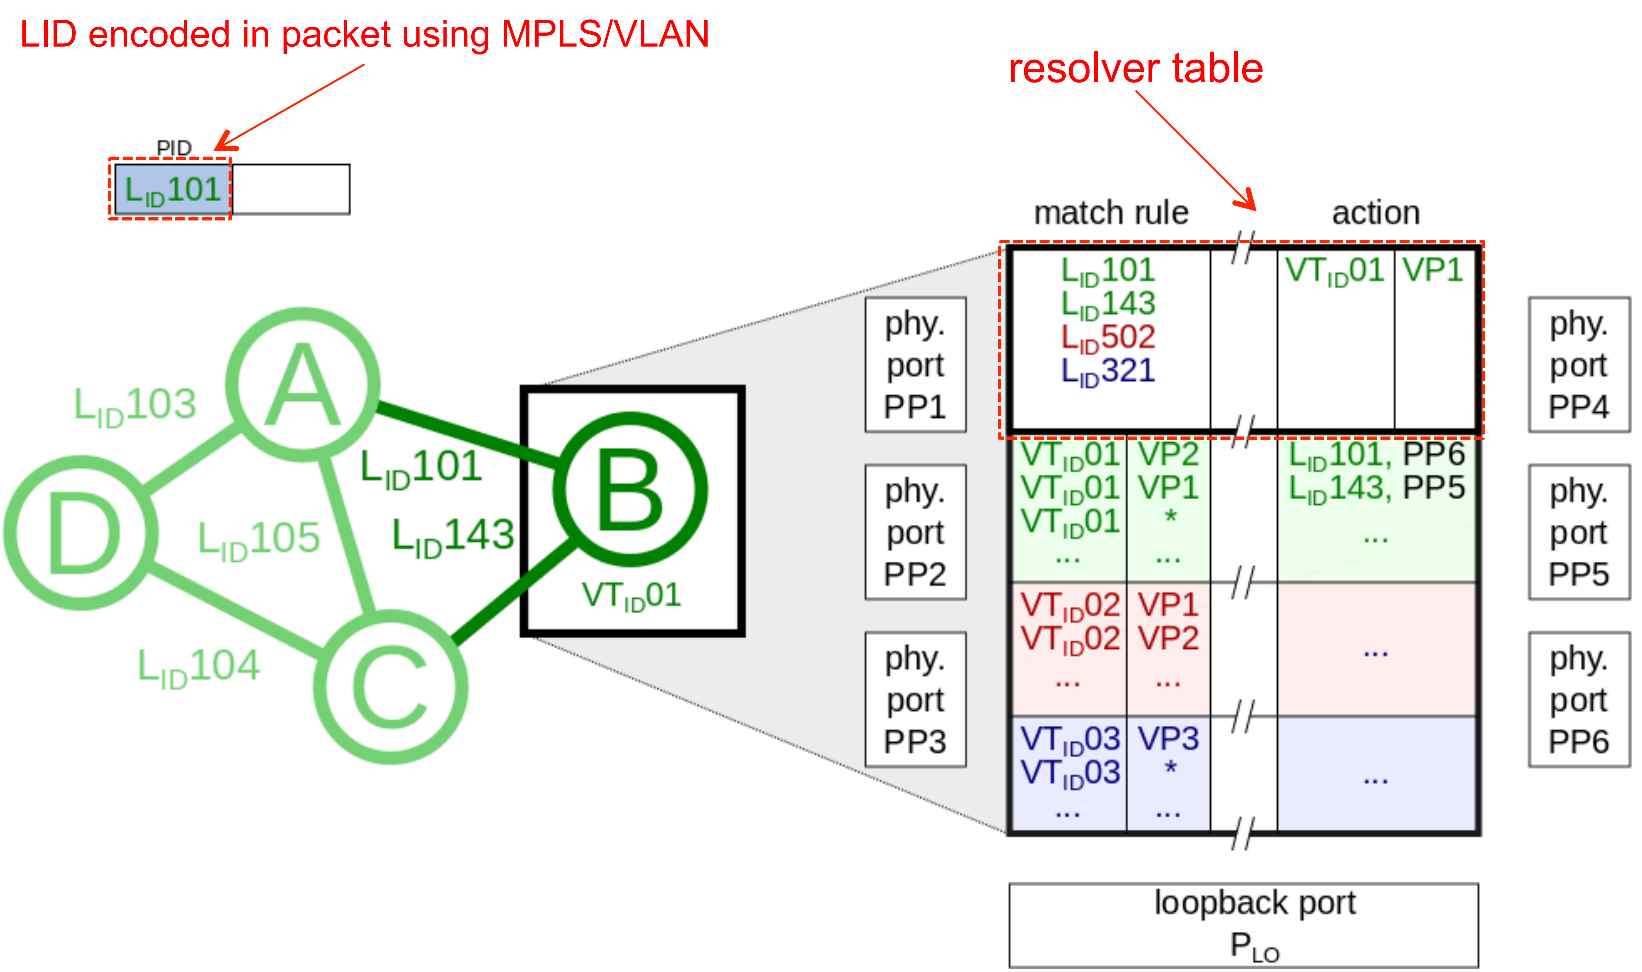

Packet Forwarding

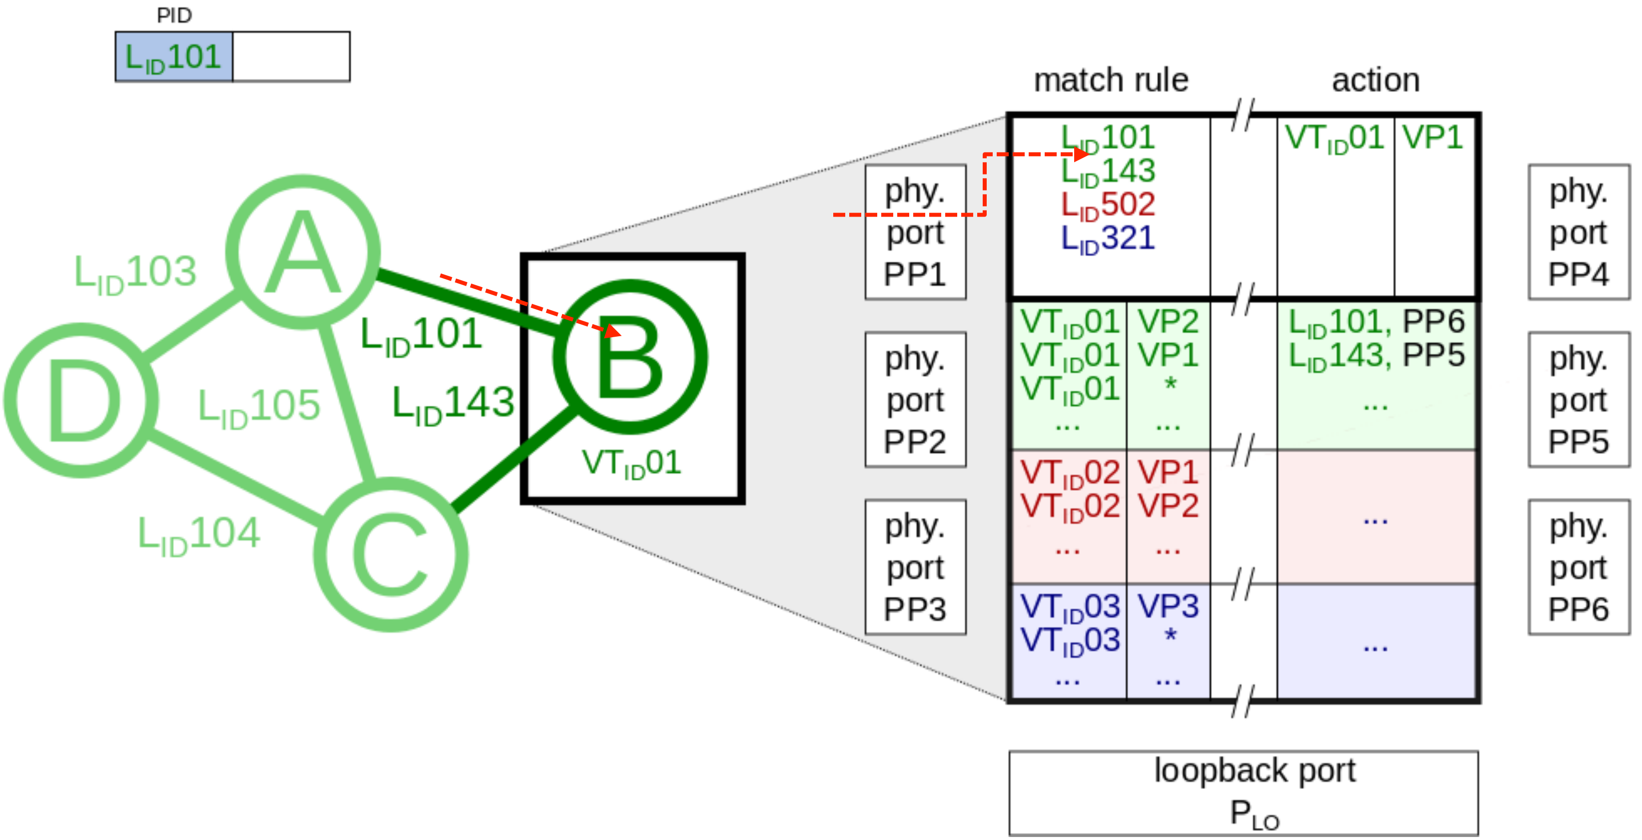

Packet Forwarding

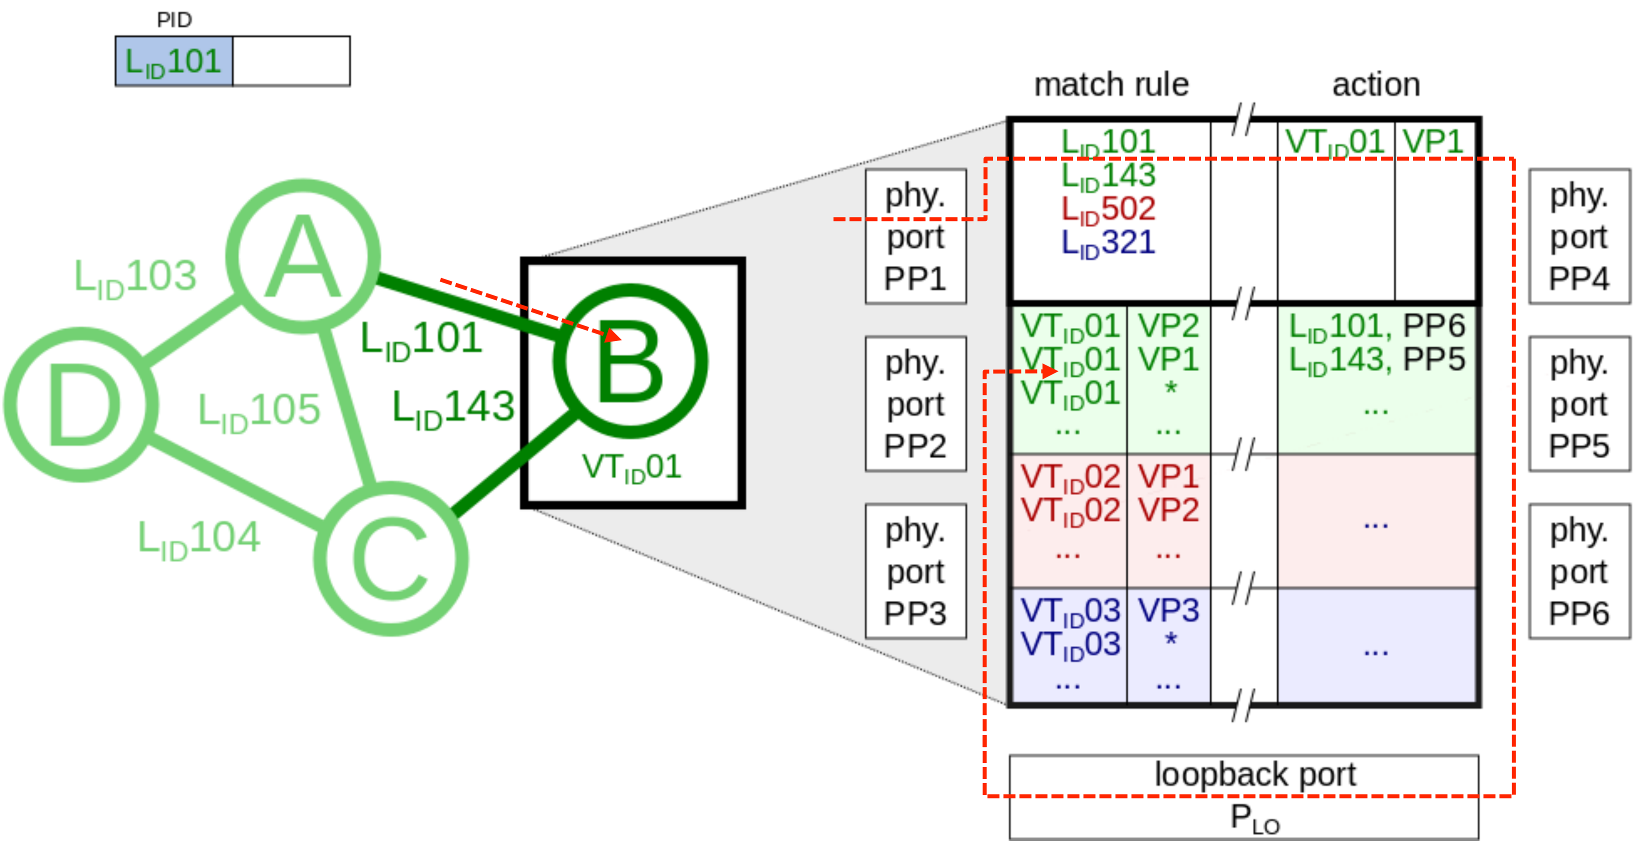

Packet Forwarding

Institut für Kommunikations-Technik

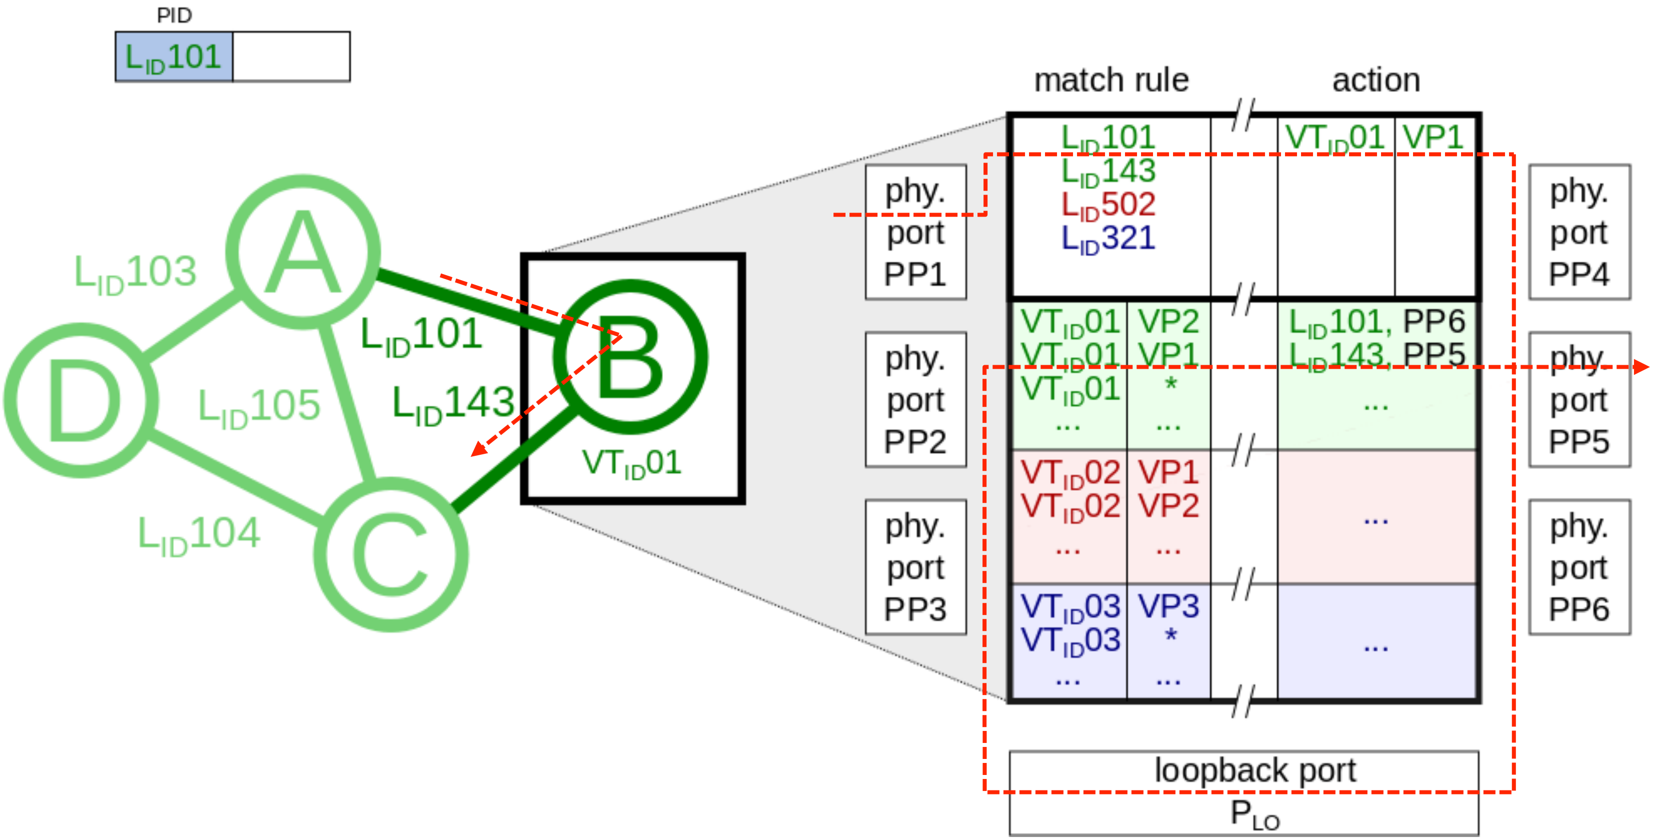

- Easier to deploy with multiple flow tables

Virtual SDN Setup

1. vSDN topology mapping

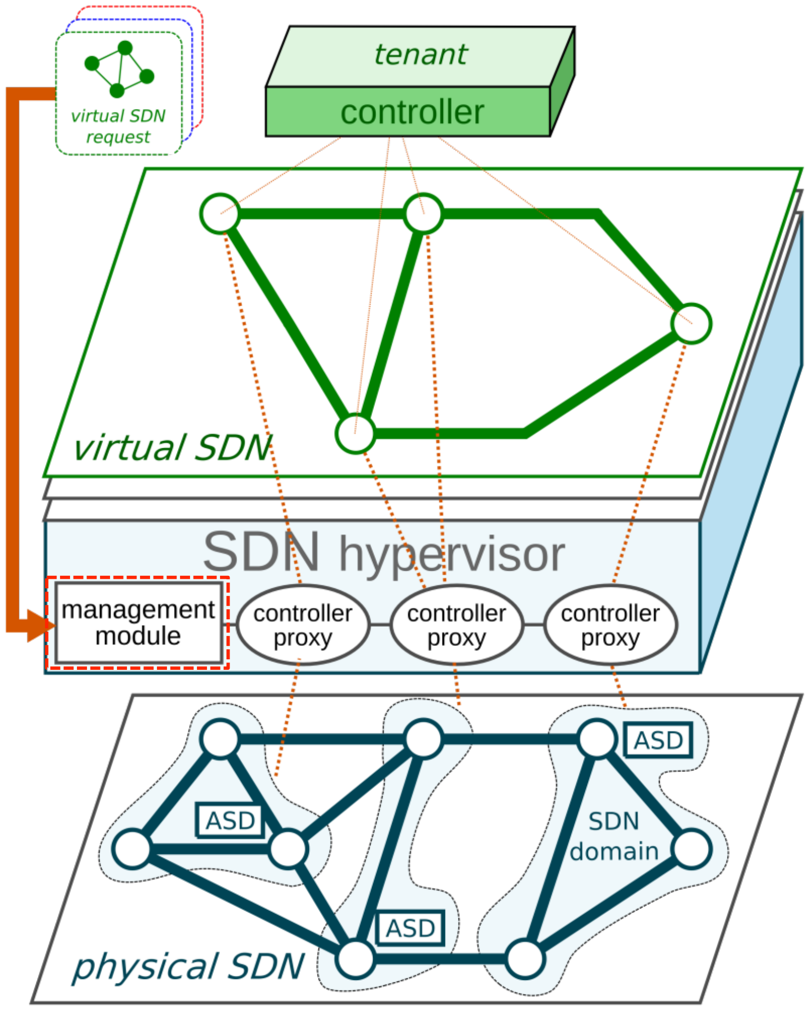

Virtual SDN Setup

1. vSDN topology mapping
2. Selection of identifiers

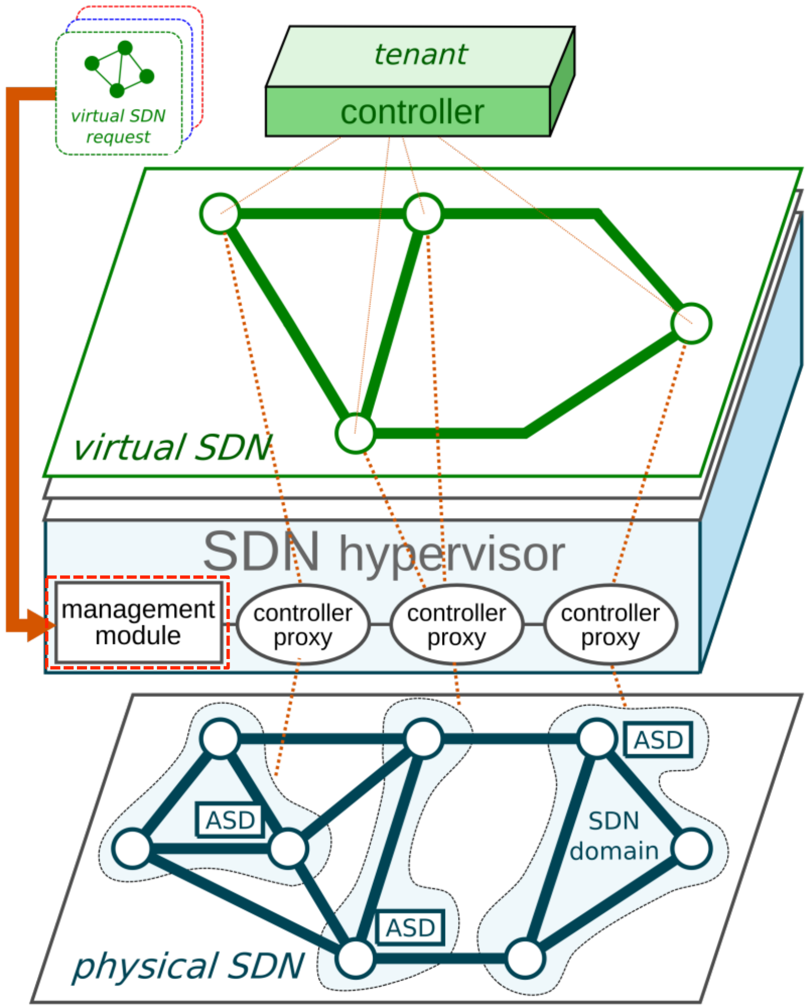

Virtual SDN Setup

Institut für
Kommunikations-
Technik

1. vSDN topology mapping
2. Selection of identifiers
3. Assignment of vSDN resources and identifiers to corresponding CPXs

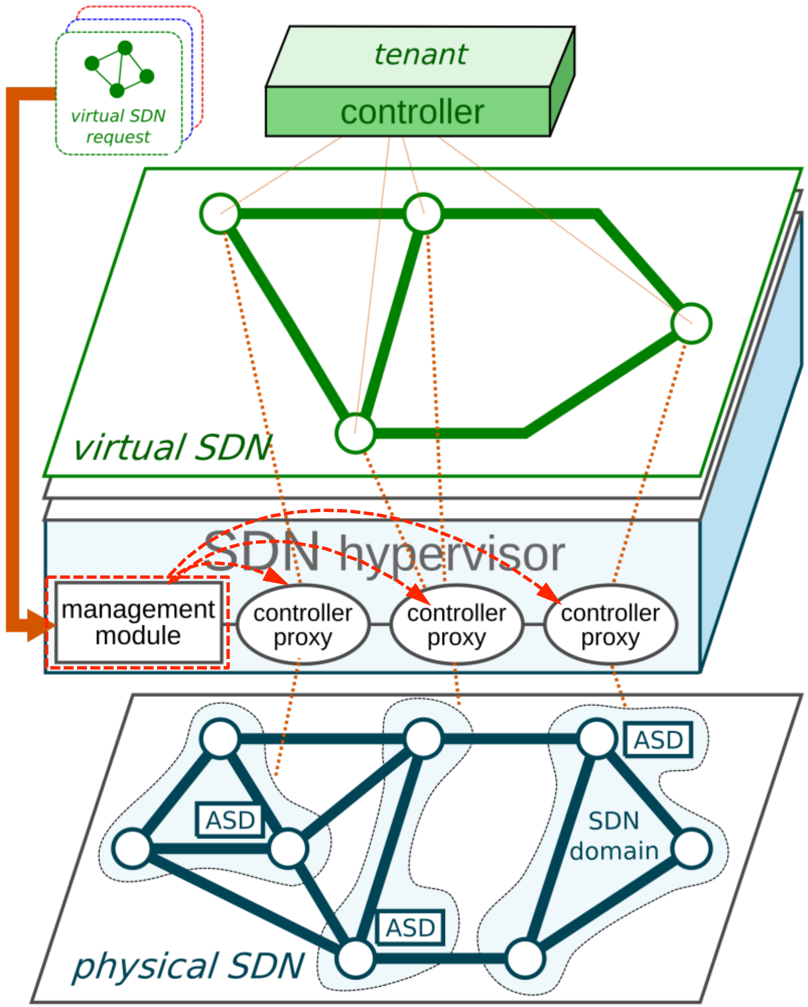

Virtual SDN Setup

1. vSDN topology mapping
2. Selection of identifiers
3. Assignment of vSDN resources and identifiers to corresponding CPXs
4. “Infrastructure” flow entry installation
   - Packet forwarding at intermediate nodes

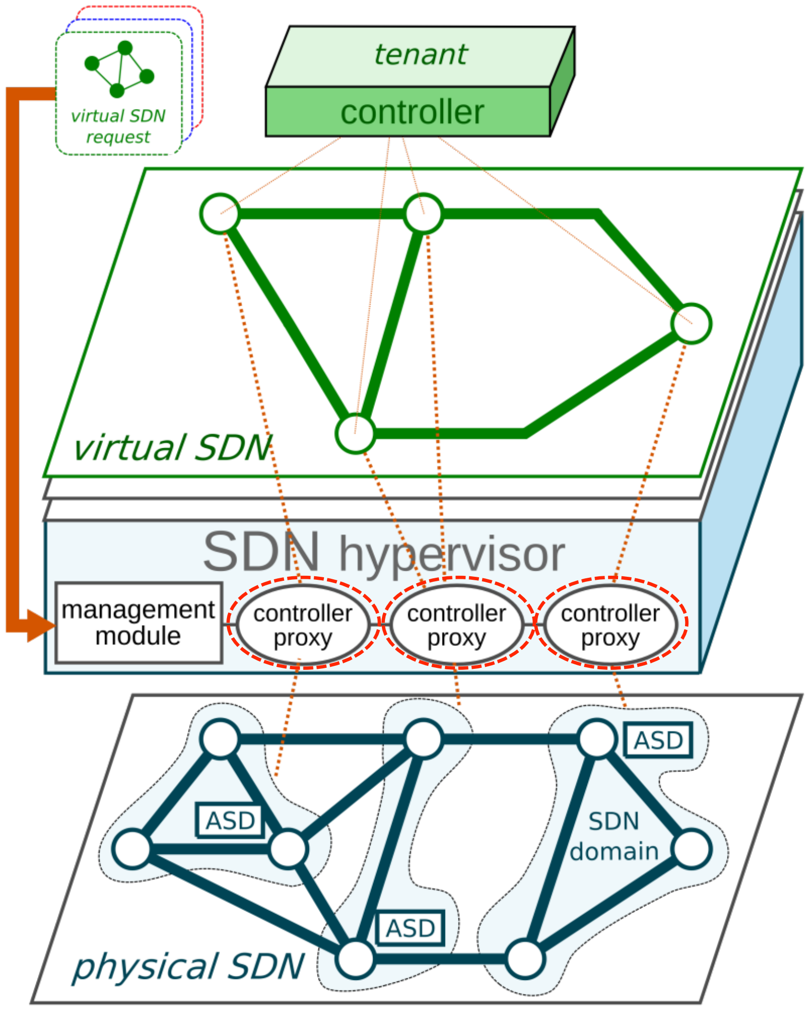

Arbitrary vSDN Topology Deployment

Institut für Kommunikations-Technik

- Mapping multiple virtual switches onto the same physical switch
  - Multiple lookups on a single flow table using a loopback interface

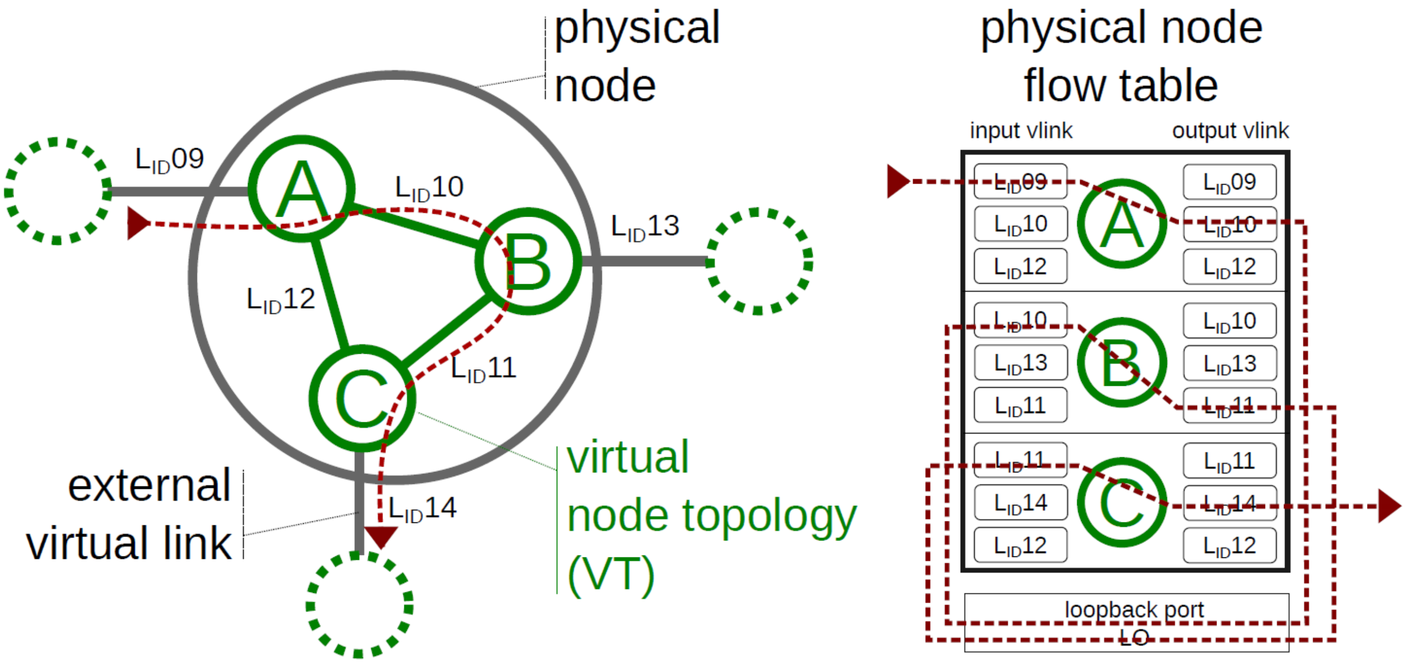

Control Message Translation

- Transparent translation of control messages:
  - Translation of references between logical and physical resource identifiers
  - Policy control to prevent access to unauthorized vSDN resources

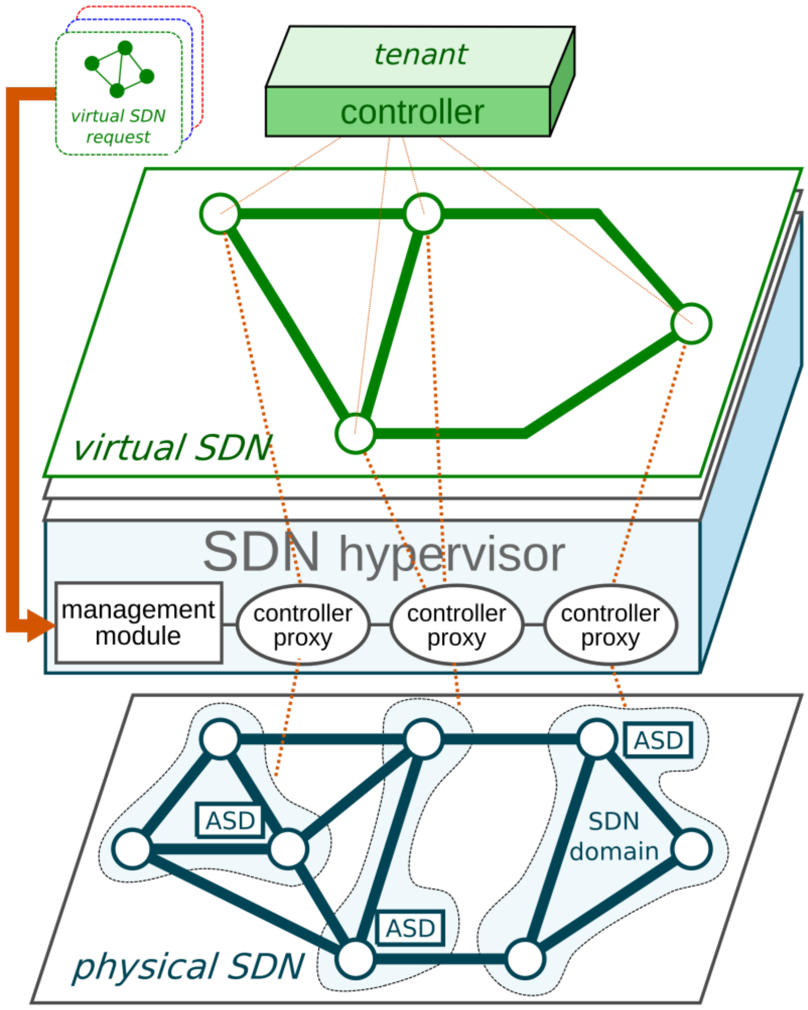

Control Message Translation

- Transparent translation of control messages:
  - Translation of references between logical and physical resource identifiers
  - Policy control to prevent access to unauthorized vSDN resources

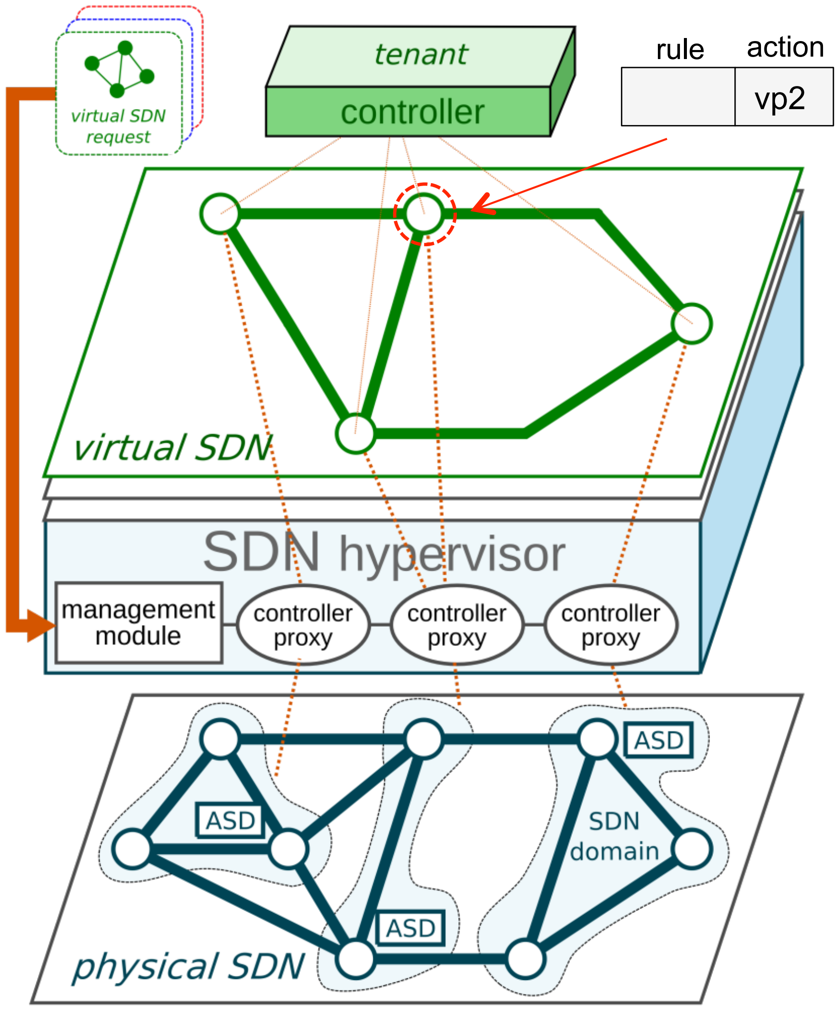

Control Message Translation

- Transparent translation of control messages:
  - Translation of references between logical and physical resource identifiers
  - Policy control to prevent access to unauthorized vSDN resources

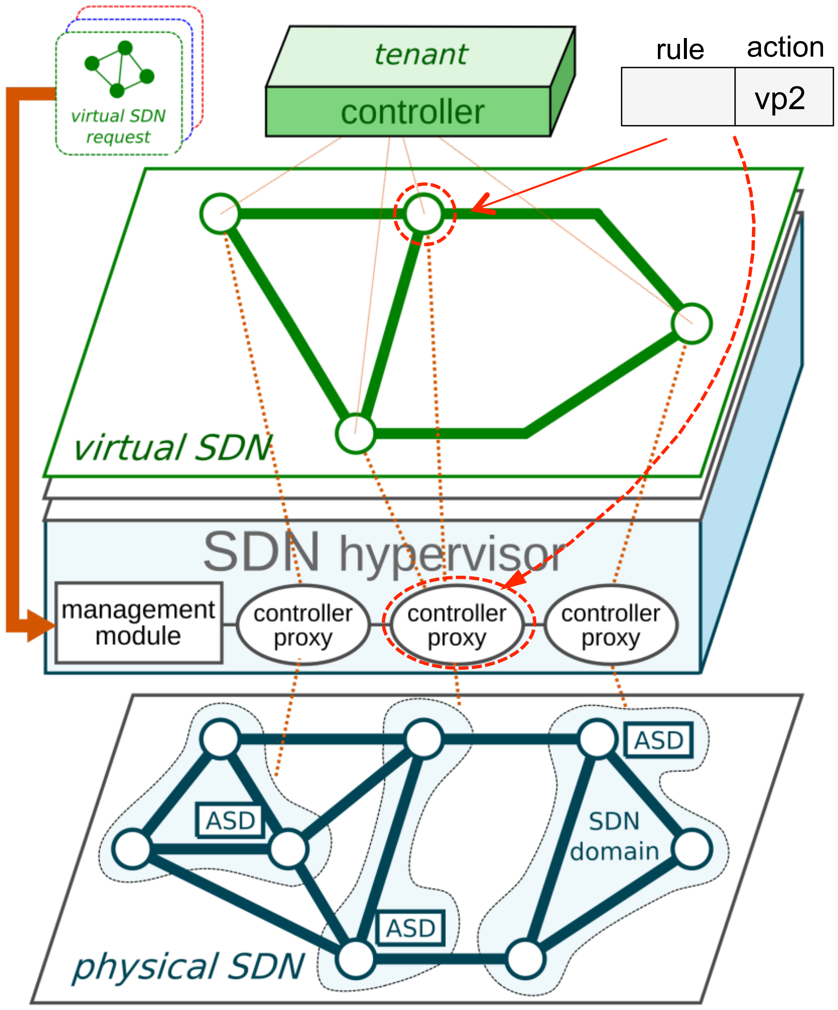

Control Message Translation

- Transparent translation of control messages:
  - Translation of references between logical and physical resource identifiers
  - Policy control to prevent access to unauthorized vSDN resources

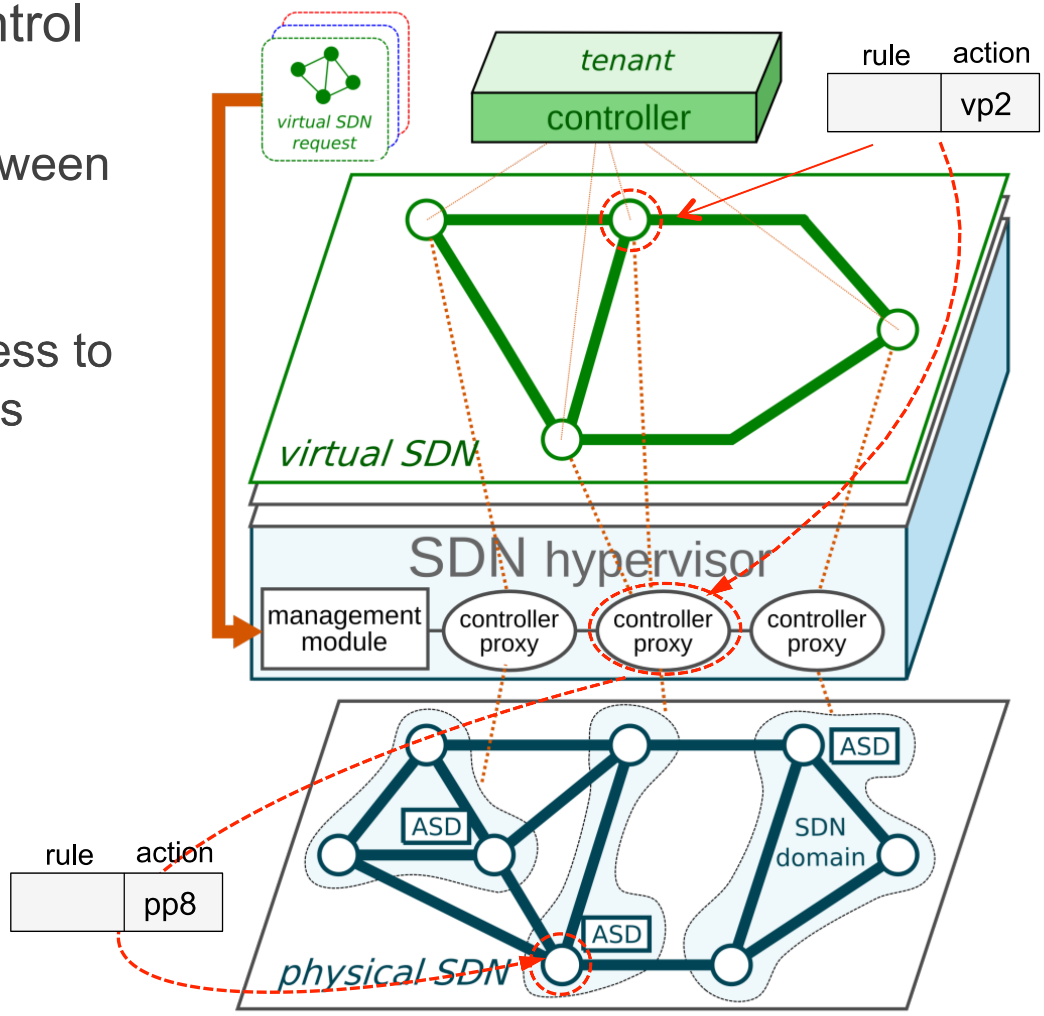

Implementation

SDN Hypervisor Requirements

- OpenFlow:
  - Multiple tags (e.g., VLAN/MPLS) for scalability
  - Arbitrary masking for VLAN and MPLS tags
- Switching hardware:
  - Loopback interfaces or multiple flow tables for mapping multiple virtual nodes onto a single switch
  - Multiple queues per port for bandwidth isolation

Prototype Implementation

- Data Plane:
  - SW datapath:
    - OpenvSwitch (hosted in a server with quad-core Xeon CPUs @2.27GHz)
  - Switch data path:
    - OF Pronto 3290 switch with 48 x 1G ports

- Hypervisor:
  - vSDN embedding
  - Control message translation (FlowVisor)
  - Flow cache management (NOX)

Conclusions

Conclusions and Future Work

- Distributed SDN hypervisor:
  - vSDN embedding
  - vSDN deployment
  - Transparent vSDN operation and configuration

- Future work:
  - Interplay between vSDN embedding and SDN segmentation
  - CPX collaboration for network-wide resource management

Thank you!

Panagiotis Papadimitriou

E-mail: panagiotis.papadimitriou@ikt.uni-hannover.de
WWW: http://www.ikt.uni-hannover.de/