On SDN Research Topics

24/02/2014

Prof. Christian Esteve Rothenberg
FEEC/UNICAMP

chesteve@dca.fee.unicamp.br

Agenda

- SDN in the WAN : Software-Defined IP Routing
- High Performance SDN Stacks
- High-Available SDNs
- SDN & Electronics and Optics

Christian Esteve Rothenberg

- PhD (FEEC/UNICAMP, 2010), MSc (Darmstadt University, 2006), Telecom Eng. (Universidad Politécnica de Madrid, 2004)
  - (2008) Visiting researcher at Ericsson Research Nomadic Lab, participated in EU Publish/Subscribe Internet Routing Paradigm (PSIRP)
- Assistant Professor at FEEC/Unicamp since August 2013
- Research Scientist at CPqD (2010-2013)
- ONF Research Associate since April 2013
- Technical Lead of
  - RouteFlow (Virtual IP Routing Services over SDN)
  - Ofsoftswitch13 (OpenFlow 1.3 controller, softswitch, and testing)
  - Mini-CCNx

Research Interests

- Network Architectures
- Information-Centric Networking
- Data Center Networks / Cloud
- OpenFlow / SDN
- Network Functions Virtualisation
- ...

Citation indices

	All	Since 2009
Citations	665	656
h-index	12	12
i10-index	14	14

Citations to my articles

Select: All, None | Actions | Show: 20 | 1-20 Next >

Title / Author	Cited by	Year
LIPSIN: line speed publish/subscribe inter-networking P Jokela, A Zahemszky, C Esteve Rothenberg, S Arianfar, P Nikander ACM SIGCOMM Computer Communication Review 39 (4), 195-206	289	2009
Theory and practice of bloom filters for distributed systems S Tarkoma, CE Rothenberg, E Lagerspetz Communications Surveys & Tutorials, IEEE 14 (1), 131-155	49	2012
Towards a new generation of information-oriented internetworking architectures C Esteve, F Verdi, M Magalhães First Workshop on Re-Architecting the Internet, Madrid, Spain	41 *	2008
A review of policy-based resource and admission control functions in evolving access and next generation networks CE Rothenberg, A Roos Journal of Network and Systems Management 16 (1), 14-45	37	2008
Revisiting routing control platforms with the eyes and muscles of software-defined networking CE Rothenberg, MR Nascimento, MR Salvador, CNA Corrêa, ... Proceedings of the first workshop on Hot topics in software defined networks ...	33	2012
Virtual routers as a service: the routeflow approach leveraging software-defined networks MR Nascimento, CE Rothenberg, MR Salvador, CNA Corrêa, ... Proceedings of the 6th International Conference on Future Internet ...	32	2011
Forwarding anomalies in Bloom filter-based multicast M Sarela, CE Rothenberg, T Aura, A Zahemszky, P Nikander, J Ott INFOCOM, 2011 Proceedings IEEE, 2399-2407	27	2011
Quagflow: partnering quagga with openflow MR Nascimento, CE Rothenberg, MR Salvador, MF Magalhães ACM SIGCOMM Computer Communication Review 40 (4), 441-442	22	2010
Self-routing denial-of-service resistant capabilities using in-packet Bloom filters		

SDN in the WAN

Research Topics on Software-Defined IP Routing

Motivation

- Combining flexible (open-source) IP routing stacks with high-performance (programmable) hardware
- Augmenting traditional IP control planes with centralized views and flow programmability
- A migration path to SDN, allowing SDN islands talk to legacy networks

- Further reading:

Revisiting Routing Control Platforms with the Eyes and Muscles of Software-Defined Networking

Christian E. Rothenberg, Marcelo R. Nascimento, Marcos R. Salvador
Telecomm. Research and Development Center (CPqD)
Campinas - SP - Brazil
esteve@cpqd.com.br

Carlos N. A. Corrêa, Sidney C. de Lucena
Federal University of the Rio de Janeiro State (UniRio)
Rio de Janeiro - RJ - Brazil
carlos.correa@uniriotec.br

Robert Raszuk
NTT MCL
San Mateo - California - USA
rr@nttmcl.com

Software Defined IP Routing

High cost
Specialized config.
Closed source
Slow innovation

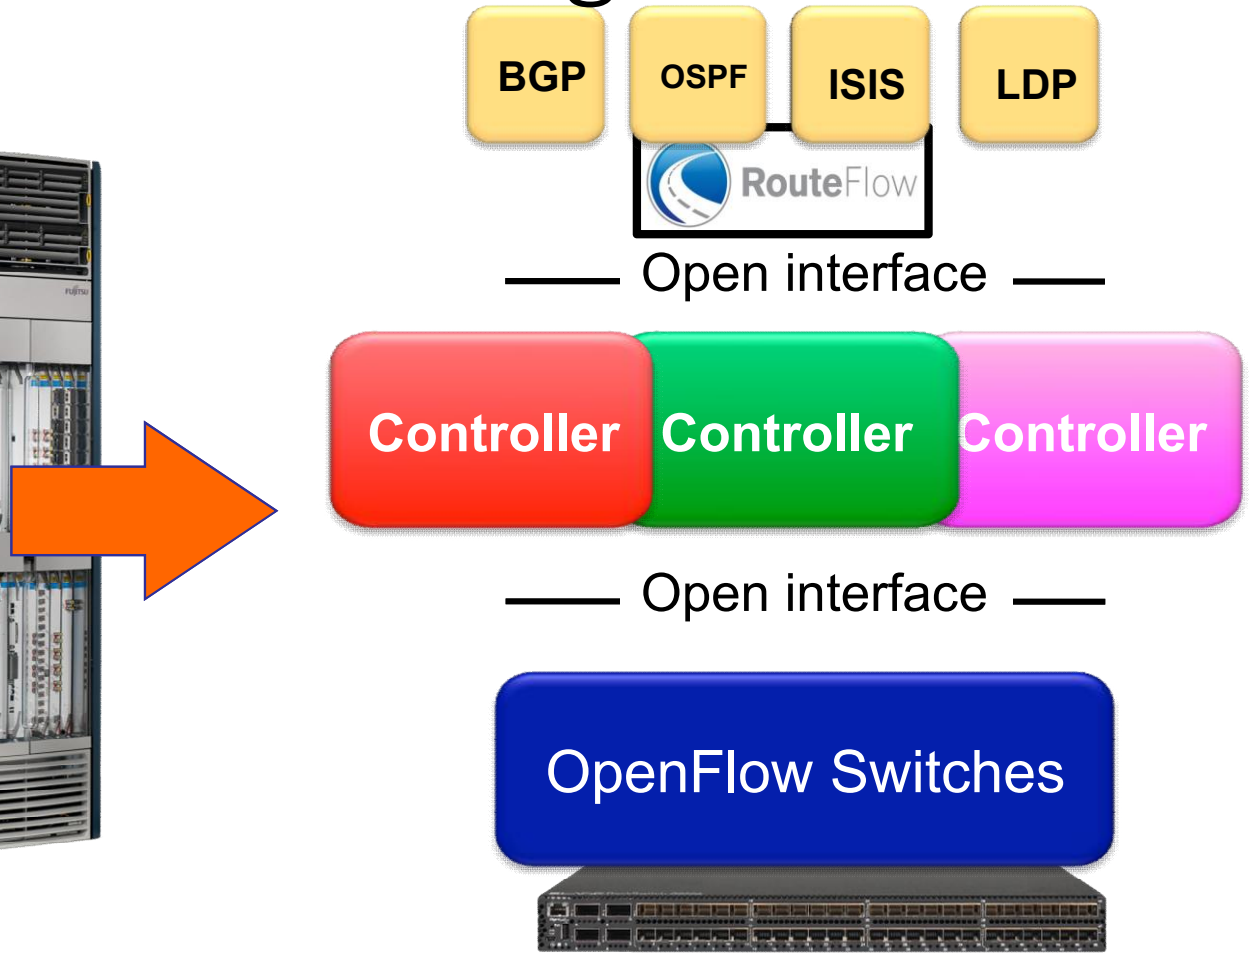

Low cost (commodity)
Multi-vendor modularity
Open source
Rapid innovation

Source: McKeown

Research on SDN Route Control

- Early work on Routing Control Platforms (RCP)
  [Ramjee 2006, Feamster 2004, Van der Merwe 2006, Wang 2009]
  - In operation at AT&T, considered a differentiator for "dynamic connectivity management".
- Research Question:
  - Re-examine the concept of RCP with the visibility (i.e., network-wide, multi-layer, flow and topology maps, full RIBs) and direct control capabilities (i.e., actual FIB installation, rich matching and instruction set) of the SDN abstraction set and the specifics of the OpenFlow choice
- RouteFlow glues virtualized IP routing stacks with OpenFlow
- RouteFlow acts as a new indirection layer for
  - routing protocol messages (e.g. BGP session terminates in servers)
  - RIB-(to-FIB)-to-OpenFlow transformations

Routing-centric research use cases

- Engineered path selection
  - Think Google WAN, performance-based routing, etc.
- Optimal best path reflection
  - Per ingress/customer [draft-ietf-idr-bgp-optimal-route-reflection-01]
- Path protection with prefix independent convergence
  - Hierarchical FIBs w/ OF 1.X Tables + LFA route-precomputation
- Security
  - Data plane blackholes and middlebox injections,
  - Secure Inter-domain routing ideas (crypto intense S*-BGP, etc..)
- Simplifying customer multi-homing
  - Easy to set and control cost/performance/policy-based routing
- IPv6 migration
  - Flow matching for service termination in v4-v6 migration solutions

Seamless MPLS / MPLS-lite / IP Traffic Engineering

Traffic Statistics (current & historic) : sFlow + OpenFlow
Event DB (failures, alarms)
Virtual & Physical Topologies

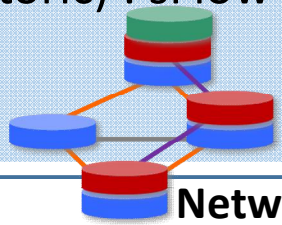

NIB

Network Information Base

TE
+ Data Mining

OF1.3
= Group Tables + Metering

Legacy control links
OpenFlow control links
Data plane links

AS 8002
R82
Edge
R1
R3
Cust. A CE
Cust. B CE

AS 1000
Control Plane Core:
RR + RFCP
Data Plane Core:
MPLS Cloud

AS 8003
R83
Edge
DP 0x1
DP 0x2
AS 1002
R12
Cust. B CE

Recommended reading: Google B4 paper at SIGCOMM 2013

Software-Defined Internet eXchanges

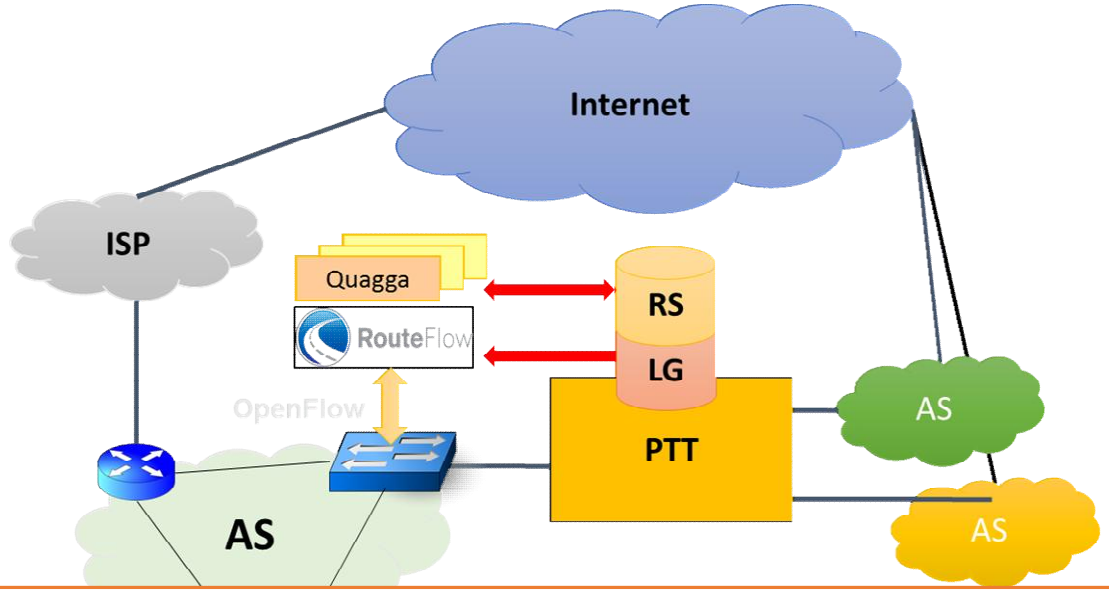

Recommended reading: Feamster et al., SDX: A Software Defined Internet Exchange at ONS13

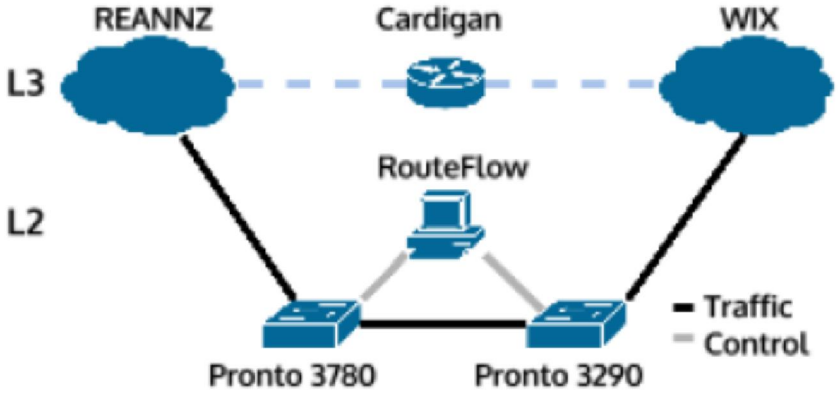

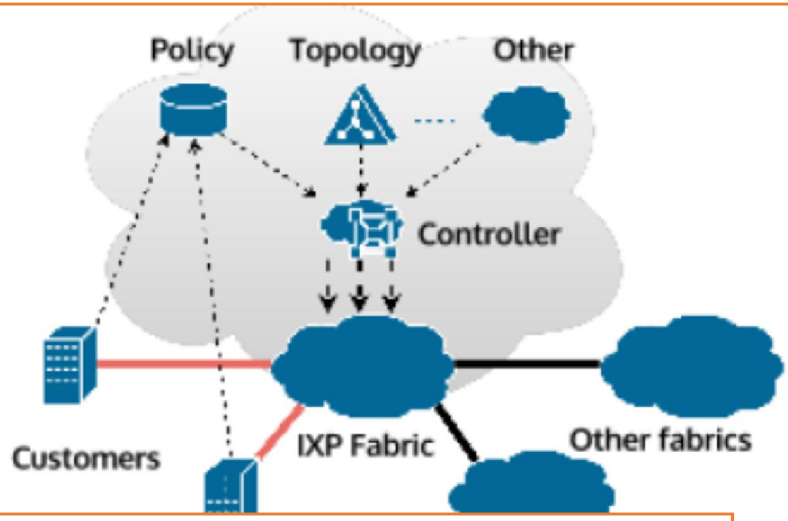

Further info on Cardigan: RouteFlow pilot deployment 9-month at NZ IXP.
Joint work with Josh Bailey (Google), REANNZ, University of Wellington, etc.

Next Generation Peering in SDN

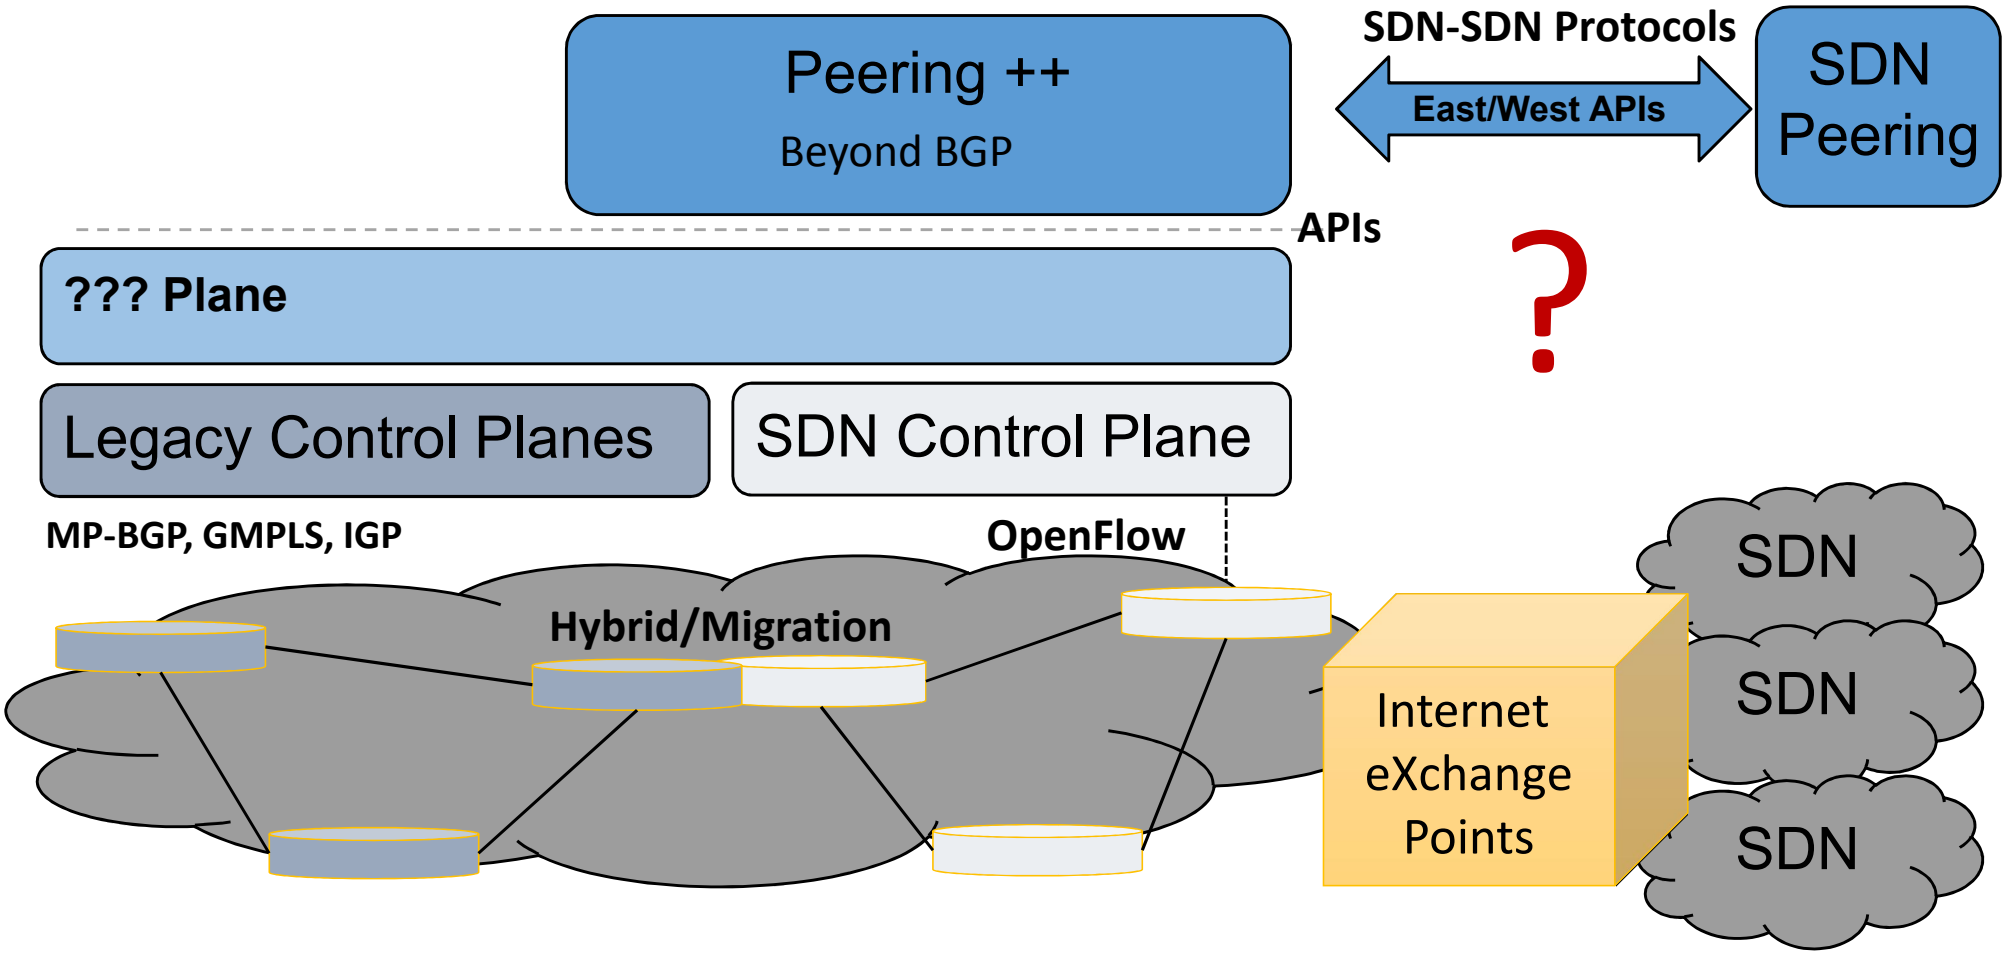

High Performance SDN Stacks

SDN Stack

Research on useful abstractions to optimize target app's performance:

- Bandwidth (Event rate)
- Latency (Event responsiveness)
- ???

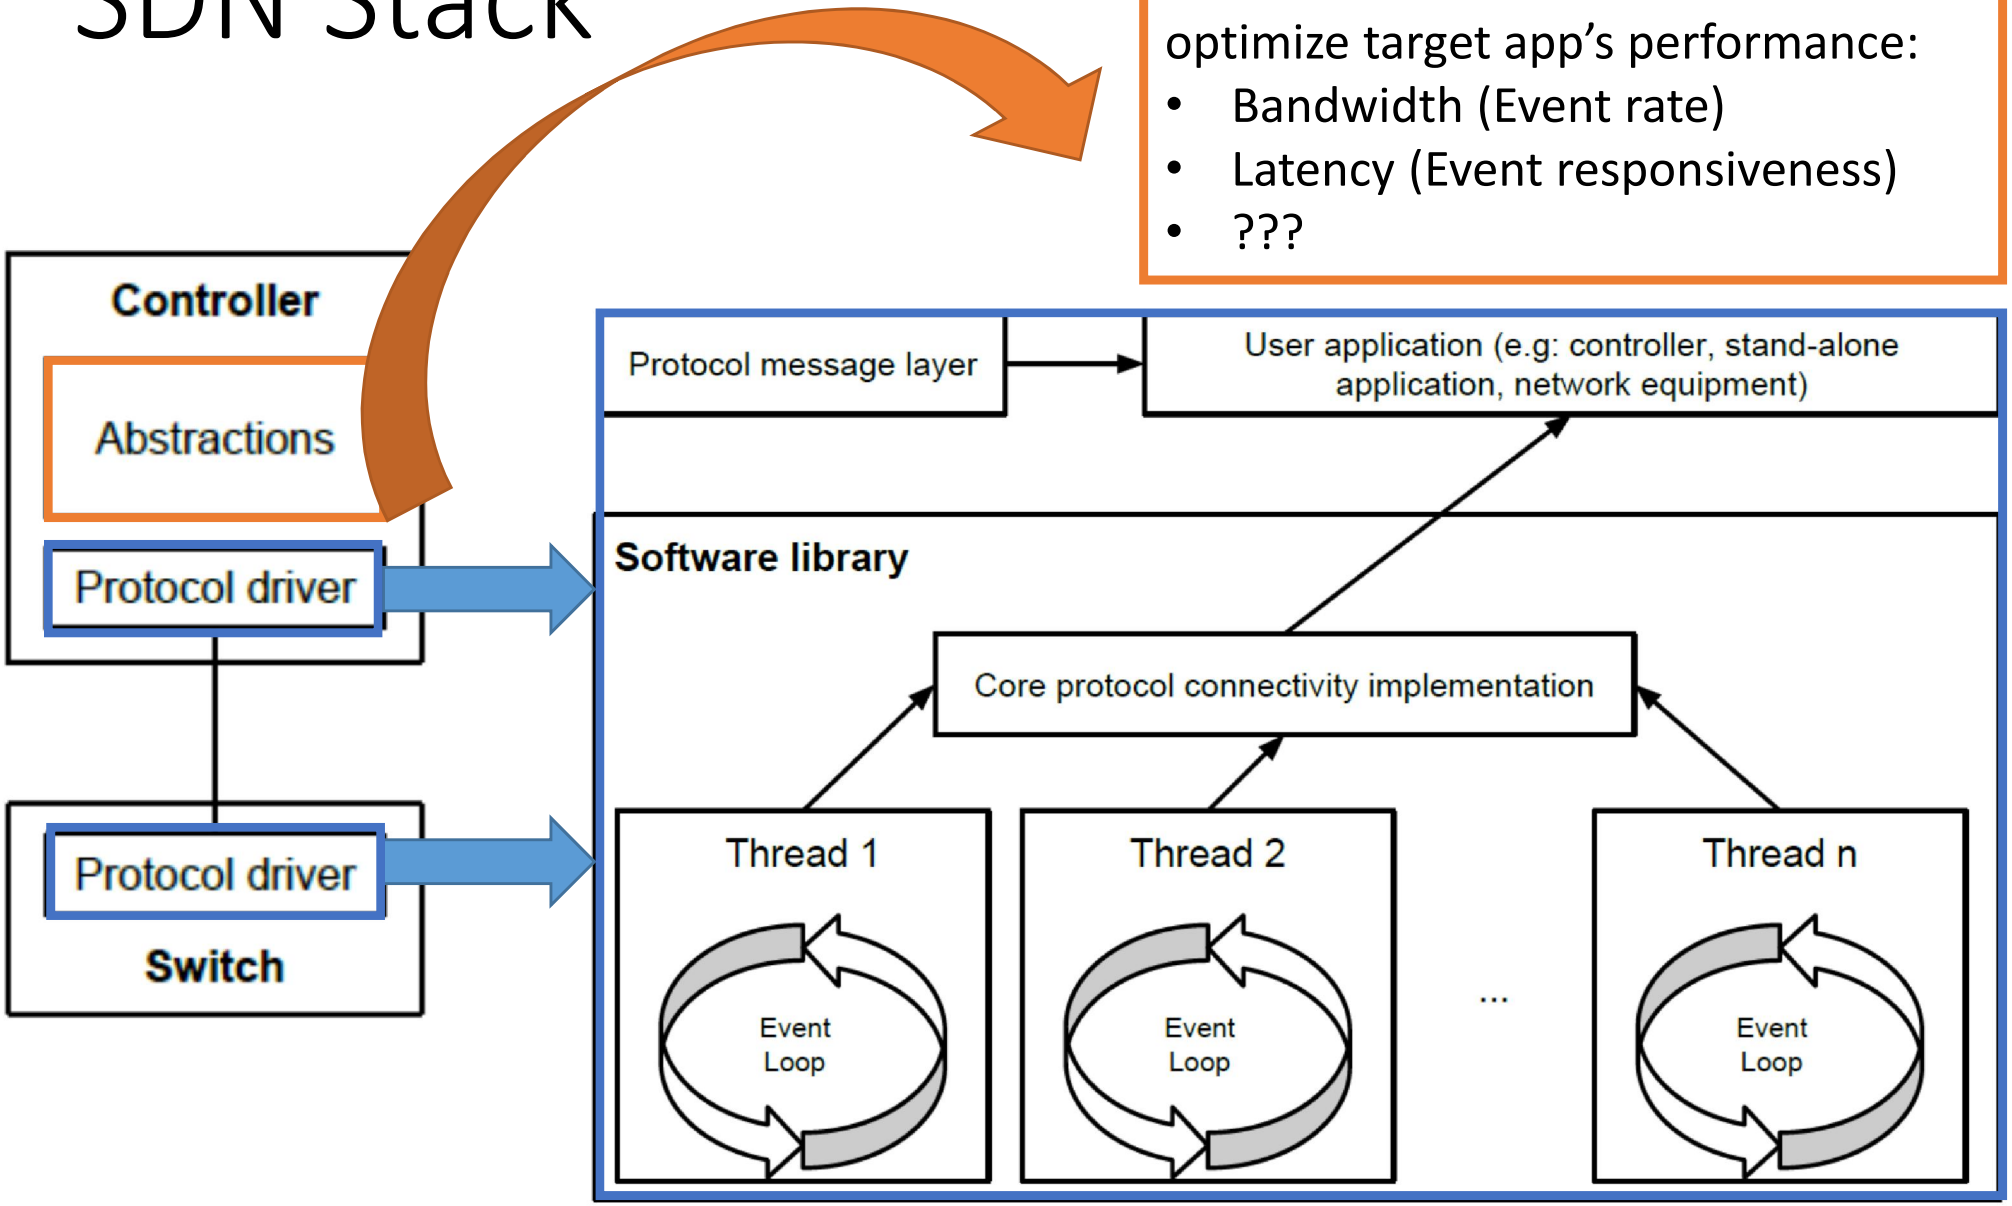

*Joint work with Prof. Fabio Verdi (Ufscar) and MSC-candidate Allan Vidal (CPqD Jr. Researcher)

libfluid ONF Driver Implementation

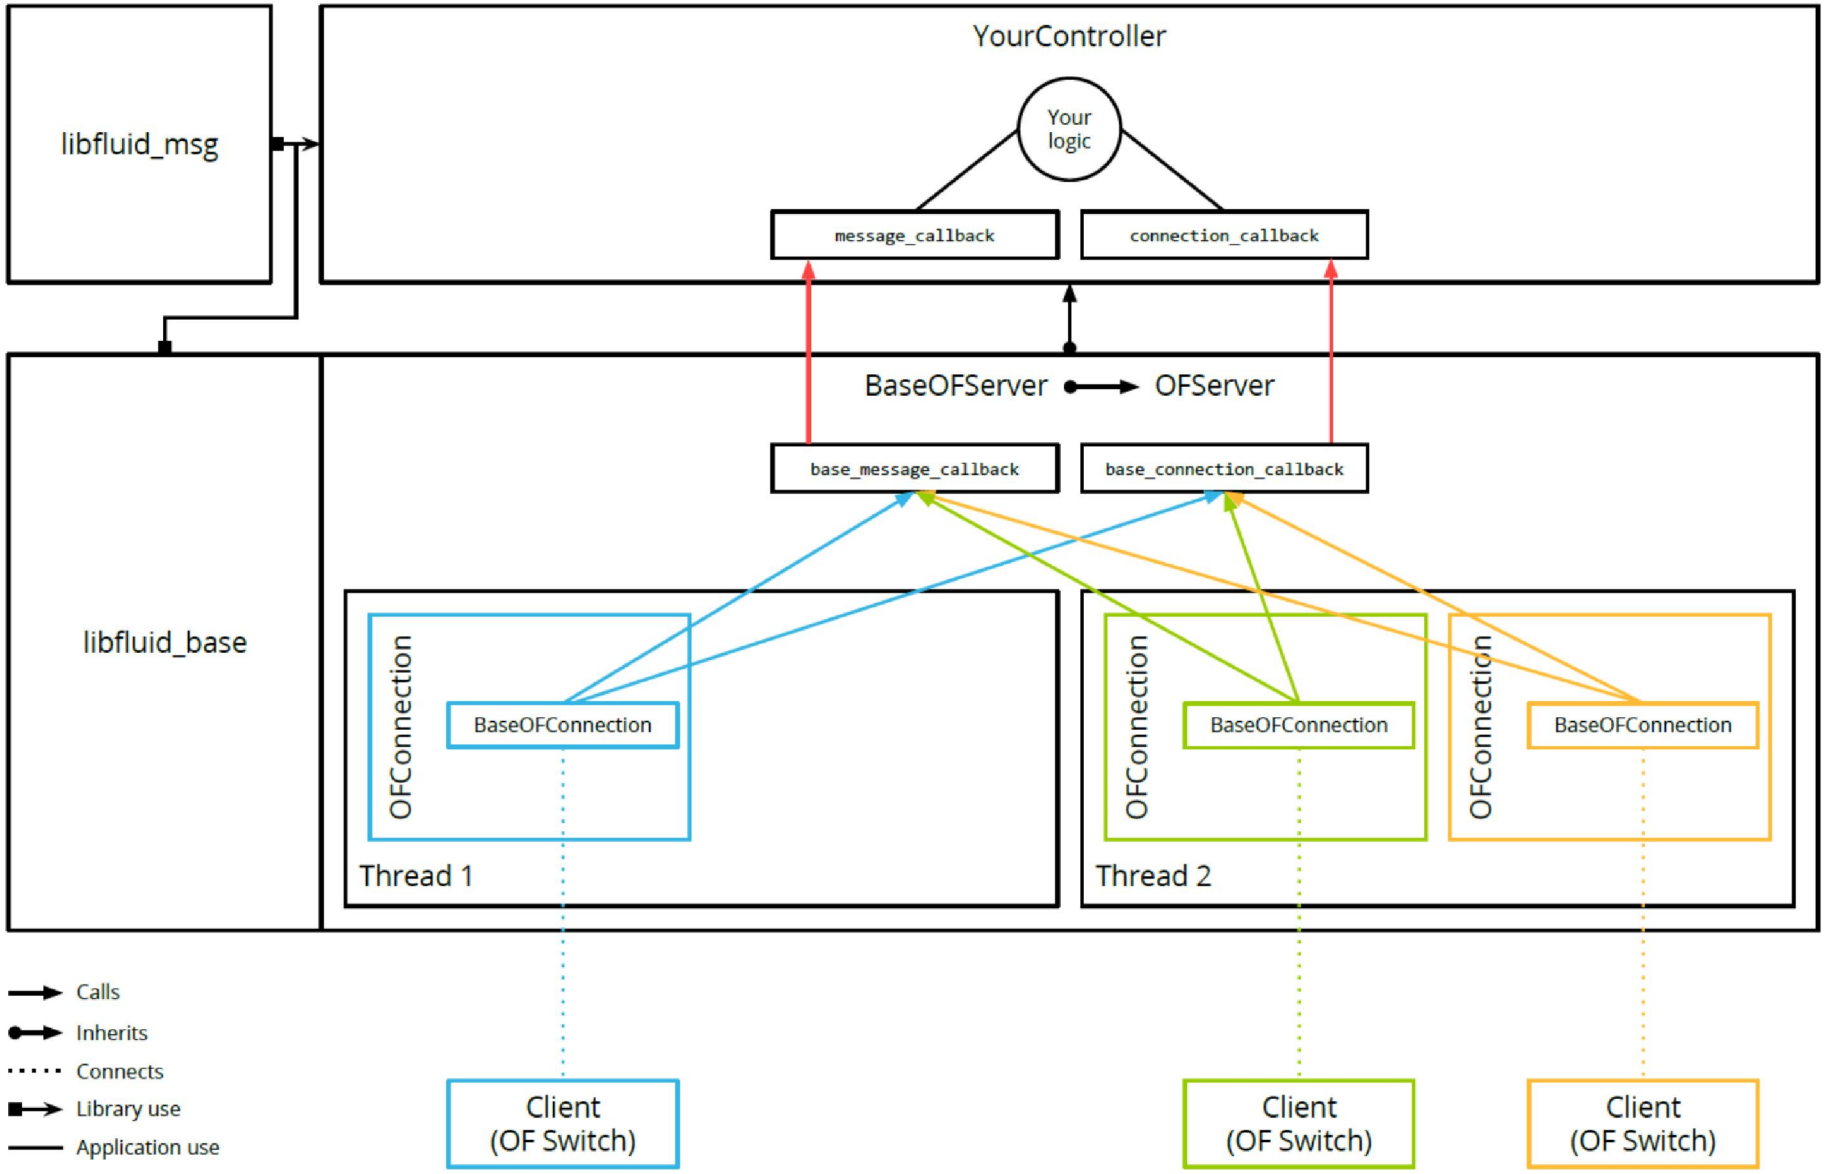

*Joint work with CPqD Jr. Researchers Eder Fernandes, Allan Vidal, and Marcos Salvador (RNP)

High-Available SDNs

A critical issue for the success of SDN technology

Motivation

- High availability of is the number one purpose of communication networks
- Logically-centralized split architectures such as OpenFlow/SDN are commonly questioned about their actual capabilities of being resilient to faults
- Any new networking technology must, at least, yield the same levels of availability as alternative and legacy technologies.

- Related work so far has focused on point solutions to some flavour of SDN architectures, mostly OpenFlow-only.
- No prior work is based on theoretical foundations that provide generally applicable architectural recommendations with proof of concept implementations experimentally validated.

Convergence in SDN Network

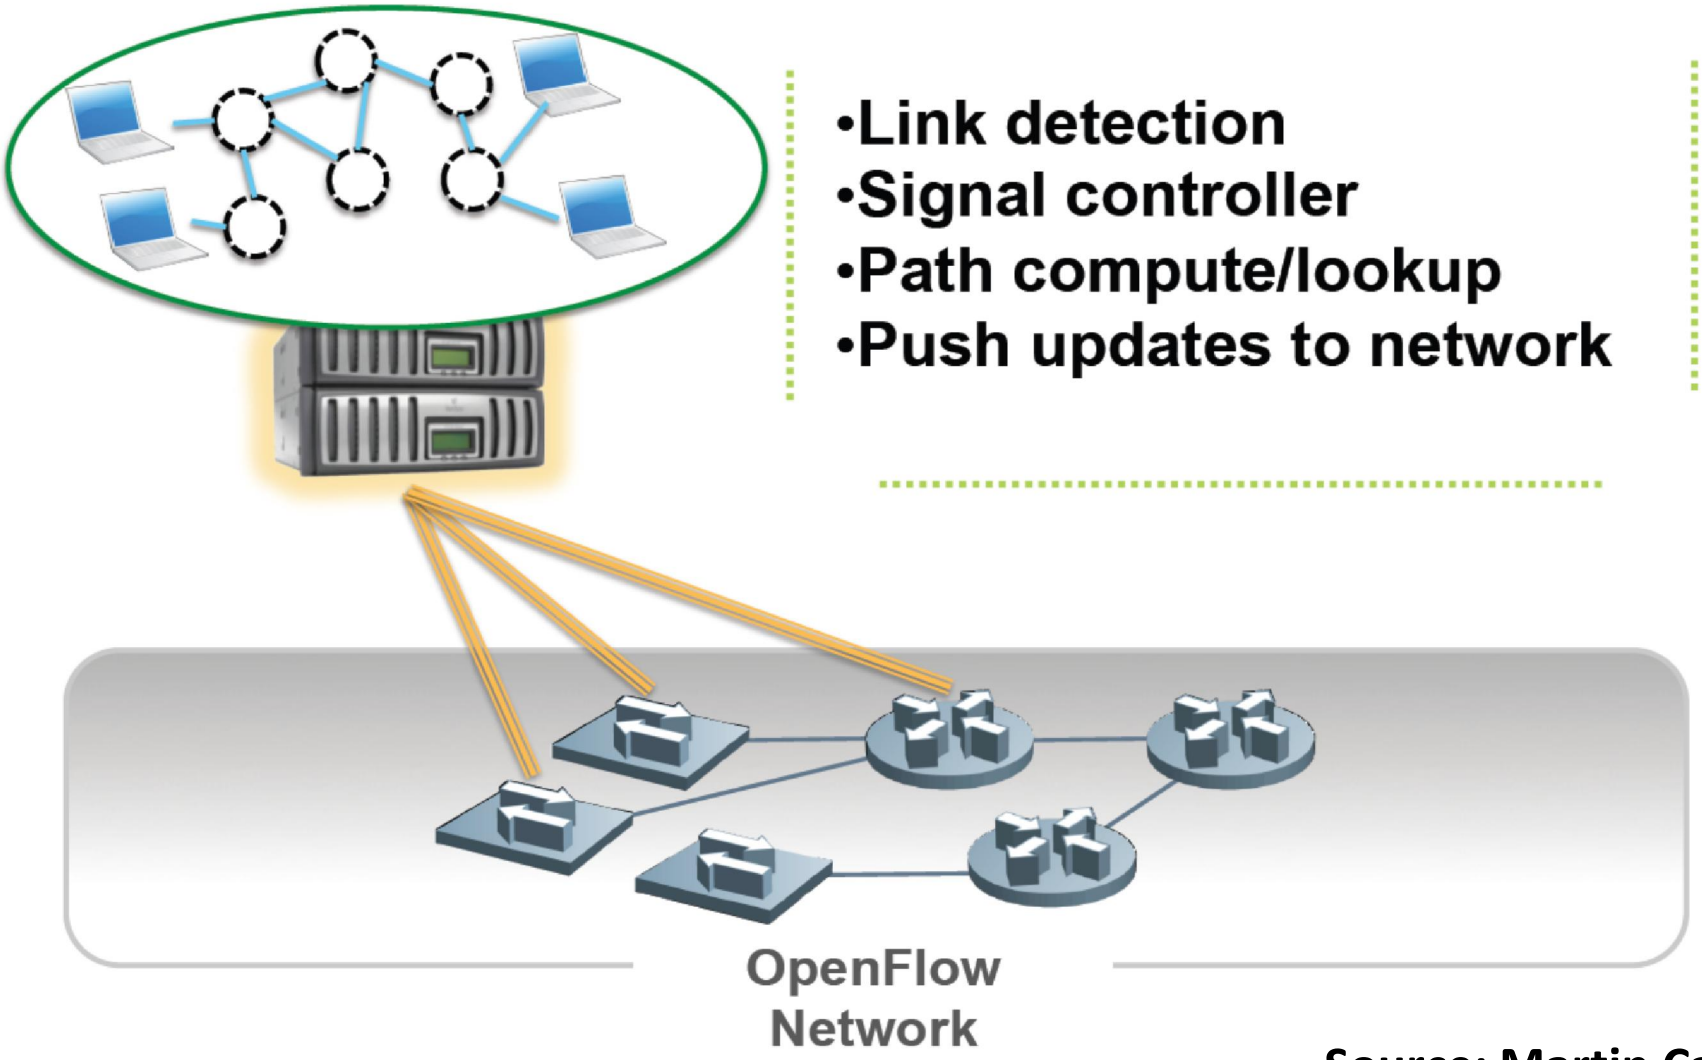

Source: Martin Casado
http://networkheresy.com/

Convergence Traditional IGP vs SDN

Traditional IGP

- Link detection
- Flood link state
- FIB computation
- Push FIB to hardware

SDN

- Link detection
- Signal controller
- Path compute/lookup
- Push updates to network

SDN Can Perform as Well or Better

SDN Optimizations

- Incremental path computation
- Pre-compute backup paths
- Zero-hop control network
- Only update affected switches
- Use OOB multicast for control dissemination

Source: Martin Casado
http://networkheresy.com/

Approach

Resilience of an OpenFlow network depends on

- fault-tolerance in the data plane (as in traditional networks)
- high availability of the (logically) centralized control plane functions, including the control-to-data plane network

Research Objectives

1. Derive the theoretical models of availability in split architectures
2. Analyze and model the solution space to provide fault-tolerance to all architectural components.
3. Design a resilient architecture considering end-to-end requirements and all layers/domains in the SDN stack
4. Implement proof of concept prototypes and validate them in real hardware testbeds with real networking traffic.

Early results and ongoing work

Towards high-availability in SDN

High Availability

Shadow Virtual Control Planes
Backup VMs

Distributed DB Cluster
State-less RFServer

Controller Cluster Architecture

- **M:N OpenFlow Controllers: Master / Slave / Equal**
- **Distributed Environment**
  - **Group of Switches / Network Domains**

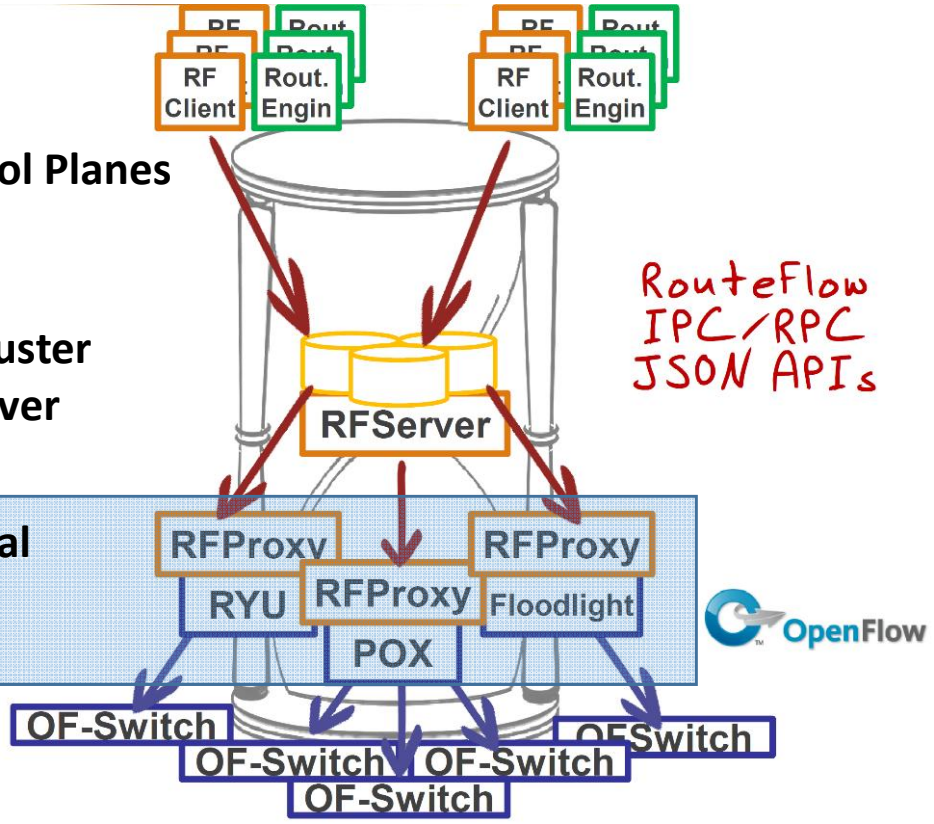

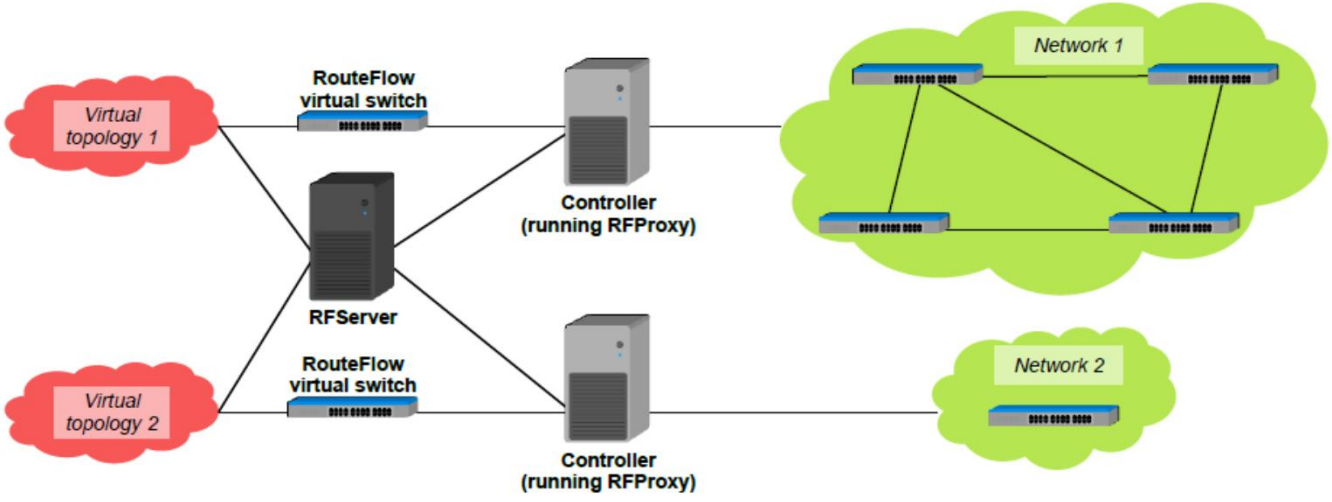

SlickFlow

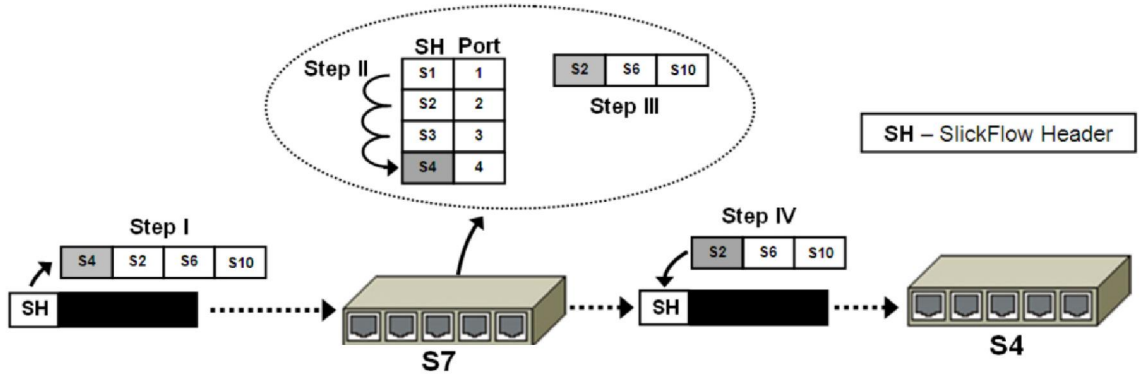

SlickFlow: Resilient Source Routing in Data Center Networks

- Joint work with Prof. Martinello (UFES). In IEEE LNC'13

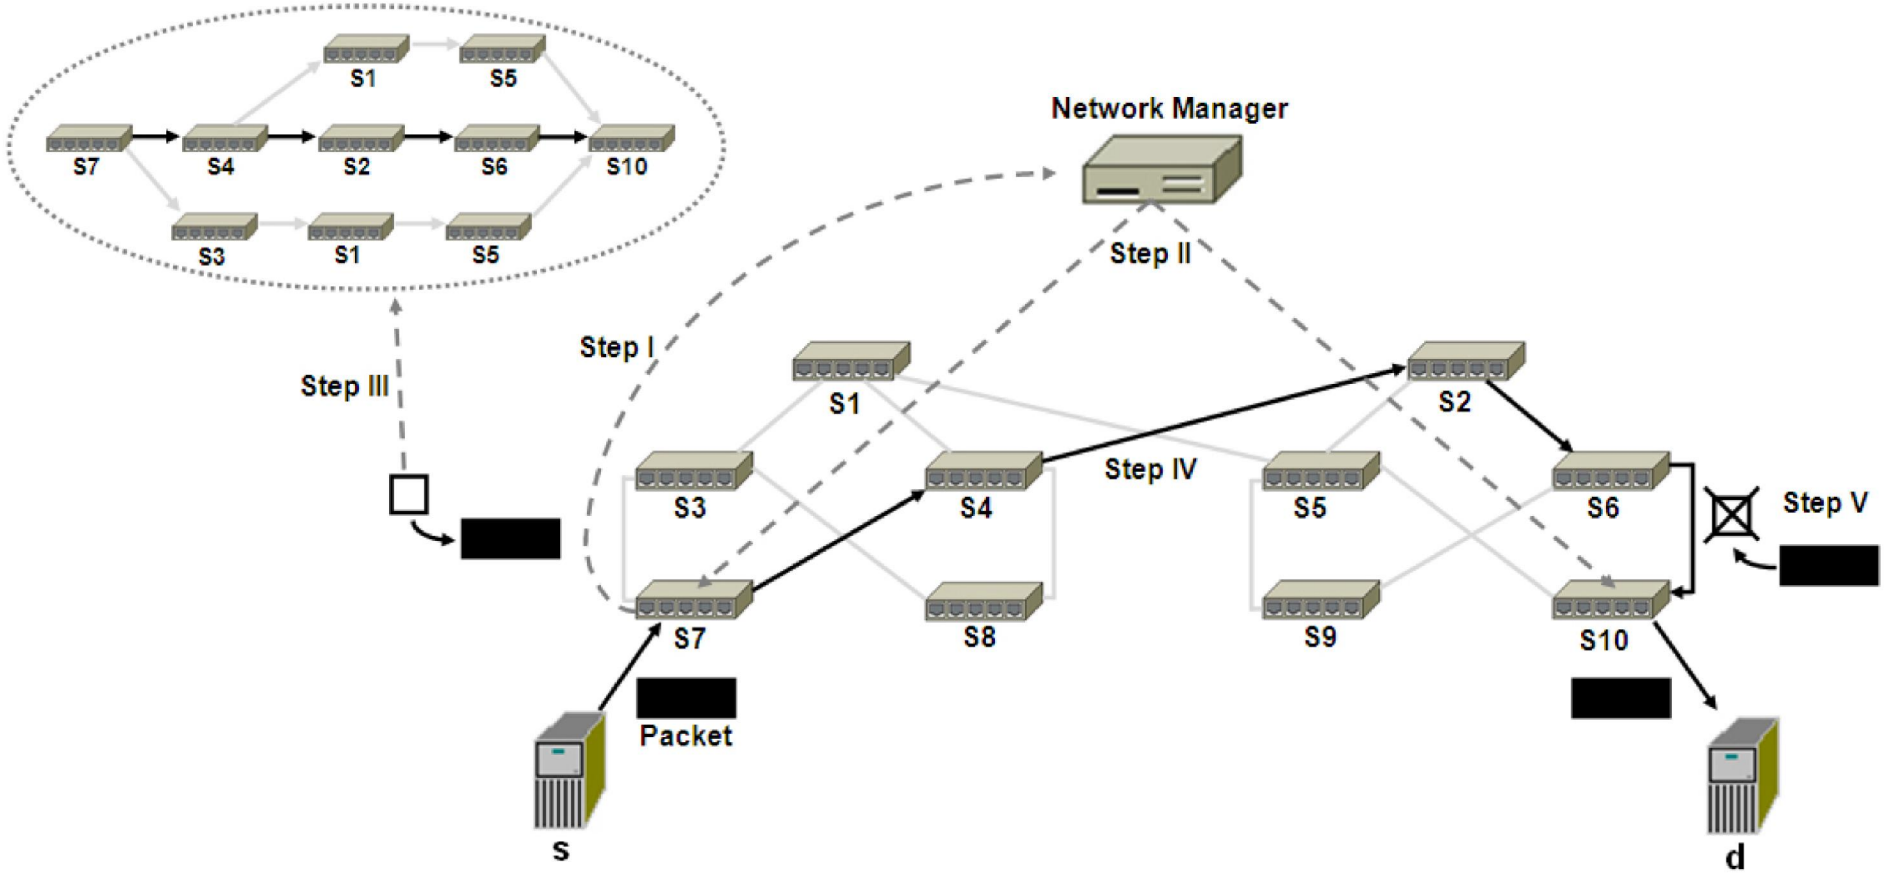

SDN and the evolution of Optics & Electronics

Abstractions and Datapath Programmability

Relative Costs of Optics & Electronics

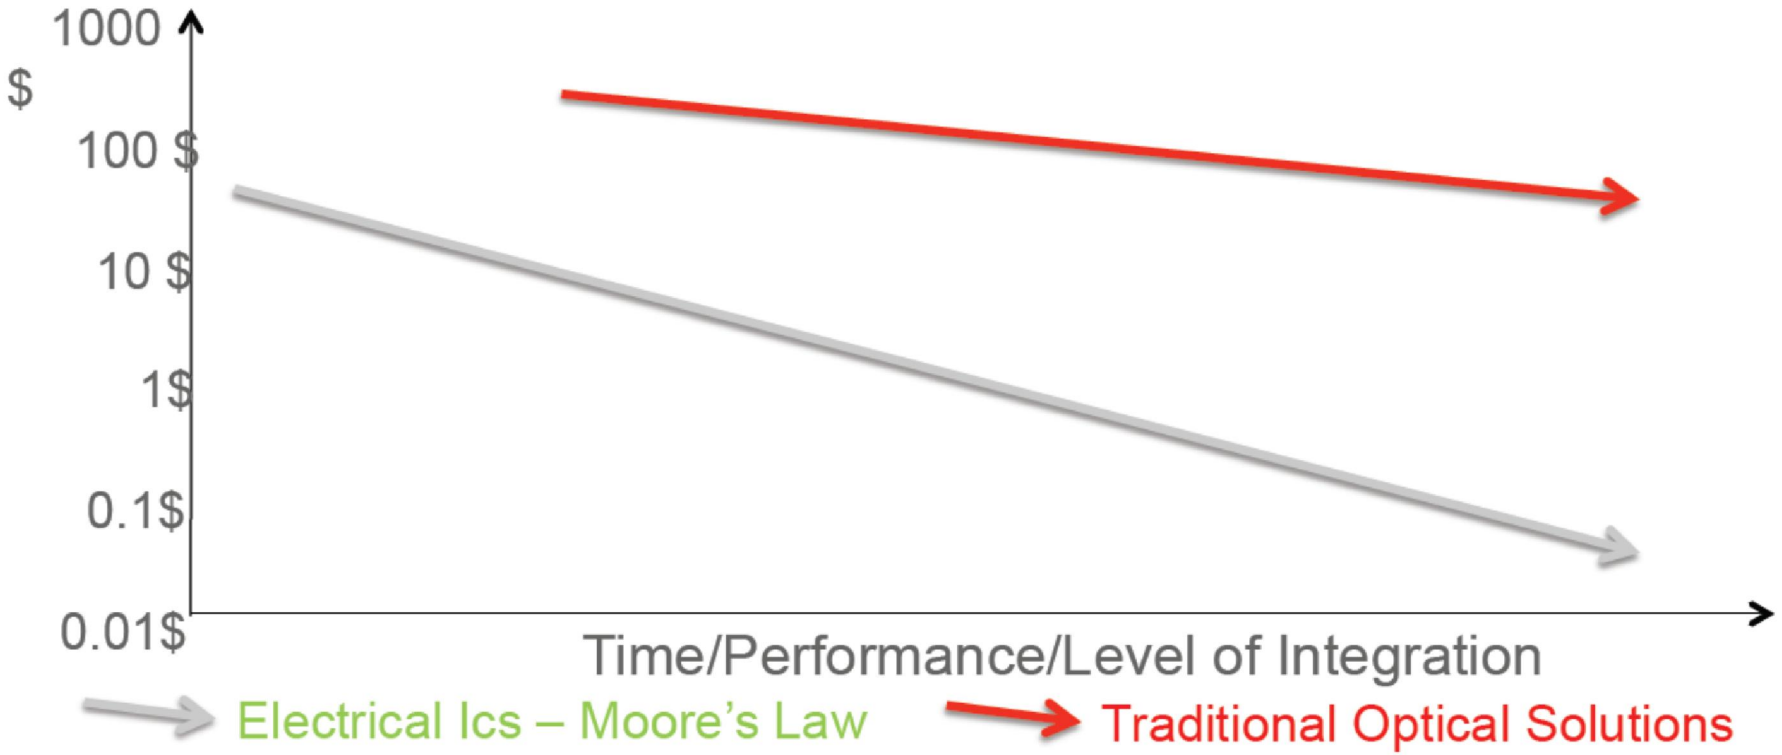

Optics ~25 years behind Electronics – level of integration, manufacturability etc & Gap is widening due to orders of magnitude larger investment in infra-structure

Source: Richard A Steenbergen (ras@gtt.net)

The increasing role of software

- So, my network device should be cheaper over time. But where is the money going?
  - If the hardware needed to forward Terabits is actually commodity (from merchant silicon to merchant optics?)
- That means what you're actually buying is software.
  - Software is hard.
  - Routing protocols, CLIs, network management platforms, and feature after feature after feature after feature…
- Software is what you're actually buying.
  - The hardware is just a delivery vehicle, so you don't feel so bad for spending millions on invisible electrons.
  - But the software is what you actually care about.

Source (adapted from):
https://www.nanog.org/sites/default/files/wednesday.general.temkin.panel.pdf

Potential Datapath Evolution

Exploring Electrical-Optical Integration

& Datapath Programmability

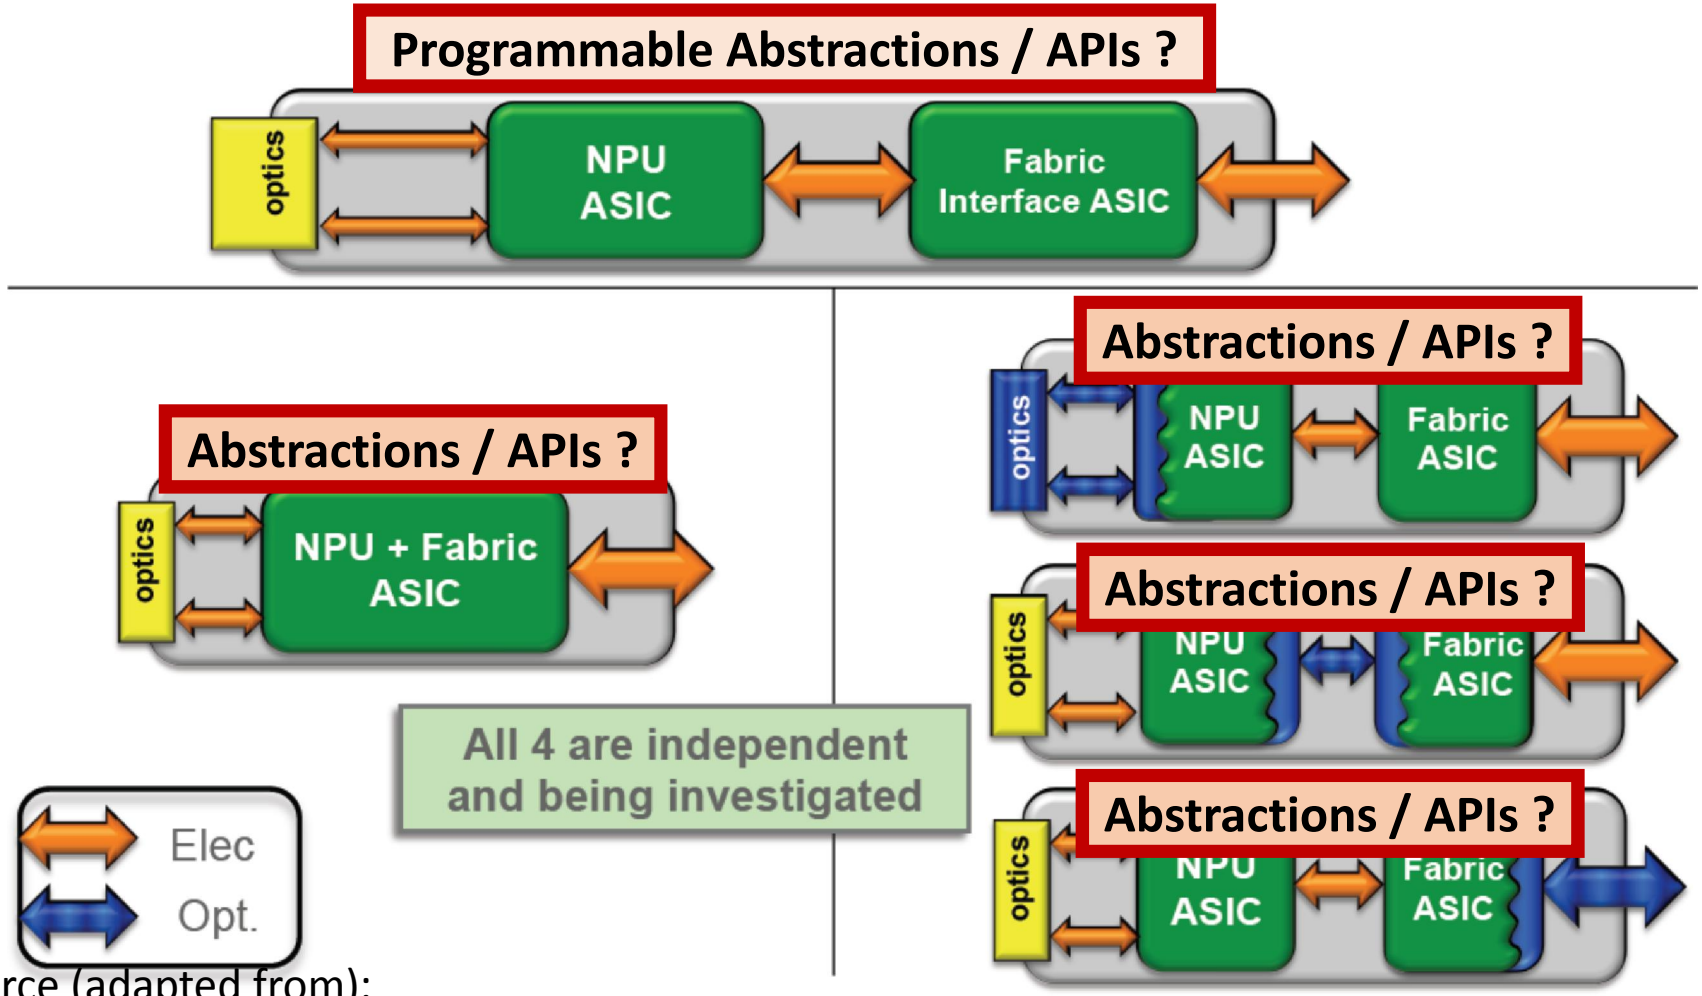

Source (adapted from):
https://www.nanog.org/sites/default/files/wednesday.general.temkin.panel.pdf

Thanks!
Questions?

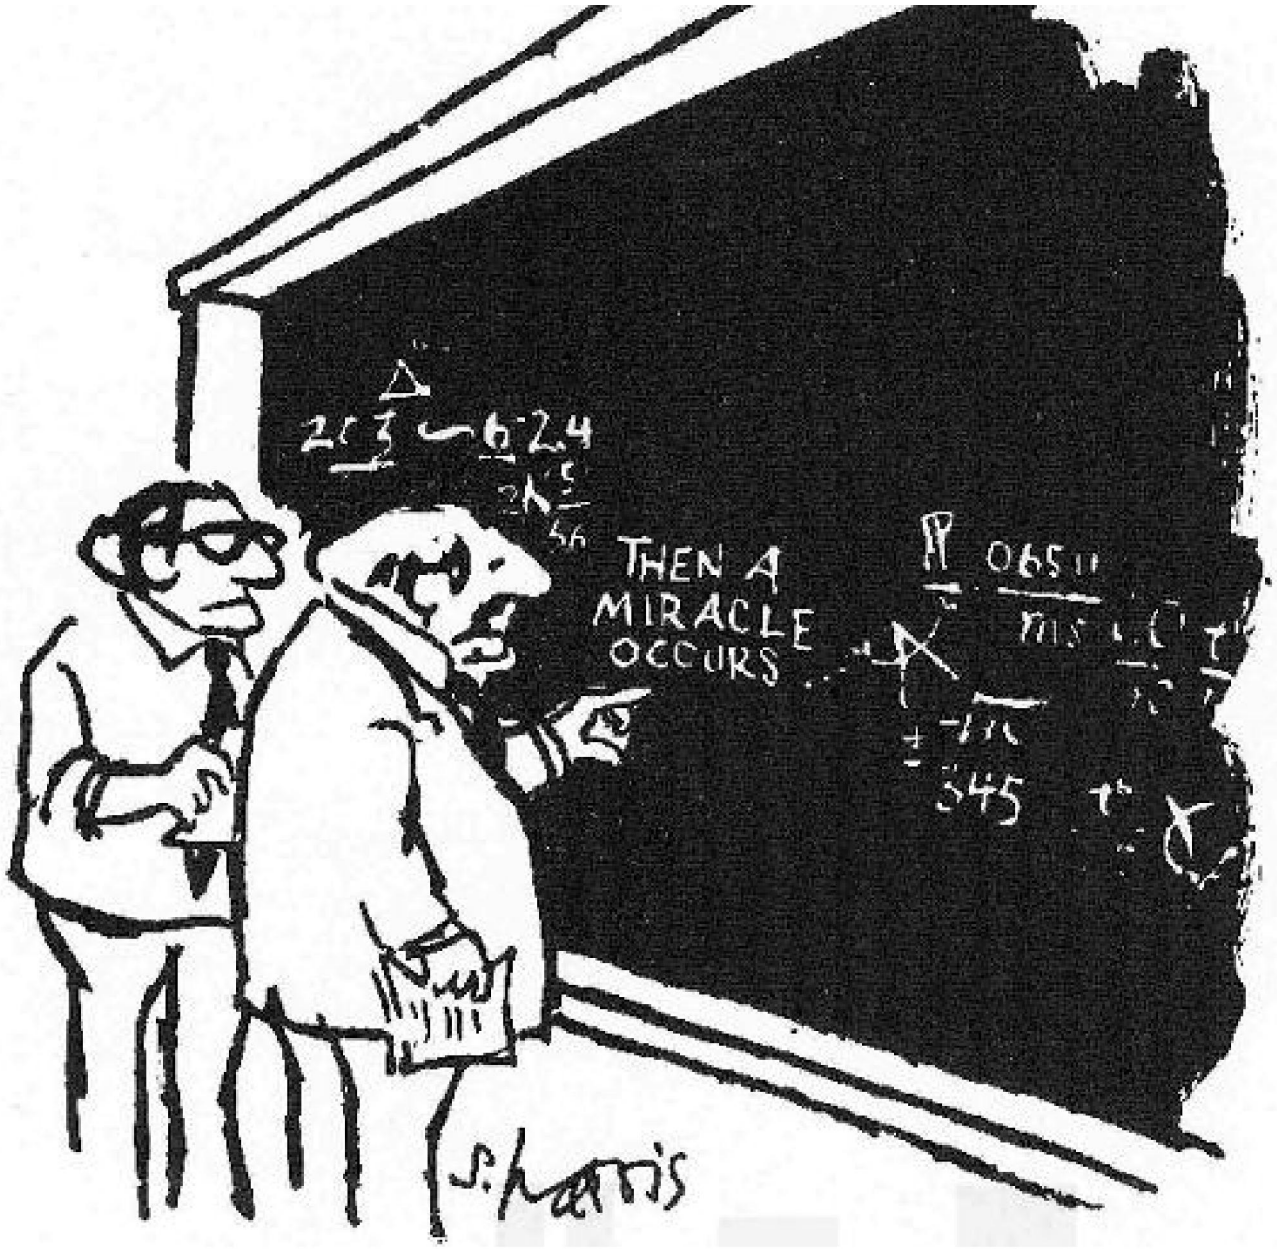

“I think you should be more explicit here in step two”