HyNS: Simulating Hybrid Networks and Beyond

Neelakandan Manihatty-Bojan*, Noa Zilberman, Andrew W. Moore
University of Cambridge, firstname.lastname@cl.cam.ac.uk

Abstract

The scalability of current data centers is limited by high power consumption, inflexible bandwidth scaling options and the cost associated with electrical packet switches (EPS) used by the network infrastructure. Hybrid switching architectures, using both EPS and circuit switches, were proposed as an alternative addressing these challenges. The evaluation of hybrid networks is challenging, as the availability of evaluation and reproduction tools is limited. This work presents HyNS: a flexible hybrid network simulation framework that is highly scalable, accurate and reflects properties of real networking hardware that is often ignored. Consequently, HyNS provides new insights into hybrid networks performance and explores new architecture schemes.

Motivation

A hybrid network uses a circuit switched network alongside an EPS (aggregation and core switches) network [1, 3, 2]. The circuit switch is used to serve long flows and the packet switch is used to serve short flows [1]. Building hybrid networks for high performance is challenging because scheduling requires understanding the complementary properties of the underlying circuit and packet switched networks and mapping the traffic accordingly. Most of the existing work (e.g. [1, 2]) focus mainly on static or fixed traffic matrices.

Data center network (DCN) fabrics have become akin to big switches, though at a higher level of complexity and scale. The similarity means that insights gained from switch architecture can be applied to DCN. Through the abstraction of hybrid network architectures to look similar to a simple switch architecture containing a circuit and packet switched network, the application can remain agnostic to the underlying network fabric.

OCS has no buffers and requires reconfiguration whenever the source-destination pair of a flow changes. This implies that traffic through the OCS needs to be buffered either at the host or the ToR switch when the OCS is reconfigured. Migrating flows between circuit and packet switched networks is a compelling idea to try in hybrid networks. While novel ideas can be explored, we lack a simulator that can precisely model hybrid networks. The design choices in hybrid networks require accurate modelling of the properties of underlying network elements, else important aspects of its behavior are missed. A simulation framework that captures the properties of real networking hardware is essential for advances in hybrid networks design and adoption.

*Student and presenter

HyNS

To date, there is no available, open-source, hybrid network simulator that accurately models current day, high end networking devices. While hybrid network simulators do exist, they are mostly proprietary, architecture specific, or gloss over lower level details that cannot be neglected in real systems. Examples for switching architecture aspects often neglected by simulators include the effect of packet size alignment to an internal bus width, header processing delay and OCS reconfiguration recovery time. All these lead to inaccurate throughput and latency simulation results. For evolving network fabrics such as hybrid networks or rack scale fabrics that integrate network devices with different properties, these have an important effect. We aim not only to explore novel data center fabric architectures, based on above considerations, but also to enable reproducible experimental results.

HyNS is an accurate network topology simulator scaling to millions of queues, providing as close as possible modelling of a real network environment. It provides a flexible, easy to use environment for modelling of different topologies, configurations and types of networks. HyNS is not a protocol-level simulator: it focuses on the hardware implementation of the network, and receives as an input both user-specified traffic loads and different pre-defined stimulus (e.g. CDF, constant bit rate, poisson). The use of the simulator is not limited to just hybrid networks. We intend to extend its use to explore efficient networking topologies in rackscale computing, comparing and assessing different approaches and existing topologies [2, 3]. HyNS will be validated using NetFPGA SUME [4], using real traffic loads.

Acknowledgments

This project is supported by the EPSRC (INTERNET Project EP/H040536/1, NAAS Project EP/K034723/1).

References

[1] FARRINGTON, N., ET AL. Helios: A hybrid electrical/optical switch architecture for modular data centers. In *SIGCOMM* (2010), ACM.

[2] LIU, H., ET AL. Circuit switching under the radar with reactor. In *NSDI* (2014), USENIX.

[3] WANG, G., ET AL. c-Through: part-time optics in data centers. *SIGCOMM CCR* (2010).

[4] ZILBERMAN, N., ET AL. NetFPGA SUME: Toward 100 Gbps as Research Commodity. *IEEE Micro* (2014).

UNIVERSITY OF CAMBRIDGE

EPSRC
Engineering and Physical Sciences Research Council

HyNS: Simulating Hybrid Networks and Beyond

Neelakandan Manihatty-Bojan, Noa Zilberman, Andrew W. Moore

Problem

1. Evaluating hybrid Data Center Networks (DCNs) is challenging.
2. Limited availability of evaluation and reproduction tools.
3. 1 & 2, limits the exploration of novel hybrid architectures and scheduling schemes.

Motivation

- ❑ Future DCNs requires more bandwidth, energy efficiency and scalability [1,3].
- ❑ Emerging hybrid switching architectures .
- ❑ Dynamic hybrid schedulers are needed [2,3].
- ❑ Need insights about hybrid network and scheduler performance.

Hybrid DCNs

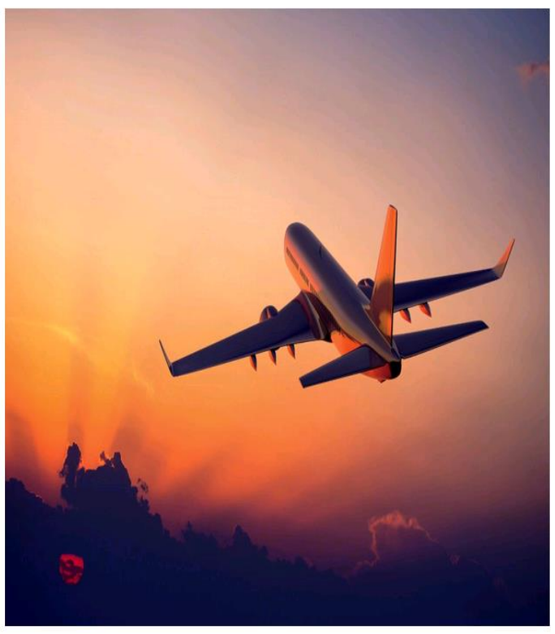

Hybrid DCN = Packet Switching + Circuit Switching

Existing Approaches

- ❑ Lacks good model of hybrid networks.
- ❑ Don't capture low level hardware details.
- ❑ Low scalability

HyNS

A flexible, accurate event driven simulator.

- ✓ Accurately models hybrid networks.
- ✓ Captures properties of low level network elements.
- ✓ Reproducible experimental results.
- ✓ User defined traffic loads, pre-defined stimulus (based on CDF, distributions).

HyNS - OVERVIEW

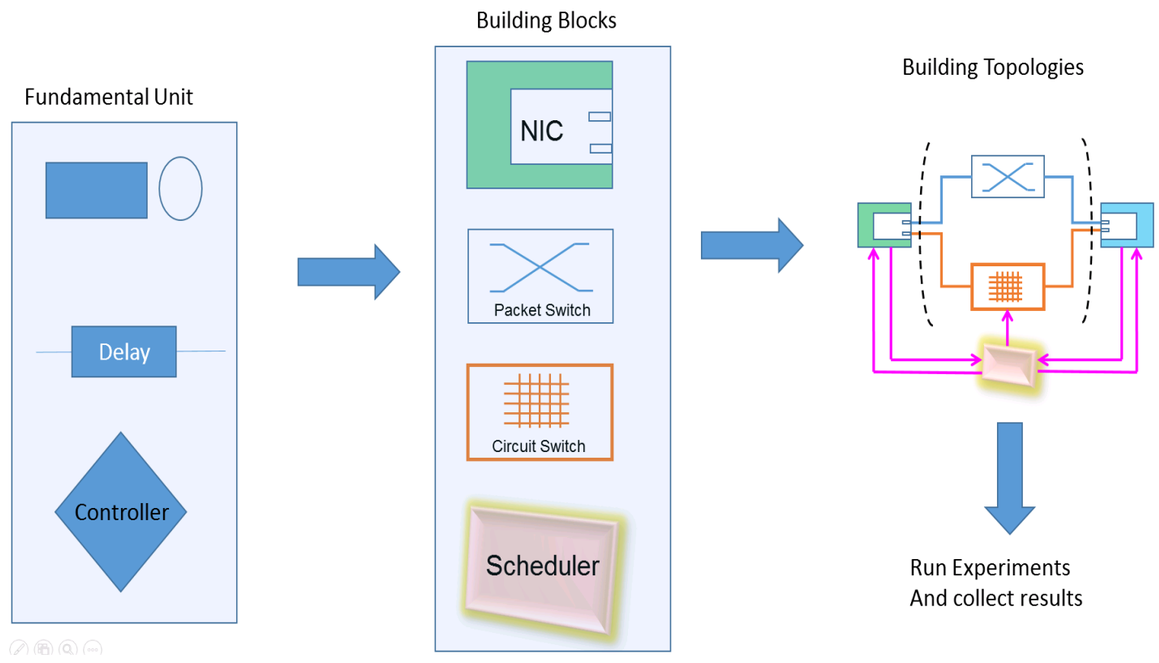

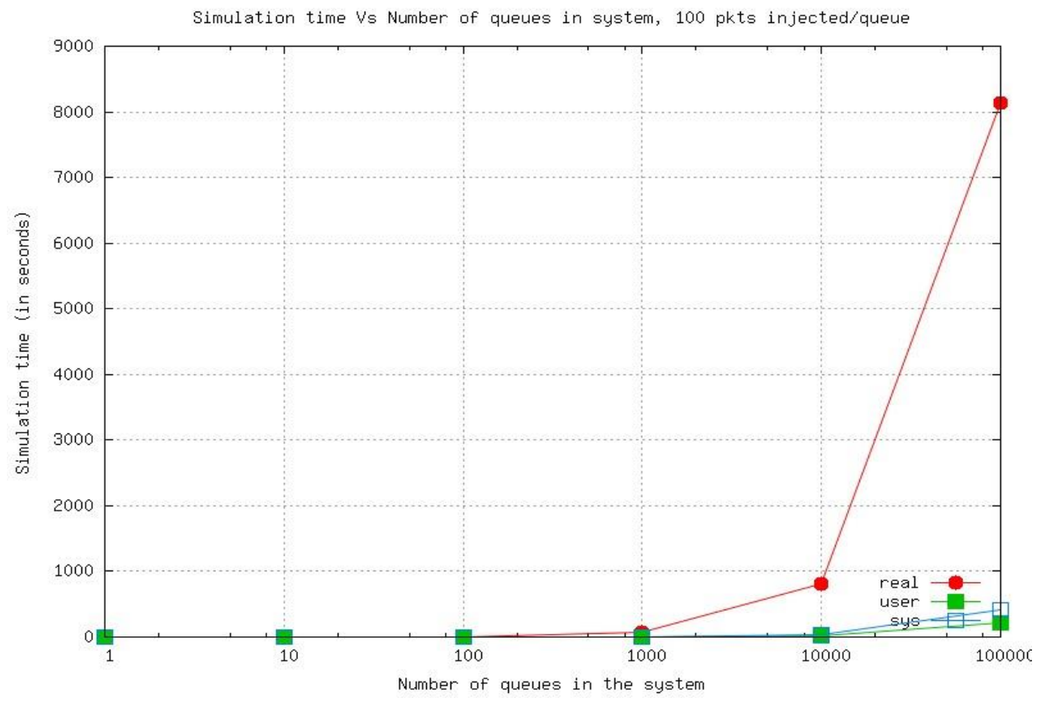

Validation Apparatus: NetFPGA-SUME [4]

- ✓ Analyze transient effects
- ✓ Real traffic workloads

Applications

- ➢ Evaluating scheduling algorithms.
- ➢ Proposing new scheduling schemes.
- ➢ Optimising buffering in the host.
- ➢ Exploring hybrid DC fabric topologies.
- ➢ Exploring rack scale fabric topologies.
- ➢ Exploring new traffic management schemes

References

[1] G. Porter et al. Integrating Microsecond Circuit Switching into the Data Center. SIGCOMM CCR, 2013.
[2] H. Bazzaz et al. Switching the Optical Divide: Fundamental Challenges for Hybrid Electrical/Optical Datacenter Networks, SOCC, 2011
[3] H. Rodrigues et al. Accurate Emulation of Fast Optical Circuit Switches, ICC, 2015
[4] N. Zilberman et al. NetFPGA SUME: Toward 100 Gbps as Research Commodity. Micro, 2014.

Further discussions

Email : firstname.lastname@cl.cam.ac.uk