Chair of Network Architectures and Services
School of Computation, Information and Technology
Technical University of Munich

HVNet

Hardware-assisted **V**irtual **Net**working for low-latency network services

Motivation

Tail-latency Experiments on Real Hardware

- Expensive
- Restrictions on creating arbitrary topologies
- Limited availability
- Can have long delivery times
- Difficult to reproduce

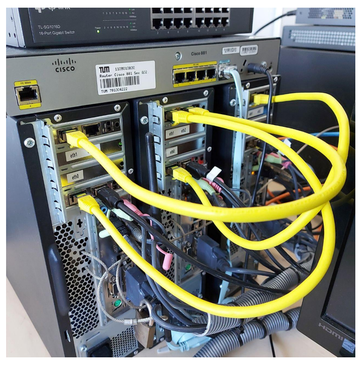

Tail-latency Experiments on Simulation, e.g., OMNet++ [5]

- Restricted in terms of usable software
- No access to hardware-level features

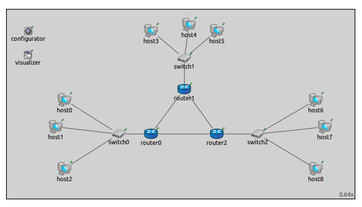

Background

Single-Root I/O-Virtualization (SR-IOV):

- Share hardware resources among VMs
- Physical Function (PF) with full control
- Lightweight Virtual Function (VF) with restricted access
- Each physical device: 1 PF, multiple VFs
- Allow direct passthrough of VFs due to be PCIe-functions

Virtualization of systems is possible using:

- Full virtualization, e.g., virtual machines (VMs)
- OS-level virtualization, e.g., container

Analyzed raise of latency is according to [2] mostly caused by:

- Interrupts raised on the same core as the virtualization
- Energy-saving mechanism during idle times
- Other applications running on the same cores as the virtualization

Measurement Setup

- Loadgen runs a packet generator (MoonGen [1]) creating UDP packets
- Device under Test (DuT) contains to be analyzed system
- Timestamper records ingress/egress traffic using passive optical traffic access points (TAPs)
  - Hardware-timestamping of entire network traffic (resolution 12.5 ns)
  - Determine worst-case latencies on a per-flow basis

Area A: Tail-latency measurements [3]

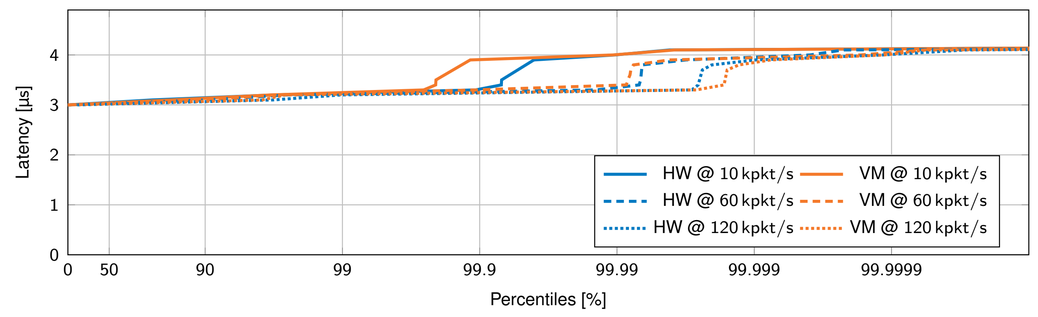

Bare-metal vs SR-IOV on VM using DPDK l2-forwarder

- Using optimizations such as to:
  - reduce interrupts
  - reduce timer-ticks
  - disable energy-saving-mechanism
  - use SR-IOV for sharing networking hardware access
- Latency between 3 and 4 µs, almost no difference between HW and VM
- Higher rates result in latency increase on a higher percentile

Area B: HVNet - Virtualization of topologies [7]

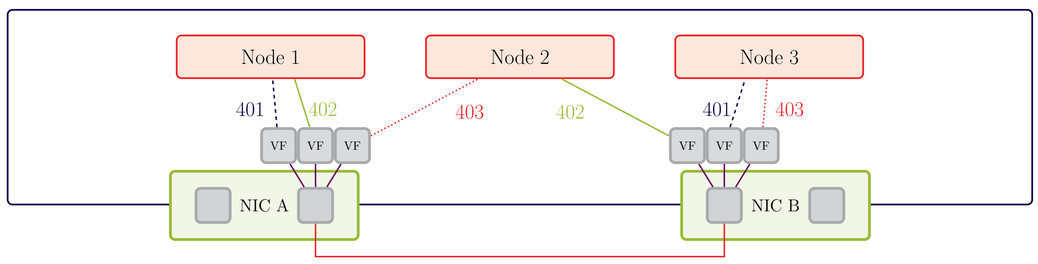

- One VLAN-ID per connection
- Split NIC into VFs
- Each packet traverses the wire per link once
- Reduce impact using optimization such as core isolation [3]

Area C: Network services and Flow measurements

HVNet vs. Mininet using 1 Mbit/s Flows:

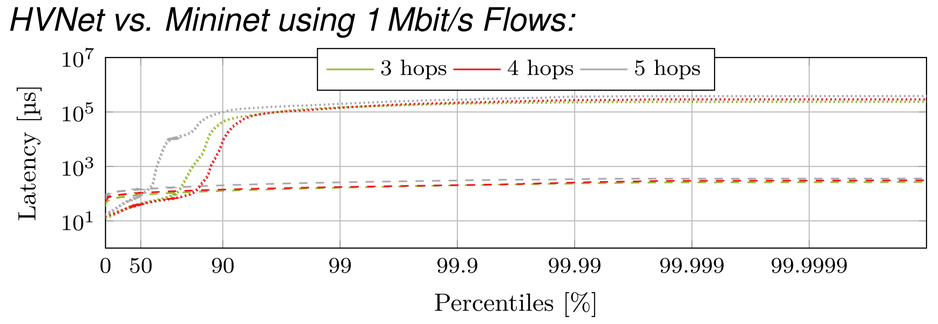

Further, we use HVNet to analyze network services such as:

- QoS-aware routing algorithms in different scenarios [6]
- TSN and TSN asynchronous traffic shaping algorithms

and use the data to predict flow-behavior, e.g., Helm et al [4].

[1] P. Emmerich, S. Gallenmüller, D. Raumer, F. Wohlfart, and G. Carle. MoonGen: A Scriptable High-Speed Packet Generator. In *Internet Measurement Conference 2015 (IMC'15)*, Tokyo, Japan, Oct. 2015.
[2] S. Gallenmüller, J. Naab, I. Adam, and G. Carle. 5G URLLC: A Case Study on Low-Latency Intrusion Prevention. *IEEE Commun. Mag.*, 58(10):35–41, 2020.
[3] S. Gallenmüller, F. Wiedner, J. Naab, and G. Carle. Ducked Tails: Trimming the Tail Latency of(f) Packet Processing Systems. In P. Chemouil, M. Ulema, S. Clayman, M. Sayit, C. Çetinkaya, and S. Secci, editors, *17th International Conference on Network and Service Management, CNSM 2021, Izmir, Turkey, October 25-29, 2021*, pages 537–543. IEEE, 2021.
[4] M. Helm, F. Wiedner, and G. Carle. Flow-level tail latency estimation and verification based on extreme value theory. In *18th International Conference on Network and Service Management (CNSM) (CNSM 2022)*, Thessaloniki, Greece, Oct. 2022.
[5] A. Varga. OMNeT++. In *Modeling and tools for network simulation*, pages 35–59. Springer, 2010.
[6] F. Wiedner, J. Andre, P. Mendes, and G. Carle. Policy-based routing for Flying Adhoc Networks. In K. R. Chowdhury and W. Jaafar, editors, *DroNet@MobiSys 2022: Proceedings of the Eighth Workshop on Micro Aerial Vehicle Networks, Systems, and Applications, Portland, OR, USA, 1 July 2022*, pages 25–30. ACM, 2022.
[7] F. Wiedner, M. Helm, S. Gallenmüller, and G. Carle. HVNet: Hardware-Assisted Virtual Networking on a Single Physical Host. In *IEEE INFOCOM 2022 - IEEE Conference on Computer Communications Workshops, INFOCOM 2022 - Workshops, New York, NY, USA, May 2-5, 2022*, pages 1–6. IEEE, 2022.

Florian Wiedner wiedner@net.in.tum.de