Chair of Network Architectures and Services
Department of Computer Engineering
Technical University of Munich

Sim2HW: Modeling Latency Offset Between Network Simulations and Hardware Measurements

Johannes Späth, Max Helm,
Benedikt Jaeger, Georg Carle

3rd Graph Neural Networking Workshop 2024
December 9, 2024, Los Angeles, USA

Chair of Network Architectures and Services
Department of Computer Engineering
Technical University of Munich

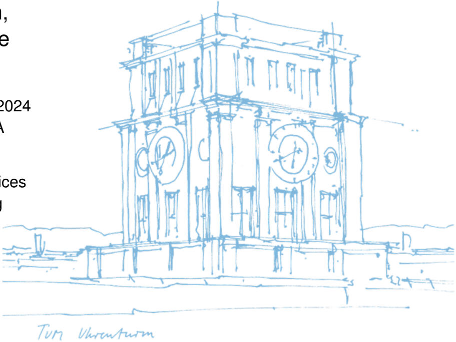

Motivation

- Network simulations are cheap, flexible, and repeatable, BUT are often not accurate
- Hardware measurements are accurate but expensive
- GNNs are well suited for modeling the behavior of computer networks
  $\rightarrow$ Might be able to bridge the gap between simulations and HW measurements

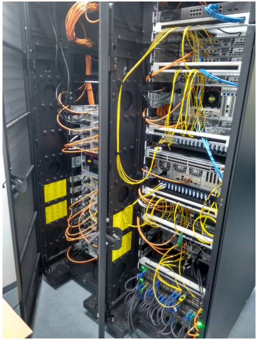

Hardware Testbed

Can we use simulations *and* a GNN to accurately approximate expensive HW measurements?

Approach Overview

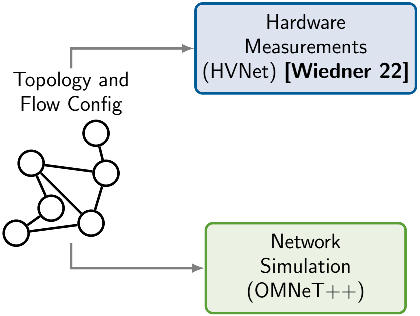

[Wiedner 22] Florian Wiedner et al. "HVNet: Hardware-Assisted Virtual Networking on a Single Physical Host". 2022.

Approach Overview

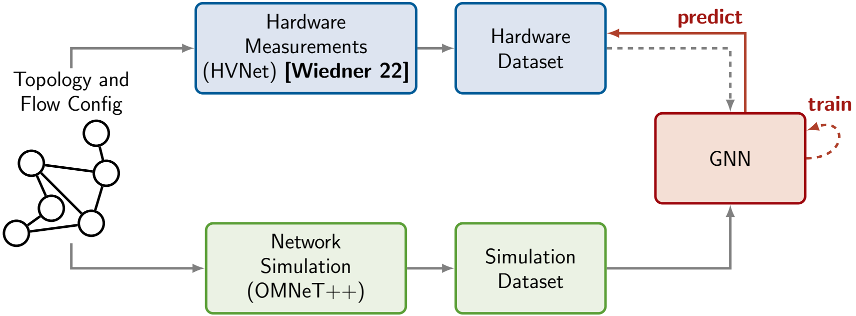

[Wiedner 22] Florian Wiedner et al. "HVNet: Hardware-Assisted Virtual Networking on a Single Physical Host". 2022.

Approach Overview

Outline

Related Work

Dataset Generation

- Hardware-Supported Measurements
- Simulations

Graph Neural Network

Results

Related Work

Work	Year	Data Source	Train Topologies	Unseen Test Topologies	Max. Network Size	Prediction Target
[Rusek 20]	2020	***Sim***	3	1	50	**Distrib.**
[Ferriol 22]	2022	***Sim, HW***	2 (*Sim*)	106 (*Sim*)	95 (*Sim*) / 8 (*HW*)	**Mean**
[Wang 22]	2022	***Sim***	2	1	24	**Mean**
[Yang 22]	2022	***Sim***	9	0?	$>$128	**Distrib.**
[Güemes 23]	2023	***HW***	$\leq$11	$<$11	8	**Mean**
[Helm 23]	2023	***HW***	87	10	15	**Distrib.**
This Work	2024	***Sim, HW***	71 (*Sim* + *HW*)	18 (*Sim* + *HW*)	15 (*Sim* + *HW*)	**Distrib.**

[Rusek 20] Krzysztof Rusek, José Suárez-Varela, Paul Almasan, Pere Barlet-Ros, and Albert Cabellos-Aparicio. "RouteNet: Leveraging Graph Neural Networks for Network Modeling and Optimization in SDN". 2020.

[Ferriol 22] Miquel Ferriol-Galmés, José Suárez-Varela, Jordi Paillissé, Xiang Shi, Shihan Xiao, Xiangle Cheng, Pere Barlet-Ros, Albert Cabellos-Aparicio. "Building a Digital Twin for Network Optimization Using Graph Neural Networks". 2022.

[Wang 22] Mowei Wang et al. "xNet: Improving Expressiveness and Granularity for Network Modeling with Graph Neural Networks". 2022.

[Yang 22] Qingqing Yang et al. "DeepQueueNet: Towards Scalable and Generalized Network Performance Estimation with Packet-Level Visibility". 2022.

[Güemes 23] Carlos Güemes-Palau, Miquel Ferriol Galmés, Albert Cabellos-Aparicio, Pere Barlet-Ros. "Building a Graph-Based Deep Learning Network Model From Captured Traffic Traces". 2023.

[Helm 23] Max Helm and Georg Carle. "Predicting Latency Quantiles Using Network Calculus-Assisted GNNs". 2023.

Hardware-Supported Measurements

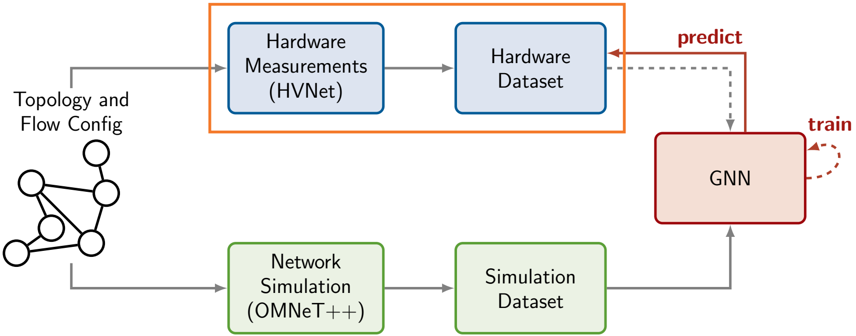

Setup by **[Gallenmüller 23]** and **[Wiedner 22]**

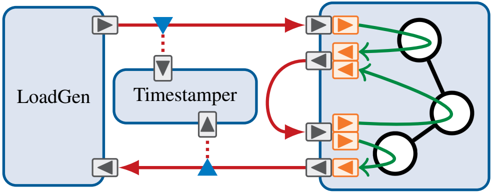

- Grey: NIC interface
- Red: Physical cable
- Green: Traffic between VMs

- Virtualized VM network on a single physical machine

[Gallenmüller 23] Sebastian Gallenmüller et al. "How Low Can You Go? A Limbo Dance for Low-Latency Network Functions". 2023.

[Wiedner 22] Florian Wiedner et al. "HVNet: Hardware-Assisted Virtual Networking on a Single Physical Host". 2022.

Hardware-Supported Measurements

Setup by **[Gallenmüller 23]** and **[Wiedner 22]**

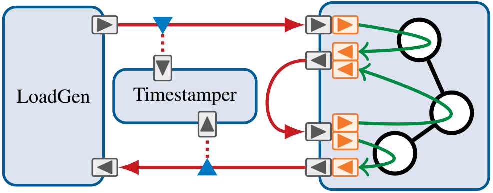

- Grey: NIC interface
- Red: Physical cable
- Green: Traffic between VMs

- Virtualized VM network on a single physical machine
- **High realism**: Traversing physical cable for every virtual link
- **High precision**: External timestamping using passive optical splitters
- Provided dataset
  - 100 random topologies
  - 14 billion latency values measured

[Gallenmüller 23] Sebastian Gallenmüller et al. "How Low Can You Go? A Limbo Dance for Low-Latency Network Functions". 2023.

[Wiedner 22] Florian Wiedner et al. "HVNet: Hardware-Assisted Virtual Networking on a Single Physical Host". 2022.

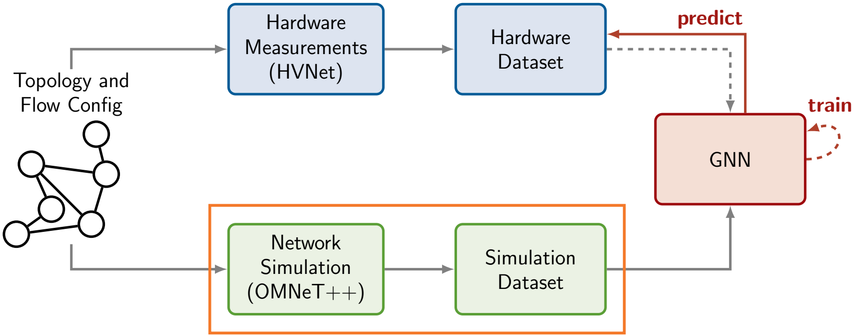
Hardware
Measurements
(HVNet)
Hardware
Dataset
predict
Topology and
Flow Config
train
GNN
Network
Simulation
(OMNeT++)
Simulation
Dataset

- Replication of the HVNet dataset in a simulation framework
  $\rightarrow$ Same 100 topologies, same flow properties
- Discrete Event Simulator **OMNeT++**[1]
- **INET** Framework[2] for simulating computer networks

---

[1] `https://omnetpp.org`

[2] `https://inet.omnetpp.org`

Simulations

- Replication of the HVNet dataset in a simulation framework
  $\rightarrow$ Same 100 topologies, same flow properties
- Discrete Event Simulator **OMNeT++**[1]
- **INET** Framework[2] for simulating computer networks
- Modelling of flow sending behavior
  - Sending distribution used for HW dataset unknown
  - Fit a *gamma* distribution over the inter-send times
  - Use the distribution in OMNeT++

---

[1] `https://omnetpp.org`

[2] `https://inet.omnetpp.org`

Dataset Properties

Properties

Metric	Min.	Max.
Network size	8	15
Number of flows	19	59
Flow length	2	9
Flow rate	1 Mbit/s	831 Mbit/s
Max. link util. per flow	0.11 %	87 %

Dataset Properties

Properties

Metric	Min.	Max.
Network size	8	15
Number of flows	19	59
Flow length	2	9
Flow rate	1 Mbit/s	831 Mbit/s
Max. link util. per flow	0.11 %	87 %

Correlation HW vs. Simulation Data

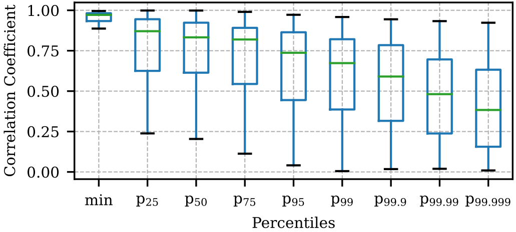

- **Pearson coefficient** between latency percentiles from HW and simulation dataset
  - 0 = no linear correlation
  - 1 = perfect correlation
- High correlation, especially for lower percentiles
- Lower correlation for high percentiles ($> p_{99.99}$) $\rightarrow$ Outliers in the HW dataset, not covered by the simulations

Graph Neural Network

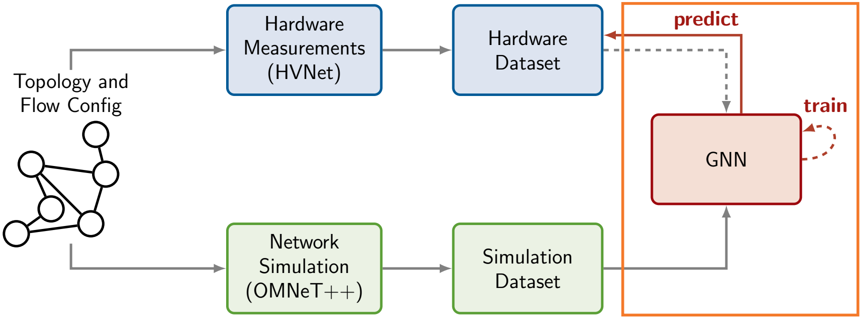

Graph Neural Network

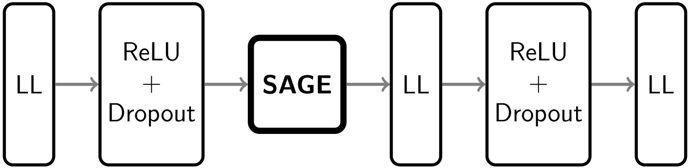

Building Blocks

- LL: Linear Layer
- ReLU + Dropout: Rectified Linear Unit and Dropout
- SAGE: GraphSAGE operator **[Hamilton 17]**

→ **Simple default architecture** (focus on effect of simulation data, not of the GNN model)

[Hamilton 17] Will Hamilton et al. "Inductive Representation Learning on Large Graphs". 2017.

Results

Predicting Latency Percentiles

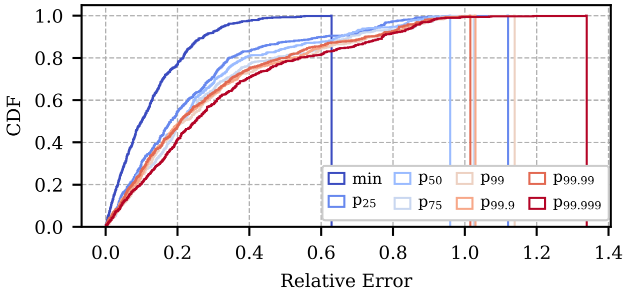

- MAPE over all prediction targets: 27.2 %
  → MdAPE: 19.8 %

Results

Predicting Latency Percentiles

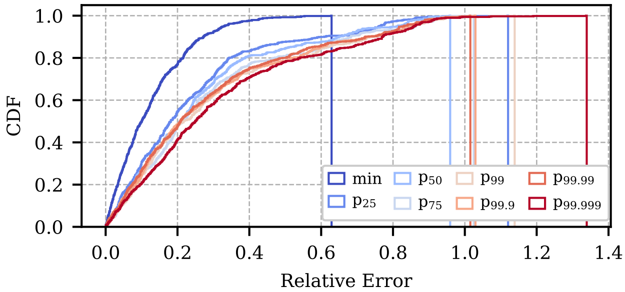

- MAPE over all prediction targets: 27.2 %
  → MdAPE: 19.8 %
- *Minimum* prediction is most accurate
- Higher percentiles are harder to predict
- "Plateau" between 40 % and 80 % rel. error

Results

Predicting Latency Percentiles

Topology with High Accuracy (13.6 % MAPE)

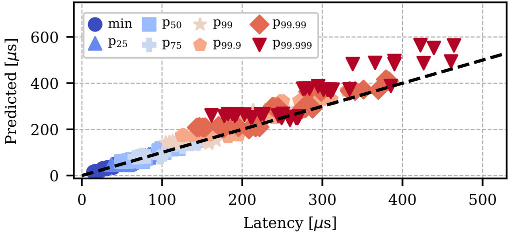

Predicting Latency Percentiles

Topology with High Accuracy (13.6 % MAPE)

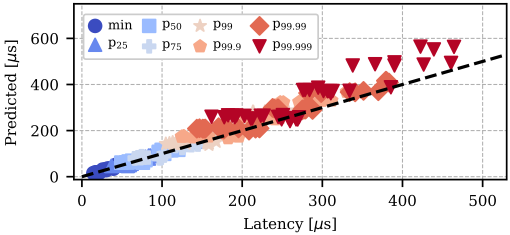

Topology with Low Accuracy (47.9 % MAPE)

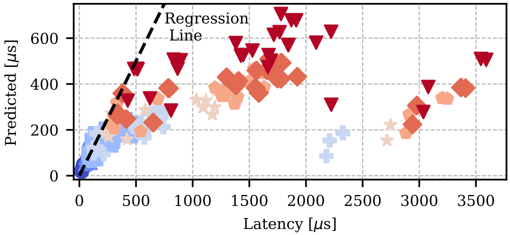

Problem: Outlier flows with 10× latency

→ Cannot be inferred from input features (incl. simulation)

Results

Feature Importance

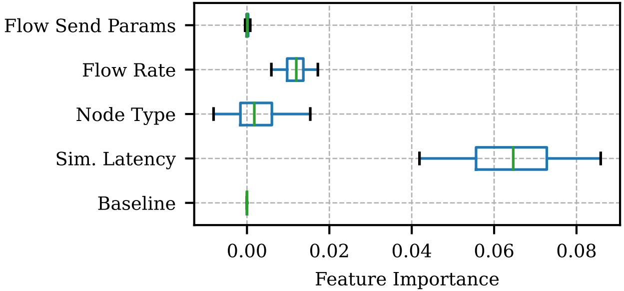

- Quantify the impact of different input features on prediction accuracy
- Simulation data has highest influence
  $\rightarrow$ Including it increases prediction accuracy

Conclusion

Approach and Dataset

- Using simulation data as additional input to GNNs
- Dataset based on HW measurements and DES
- $\sim$ 100 different network topologies
- **Limitations**: small dataset, only UDP traffic, no flow dynamics, simple GNN architecture

Results

- Simulation data improves prediction accuracy
- Minimum and lower percentiles easier to predict
- **Limitations**: outlier flows with unexpectedly high latency

Paper

Data and Code

Questions? – spaethj@net.in.tum.de