SmartNIC Architectures and their Roles in High-Speed Networking

Andreas Herkersdorf

School of Computation, Information and Technology

Technical University of Munich

MPSoC'25

Megève, June 19, 2025

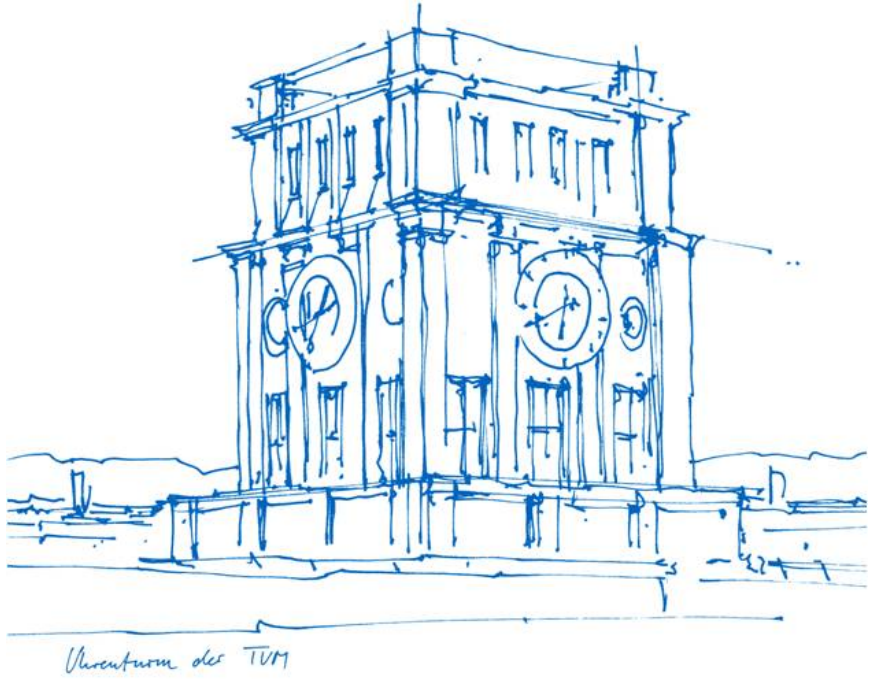

Reference back to MPSoC'24 in Kanazawa

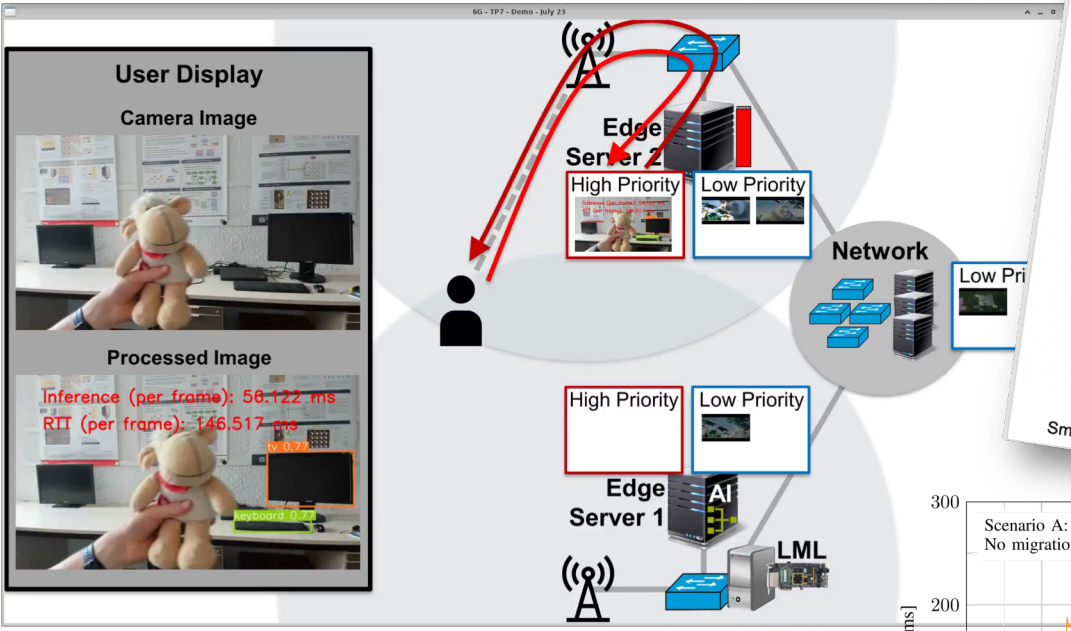

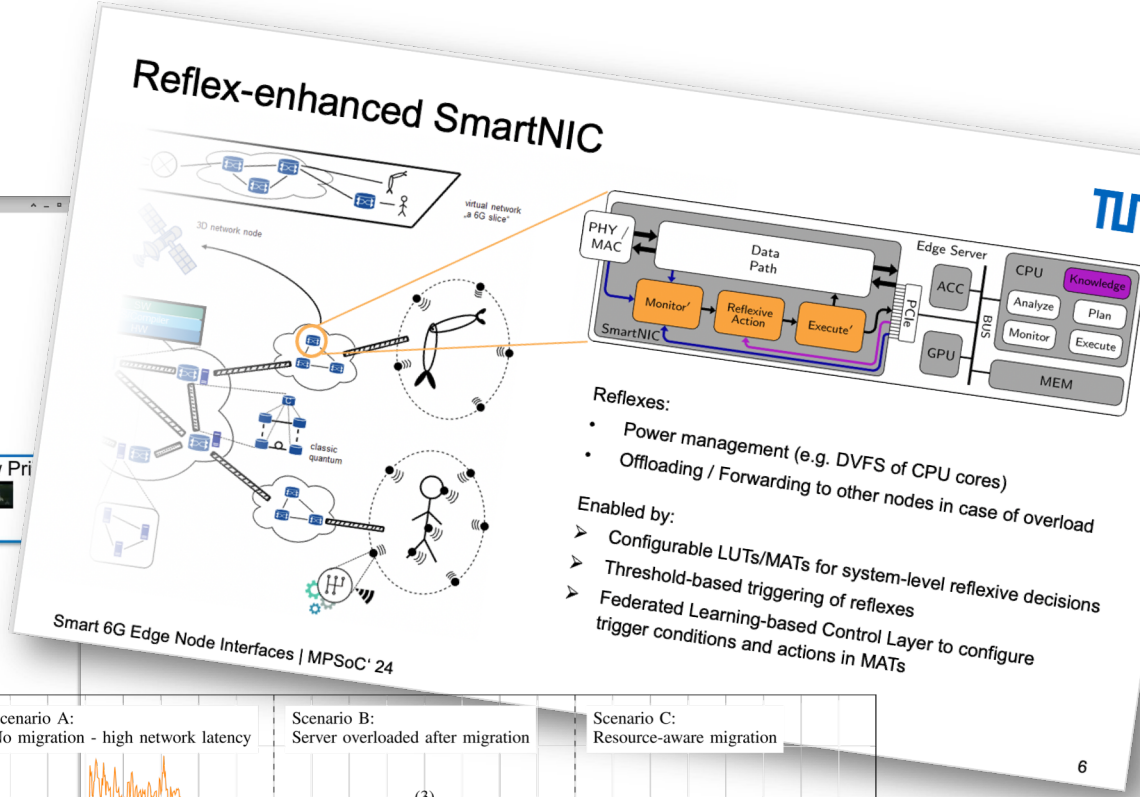

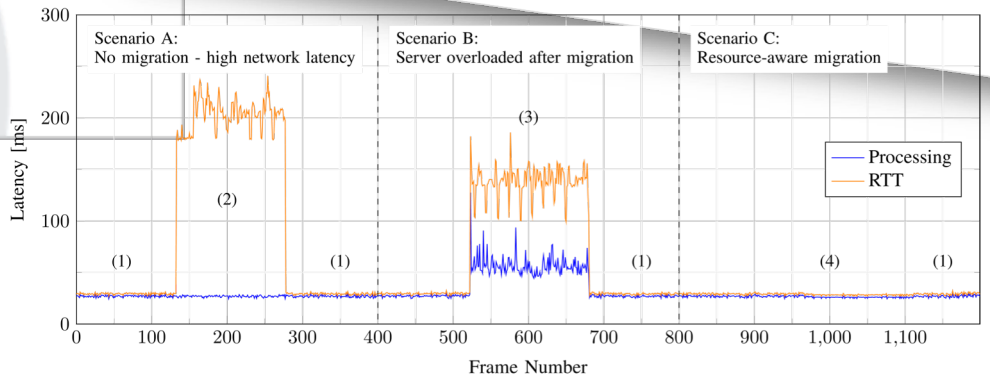

A Well-Known Story: Big Data in Motion

Feb 27, 2010

2019

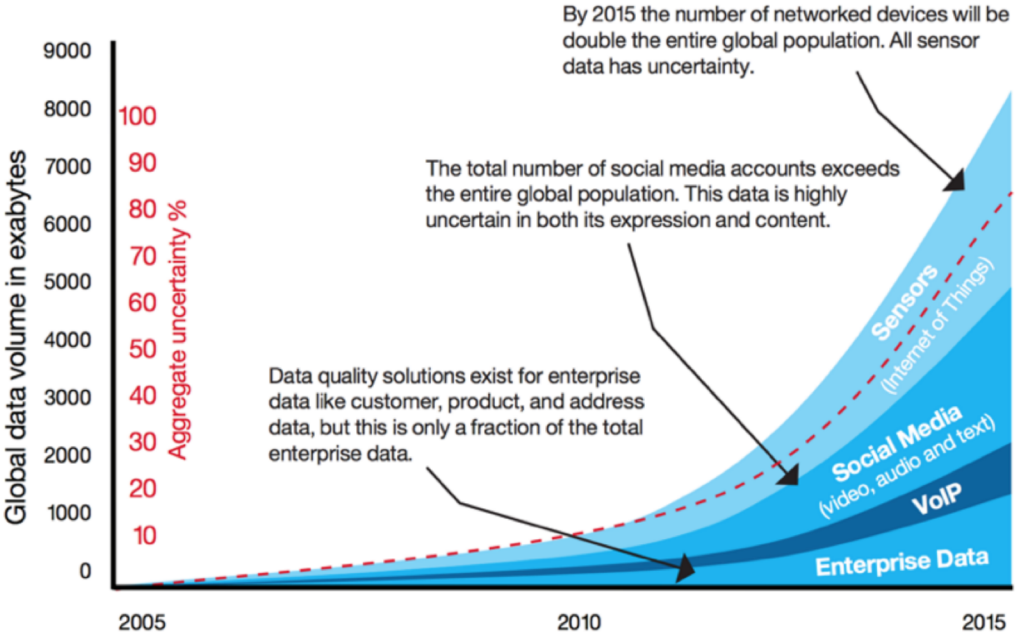

„The Promise and Prejudice of Big Data in Intelligence Community“, K. Jani, Georgia Institute of Technology, October 26, 2016

„Hey, ChatGPT …

- … how to visualize the mutual dependencies between microelectronic advancements and staggering networking demands?”
  - “… in 3D please!”

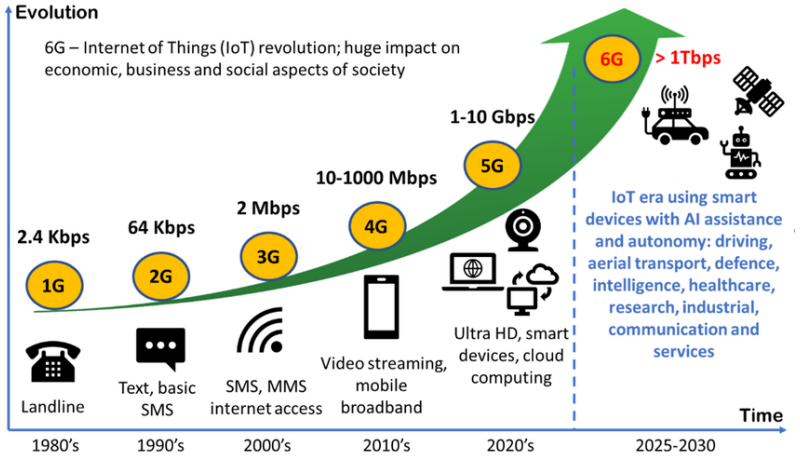

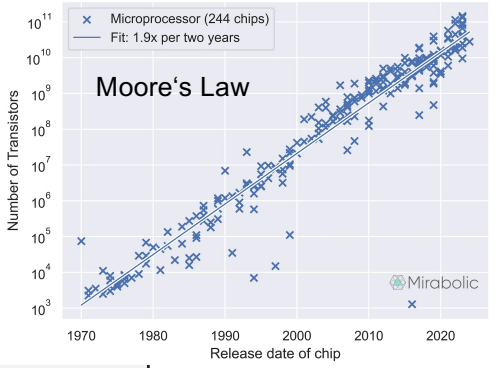

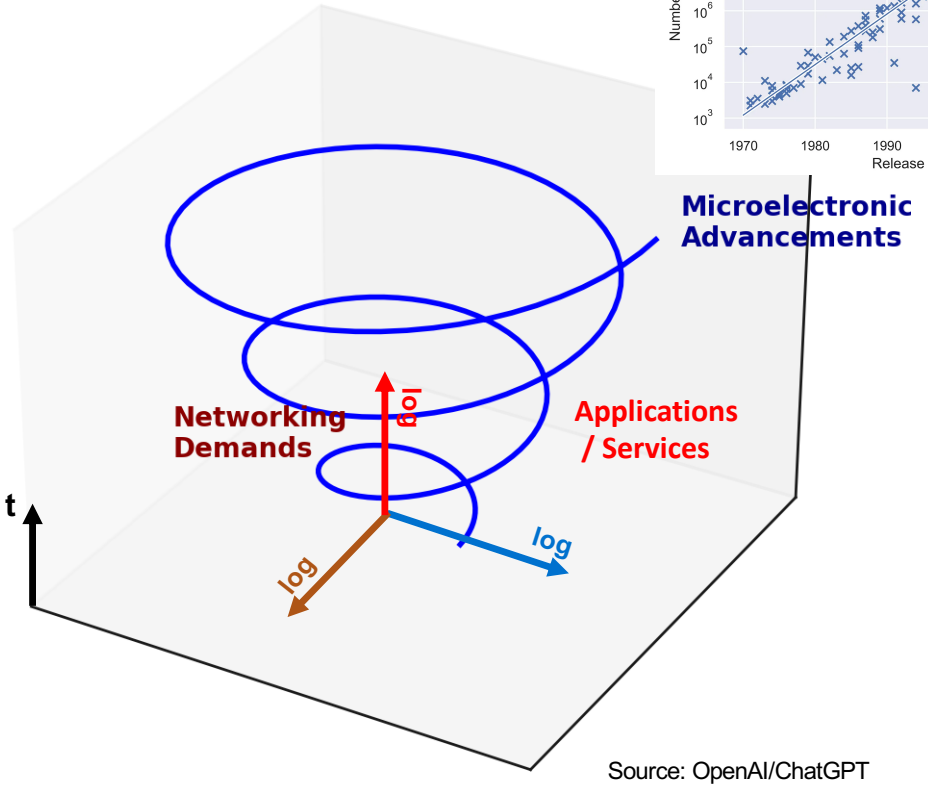

Source: OpenAI/ChatGPT

Let's quantify some challenges for high-speed NICs

The required processing capacity of NIC depends on:

Packet rate / Packet interarrival time

→ *packets per second (PPS)* $PPS = \frac{link\_speed}{packet\_size}$

Link speed	*40 Gbps*		*100 Gbps*		*800 Gbps*	
Size [Byte]	*64*	*512*	*64*	*512*	*64*	*512*
PPS [M]	*78*	*9.8*	*195*	*24*	*1,563*	*195*
1 / PPS [ns]	*12.8*	*102*	*5.1*	*41*	*0.64*	*5.1*

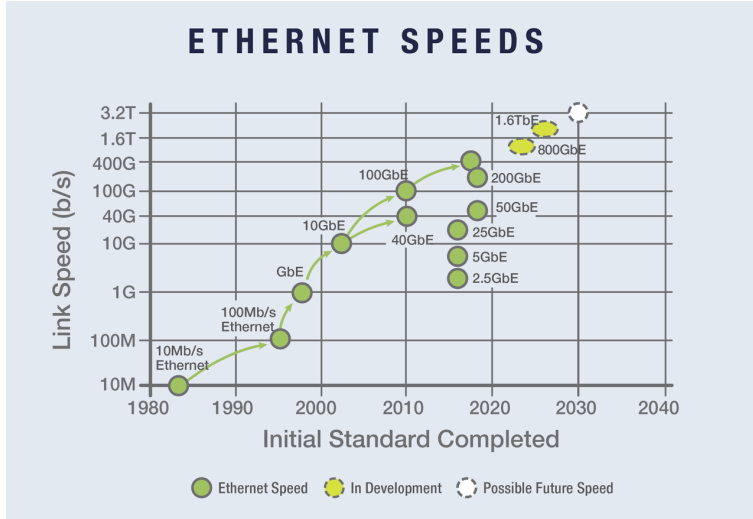

Source: https://iebmedia.com/wp-content/uploads/2023/05/Ethernet-Speeds.png

LLet's quantify some challenges for high-speed NICs

Header Processing

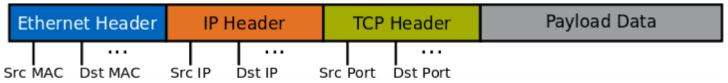

- Fixed header location allows simple parsing
- **Common use-cases:** routing, switching, IP/port-based firewalls, …

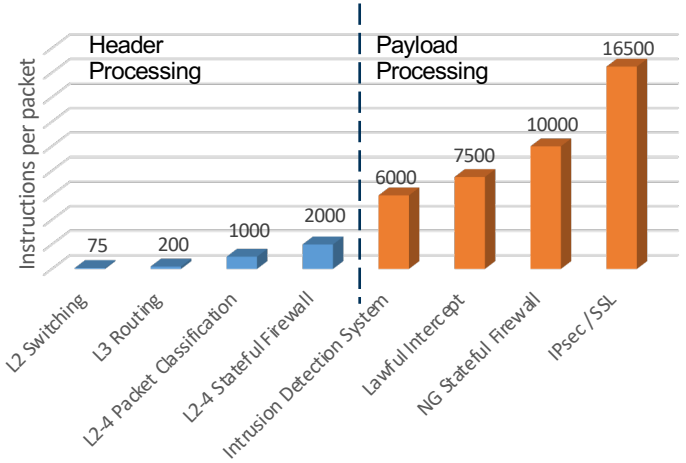

Source: Netronome

Payload Processing

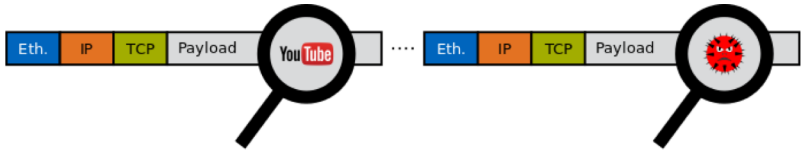

- **Common use-cases:** application/session/user identification for firewalls or bandwidth throttling, cryptology, intrusion detection, virus scanning, …

Let's quantify some challenges for high-speed NICs

The processing complexity of a networking function:

→ *instructions per packet (IPP)*

Required processing capacity = $PPS * IPP \left[\frac{instructions}{second}\right]$

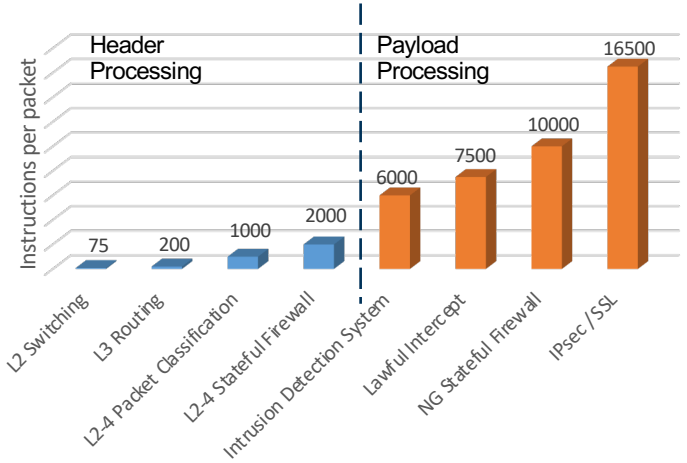

Source: Netronome

@ 100Gbps

L2 Switching (64B): $PPS * IPP = 195.313.000 \frac{packets}{second} * 75 \frac{instructions}{packet} = 14.6 * 10^9 \frac{instructions}{second}$

Intrusion Detection (512B): $146 * 10^9 \frac{instructions}{second}$ IPSec (512B) : $403 * 10^9 \frac{instructions}{second}$

Let's quantify some challenges for high-speed NICs

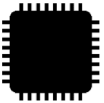

AMD Ryzen 5 5600: 6 cores @ 3.9 GHz, 65 W TDP

$$IPS = 6\ cores * f_{clk} * IPC_{core} \approx 23{,}4 * 10^9 \frac{instructions}{second}$$

$$IPC_{core} \approx 1$$

Networking Function	100Gbps
L2 Switching	$14.6 * 10^9$
Intrusion Detection	$146 * 10^9$
IPSec	$403 * 10^9$

L2 Switching ✓ Intrusion Detection ✗ IPSec ✗

IPSec @ 100Gbps would require 18 CPUs with a TDP of 1.1 kW!

Let's quantify some challenges for high-speed NICs

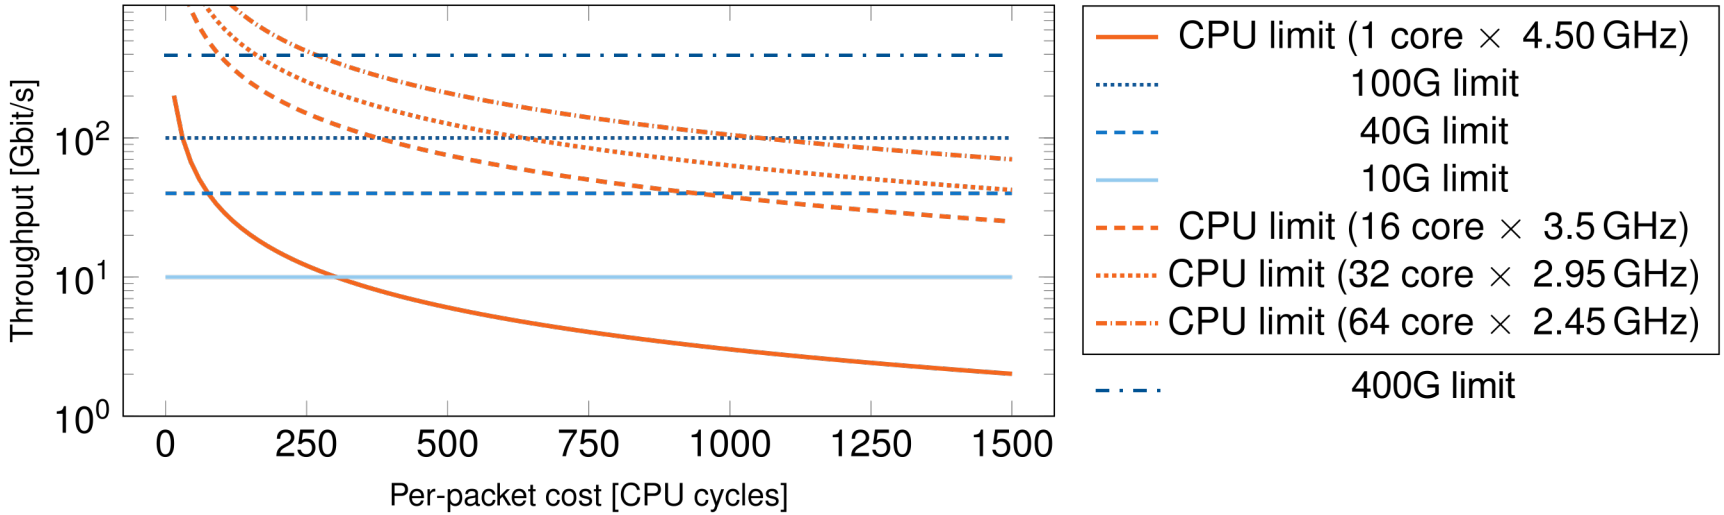

- Multi- and manycore architectures help to achieve higher throughputs
- … but complexity of network services grows faster than processor performance

Ref.: S. Gallenmüller, P. Emmerich, F. Wohlfart, D. Raumer, and G. Carle. „Comparison of Frameworks for High-Performance Packet IO". In: ANCS '15

The „Bigger Context“ of SmartNICs

With a Legacy NIC, this „fraction“ of the SW stacks „eats“ the Server-CPU compute performance!

Offload networking functions to SmartNIC resources for ever increasing networking demands

Typical functions of commercial SmartNIC:

- vSwitch offloading, NFV
- IPSec, SSL, Paket Filtering (DPI)
- Storage Protocol offload

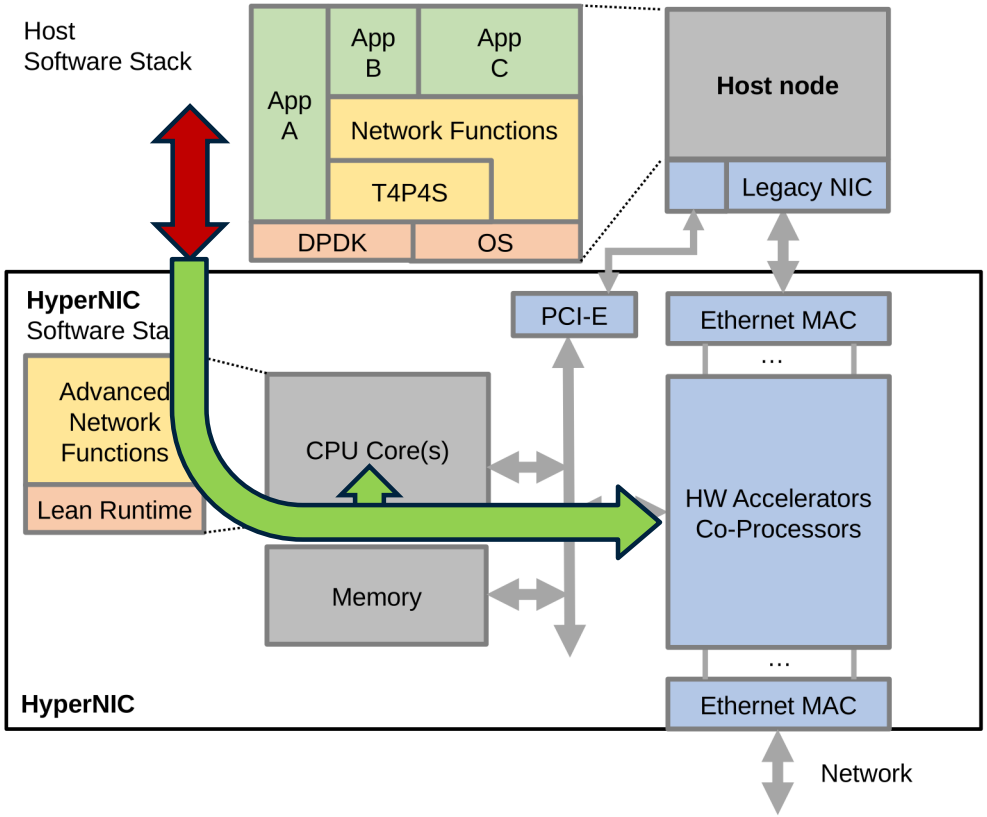

The „Bigger Context“ of SmartNICs

Performance and Energy Efficiency

- cope with very high data rates (up to hundreds of Gbps)
- lowest packet delay as possible
- low power consumption

➡ Customized **ASIC**

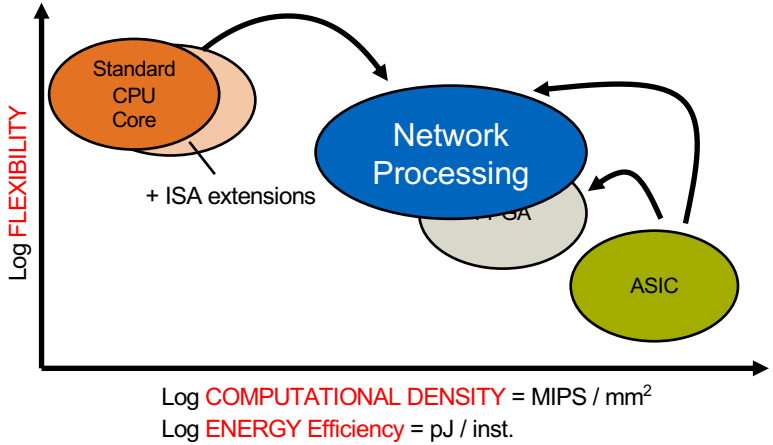

Source: T. Noll / H. Blume et al., „Model-based exploration of the Design Space for Heterogeneous System on Chip“, 2002

Flexibility

- adapt to evolving packet processing applications
- efficient resource sharing among network applications

➡ Programmable **CPU / ASIP**

State-of-the-art Network Processors

Netronome NFP-6xxx Flow Processor

- 216 programmable cores to execute software
  - 96 *packet processing cores* for stateless processing
  - 120 *flow processing cores* for stateful processing
- More than $300 * 10^9 \frac{\text{instructions}}{\text{second}}$
- 100 hardware accelerators for
  - DPI, regular expression matching
  - Cryptography
  - Hash calculation
  - Packet I/O, Queue Management
  - …
- 50 Gbps bulk cryptography
- 720 Gbps I/O

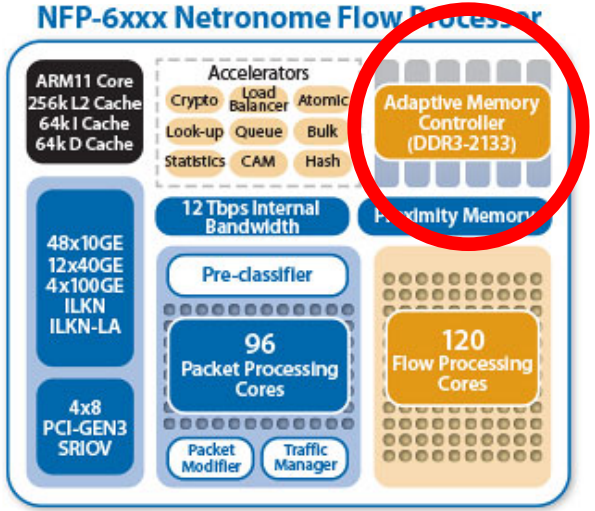

Source: Netronome

Network Processing Memory Bandwidth Requirements

IP packets (when processed “as usual”) …

- … traverse the memory interface **at least** 4 times!

- **Exceeds the peak streaming data rate of DDRx!**

$$\boldsymbol{BW_{store}} \geq 4 \cdot 100Gbps = 400\ Gbps = \boldsymbol{50GByte/s}$$

Memory	T_{acc} [ns]	$f_{I/O}$ [MHz]	Data width	BW_{max} [GByte/s]
DDR3-2133	21 - 26	1066	64 bit	17.0
DDR4-3200	25 - 30	1600	64 bit	25.6
DDR5-5600	(~25)	2800	64 bit	44.8
HBM 3		3200	1024	820

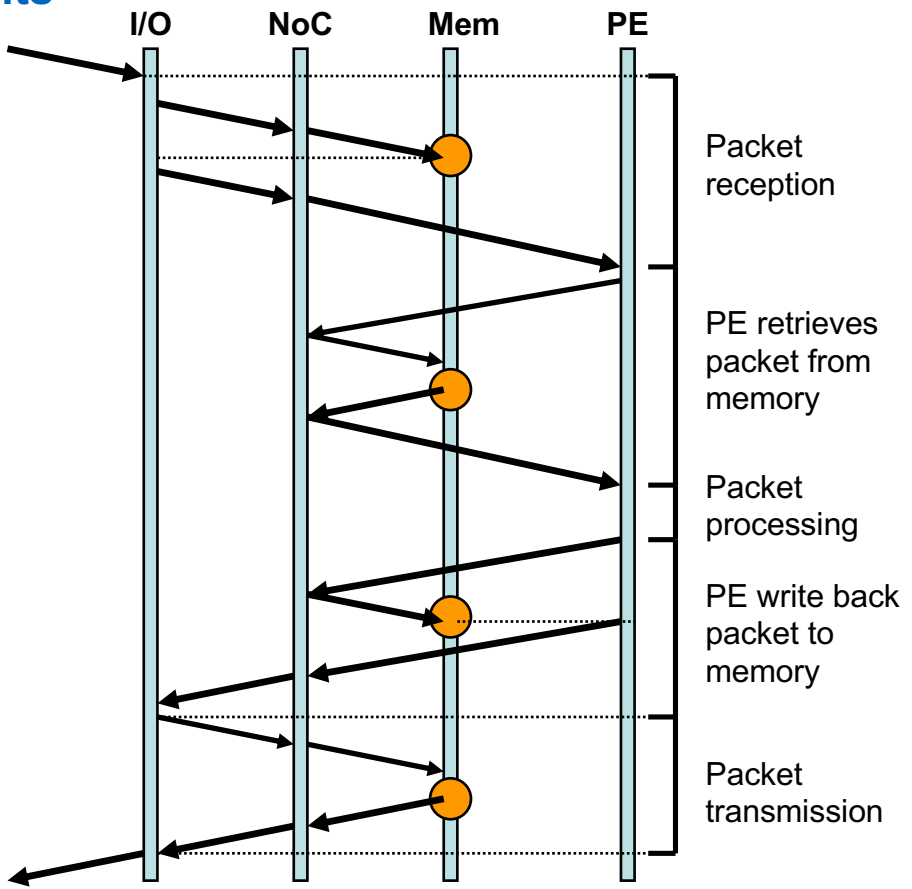

SASSMC Memory Controller Extension

- Wrapper extension to standard Memory Controller
  - Reduce average data access latency to main memory and compulsory LLC cache miss accesses to main memory
  - “On the fly” encryption / decryption of data in main memory
- Strategies for pre-fetch / write back
  - Application profiling
  - OS / Hypervisor “hints”
  - ML

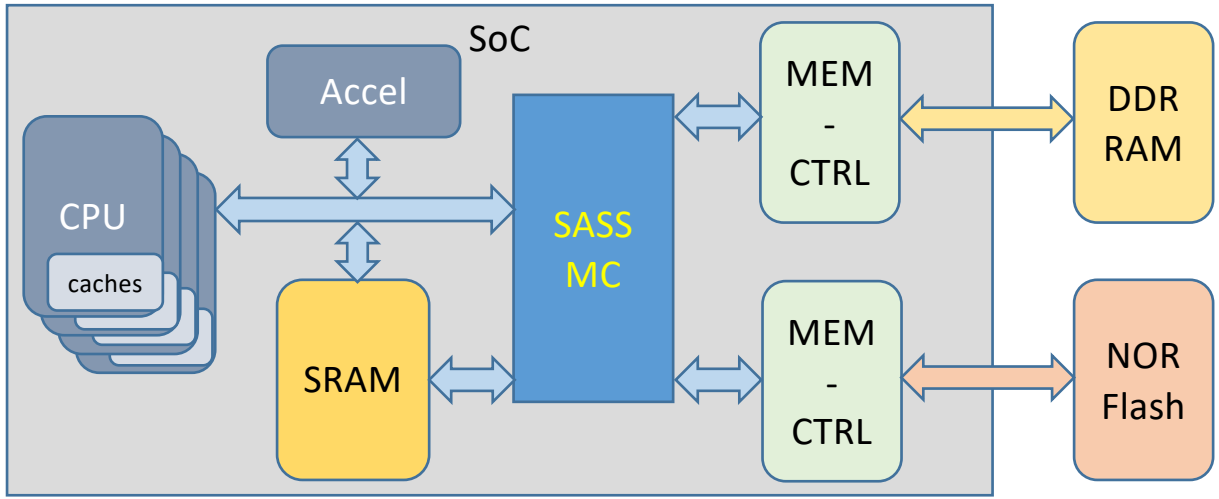

Yet another Challenge, … the BIG ONE

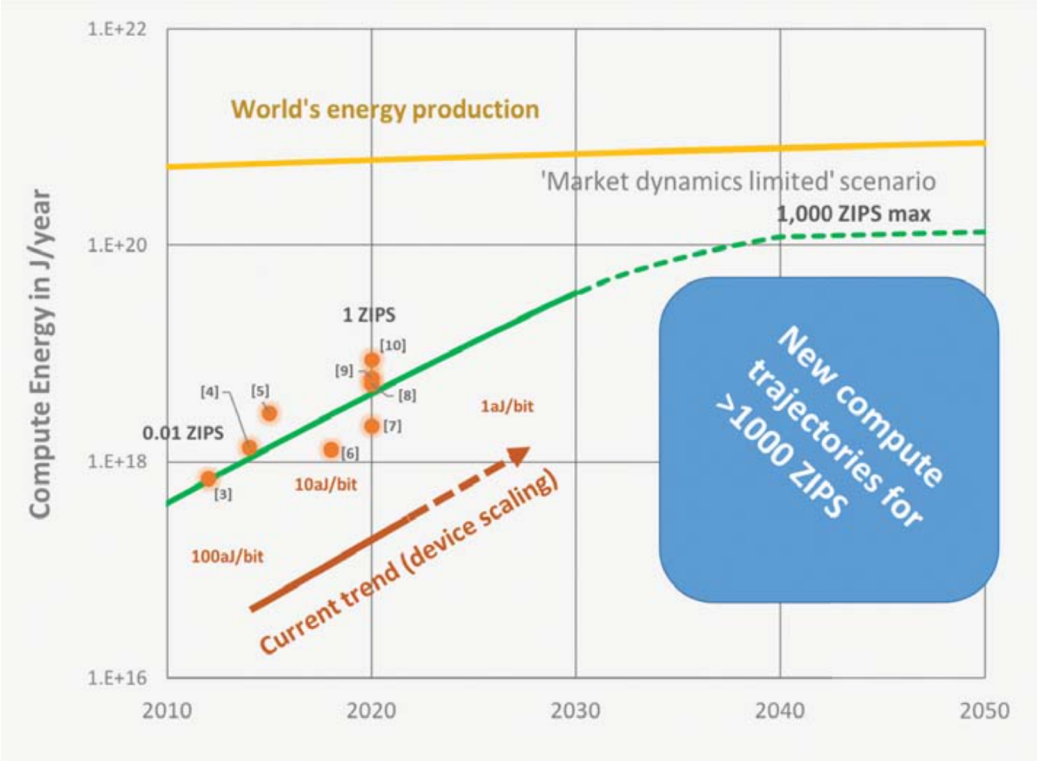

SRC, Decadal Plan for Semiconductors, Full Report, January **2021**:
https://www.src.org/about/decadal-plan/

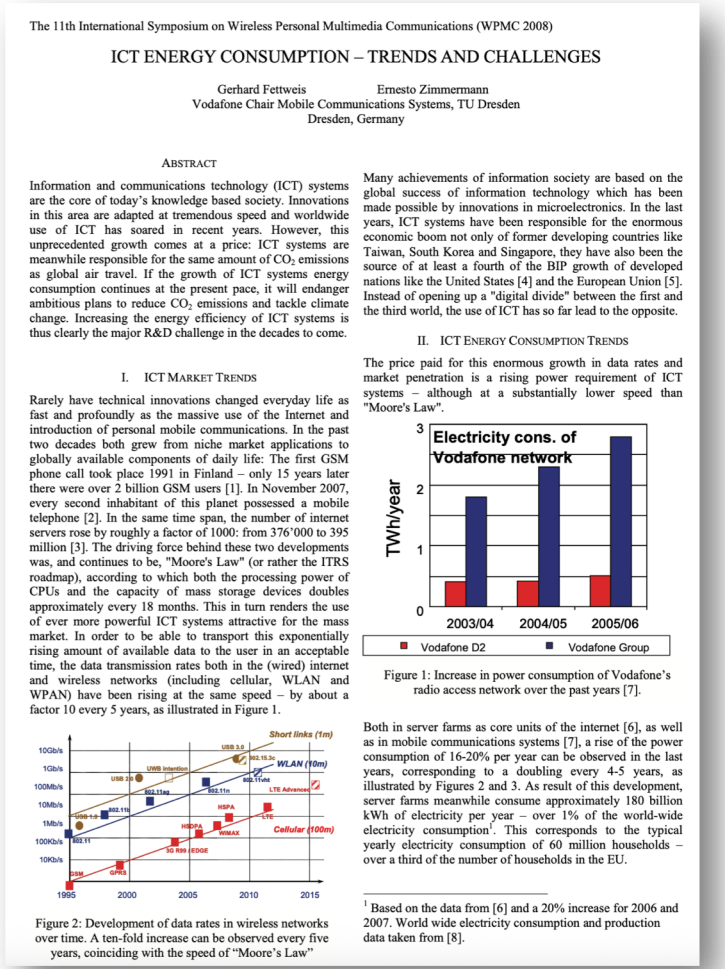

The 11th International Symposium on Wireless Personal Multimedia Communications (WPMC 2008)

ICT ENERGY CONSUMPTION – TRENDS AND CHALLENGES

Gerhard Fettweis Ernesto Zimmermann
Vodafone Chair Mobile Communications Systems, TU Dresden
Dresden, Germany

ABSTRACT

Information and communications technology (ICT) systems are the core of today's knowledge based society. Innovations in this area are adapted at tremendous speed and worldwide use of ICT has soared in recent years. However, this unprecedented growth comes at a price: ICT systems are meanwhile responsible for the same amount of CO_2 emissions as global air travel. If the growth of ICT systems energy consumption continues at the present pace, it will endanger ambitious plans to reduce CO_2 emissions and tackle climate change. Increasing the energy efficiency of ICT systems is thus clearly the major R&D challenge in the decades to come.

I. ICT MARKET TRENDS

Rarely have technical innovations changed everyday life as fast and profoundly as the massive use of the Internet and introduction of personal mobile communications. In the past two decades both grew from niche market applications to globally available components of daily life: The first GSM phone call took place 1991 in Finland – only 15 years later there were over 2 billion GSM users [1]. In November 2007, every second inhabitant of this planet possessed a mobile telephone [2]. In the same time span, the number of internet servers rose by roughly a factor of 1000: from 376'000 to 395 million [3]. The driving force behind these two developments was, and continues to be, "Moore's Law" (or rather the ITRS roadmap), according to which both the processing power of CPUs and the capacity of mass storage devices doubles approximately every 18 months. This in turn renders the use of ever more powerful ICT systems attractive for the mass market. In order to be able to transport this exponentially rising amount of available data to the user in an acceptable time, the data transmission rates both in the (wired) internet and wireless networks (including cellular, WLAN and WPAN) have been rising at the same speed – by about a factor 10 every 5 years, as illustrated in Figure 1.

Figure 2: Development of data rates in wireless networks over time. A ten-fold increase can be observed every five years, coinciding with the speed of "Moore's Law"

Many achievements of information society are based on the global success of information technology which has been made possible by innovations in microelectronics. In the last years, ICT systems have been responsible for the enormous economic boom not only of former developing countries like Taiwan, South Korea and Singapore, they have also been the source of at least a fourth of the BIP growth of developed nations like the United States [4] and the European Union [5]. Instead of opening up a "digital divide" between the first and the third world, the use of ICT has so far lead to the opposite.

II. ICT ENERGY CONSUMPTION TRENDS

The price paid for this enormous growth in data rates and market penetration is a rising power requirement of ICT systems – although at a substantially lower speed than "Moore's Law".

Figure 1: Increase in power consumption of Vodafone's radio access network over the past years [7].

Both in server farms as core units of the internet [6], as well as in mobile communications systems [7], a rise of the power consumption of 16-20% per year can be observed in the last years, corresponding to a doubling every 4-5 years, as illustrated by Figures 2 and 3. As result of this development, server farms meanwhile consume approximately 180 billion kWh of electricity per year – over 1% of the world-wide electricity consumption[1]. This corresponds to the typical yearly electricity consumption of 60 million households – over a third of the number of households in the EU.

[1] Based on the data from [6] and a 20% increase for 2006 and 2007. World wide electricity consumption and production data taken from [8].

G. Fettweis, E. Zimmermann, „ICT Energy Consumption – Trends and Challenges", WPMC **2008**.

Multi-Purpose SmartNICs

Offload compute node resources for ever increasing networking demands:

- Network and Node Resilience
  - Low-latency network coding / FEC
  - Reflex-based traffic steering
- Energy Efficiency & Power Management
  - ecoNIC-based workload pinning

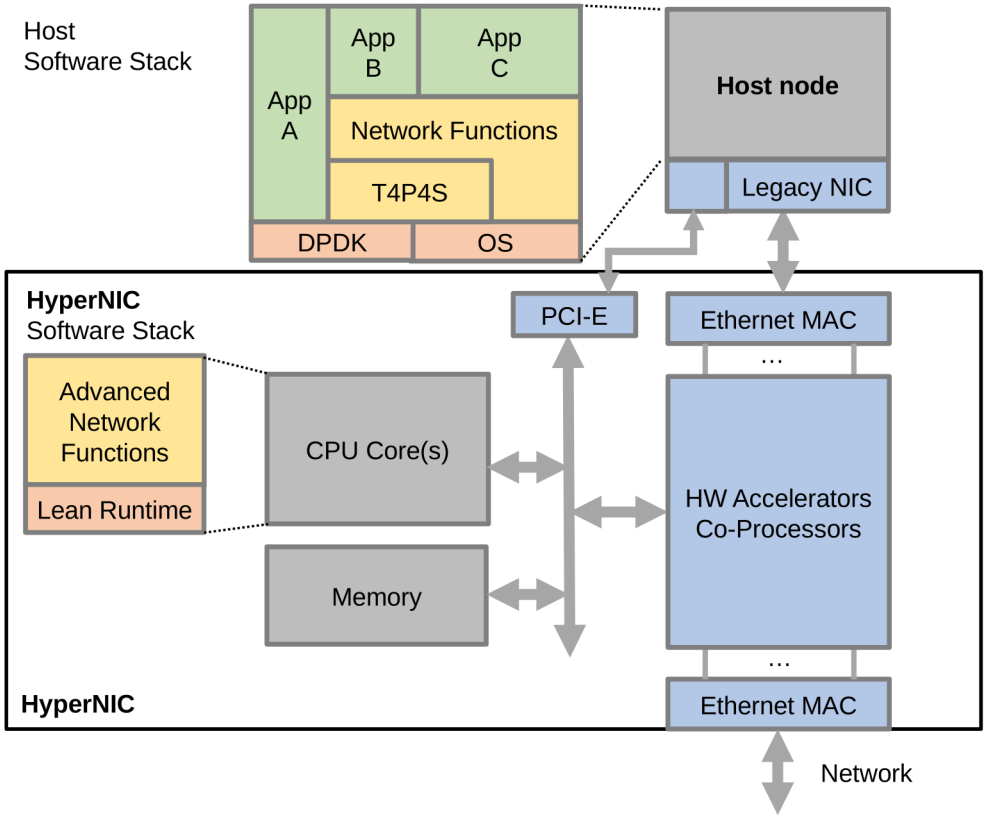

ecoNIC

- Combines traffic-dependent power management in OS / Linux with priority-aware traffic steering / pinning in SmartNIC HW

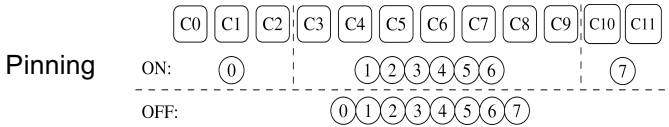

F. Biersack, M. Liess, et. al., “ecoNIC: Saving Energy through SmartNIC-based Load Balancing of Mixed-Critical Ethernet Traffic“, 27th Euromicro DSD, 2024.

Networking Testbed @ LIS

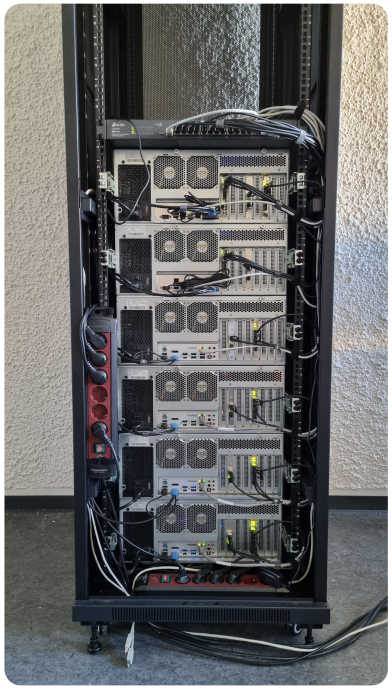

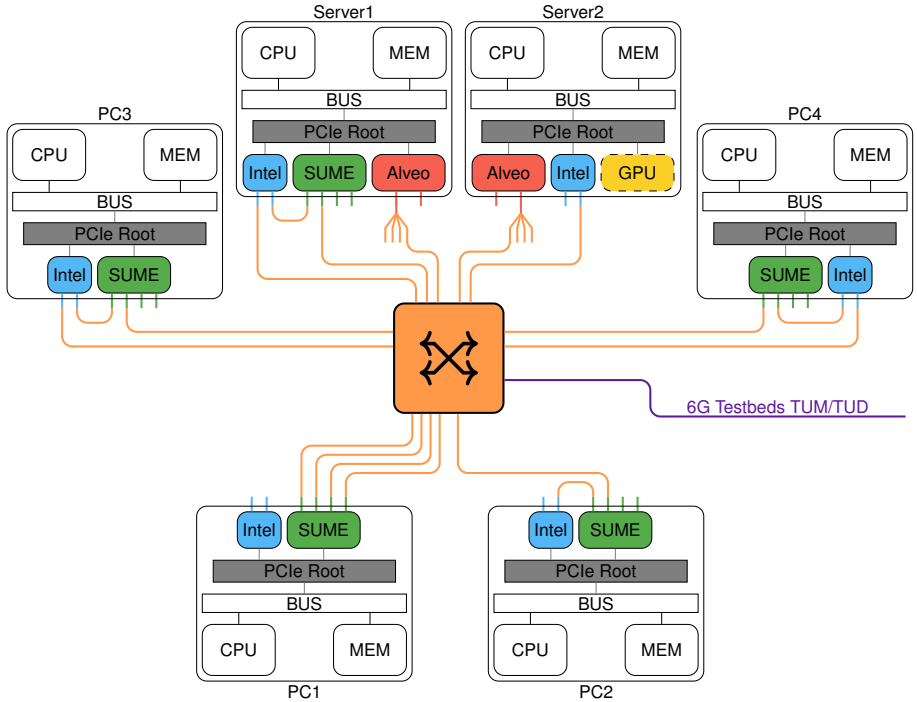

ecoNIC

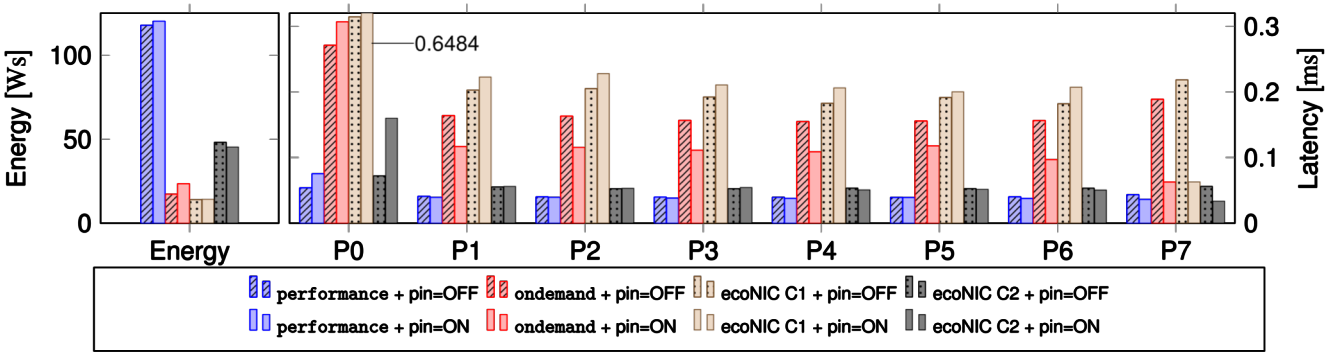

- Comparison relative to Linux Power Governors (performance, ondemand)
- C1/C2: different parameter settings for switching between power states

Take Aways …

- High-speed / high-data-rate networks pose major technical challenges on the data plane compute architecture
  - Not only on provisioning sufficient compute performance / accelerators, …
  - equally on data movement and storage
  - energy consumption

- Crucial relevance of ingress / egress wire-rate pre-/post-processing in NICs
  - Offloading heterogeneous host processing (function repartitioning)
  - Smart traffic steering and monitoring
  - Energy saving and low-latency priority services are not necessarily contradicting goals

Acknowledgements

First and foremost, many thanks to my team. For this work in particular: Franz Biersack, Shichen Huang, Oli Lenke, Marco Liess, Yuanji Ye and Thomas Wild.

We are greatful for the financial support by the *Bavarian Ministry of Economic Affairs, Regional Development and Energy* in the *6G Future Lab Bavaria* project and the *Federal Ministry of Research, Technology and Space* in the *6G-life and CeCaS* projects with project identification numbers 16KISK002 and 16ME0800K.