Chiplet Technology and the Impact of NUMA on Applications

Alex Daichendt

IN2076 Advanced Computer Architecture
School of Computation, Information, and Technology
Technical University of Munich

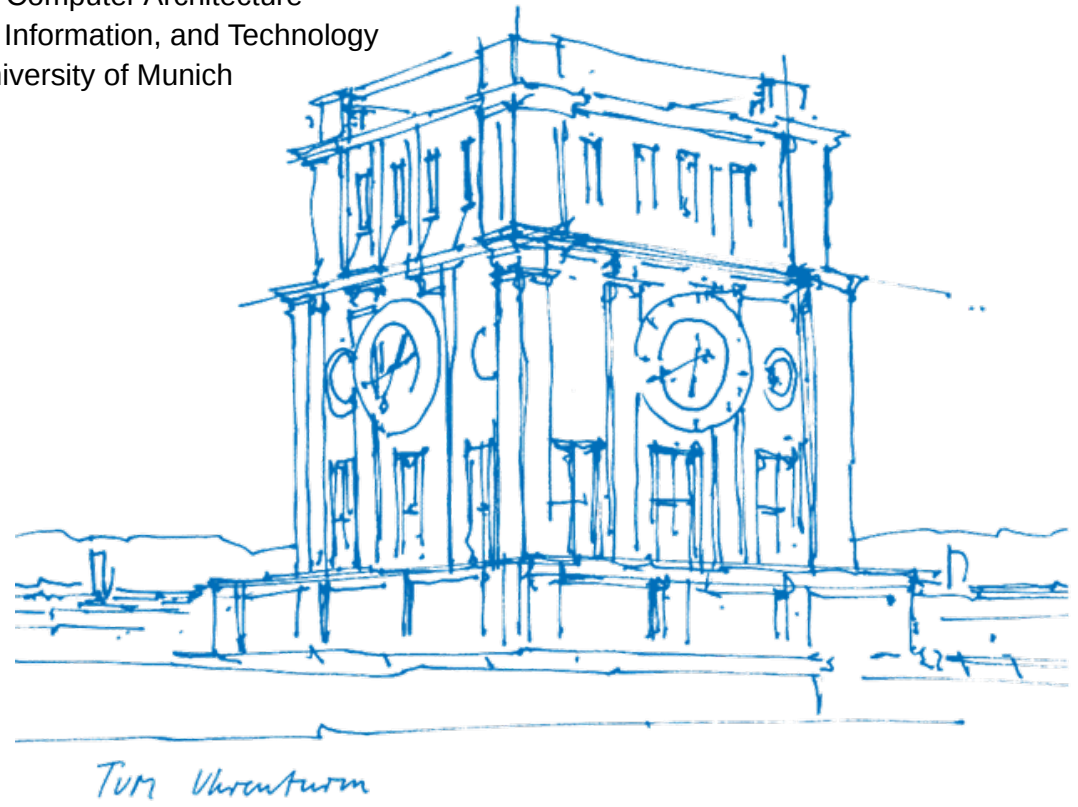

Intel Press Workshops June 2017[1]

Intel® Data Center Advantages vs. AMD Naples

Implementation Matters

INTEL

HOLISTIC DATA CENTER SOLUTIONS

- ✓ Proven System Performance and Innovation
- ✓ Architected for the Data Center
- ✓ Robust SW and HW Ecosystem

AMD

REPURPOSED DESKTOP PRODUCT FOR SERVER

- ✗ Inconsistent Performance from 4 Glued-together Desktop Die
- ✗ Poor Track Record, Inconsistent Supplier
- ✗ Lack of Ecosystem

[1] TechPowerUp; no primary source available

Why chiplets?

- Moore’s Law
- more flexibility in design

[2] [1] S. Naffziger *et al.*, “Pioneering Chiplet Technology and Design for the AMD EPYC™ and Ryzen™ Processor Families : Industrial Product”, IEEE, Jun. 2021. doi: 10.1109/isca52012.2021.00014.

Why chiplets?

- Moore’s Law
- more flexibility in design
- low production yield for monolithic dies
  → $$$

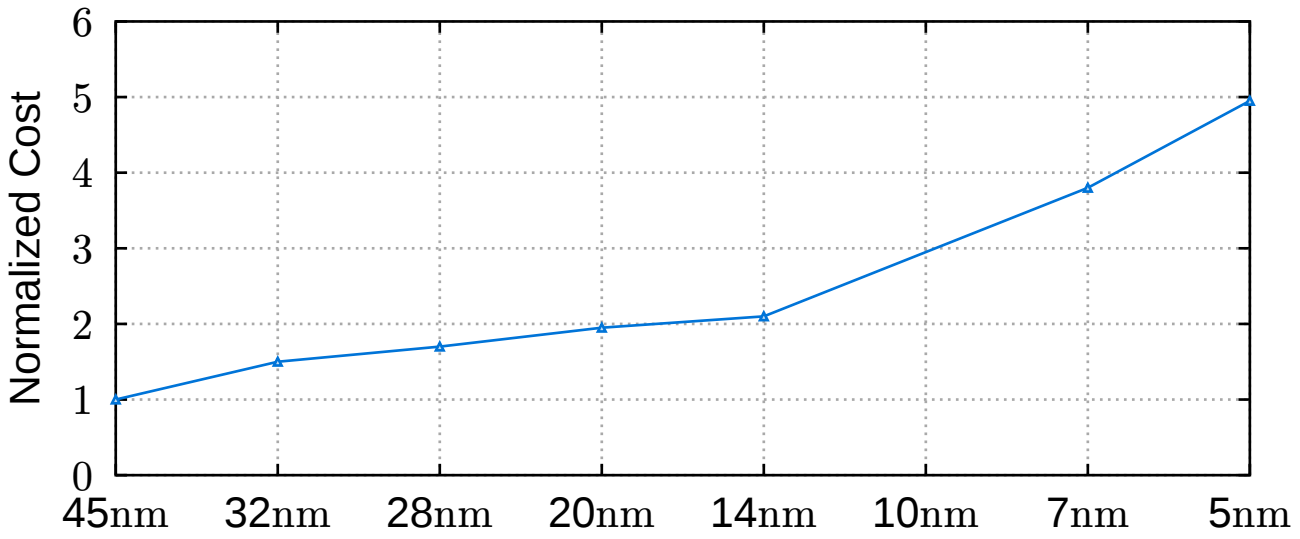

Figure 2: Normalized cost per chip vs. technology node, based on Naffziger et al.[3]

[3] [1] S. Naffziger *et al.*, “Pioneering Chiplet Technology and Design for the AMD EPYC™ and Ryzen™ Processor Families : Industrial Product”, IEEE, Jun. 2021. doi: 10.1109/isca52012.2021.00014.

AMD Naples (1st Gen. EPYC) – NUMA Toplogy[4]

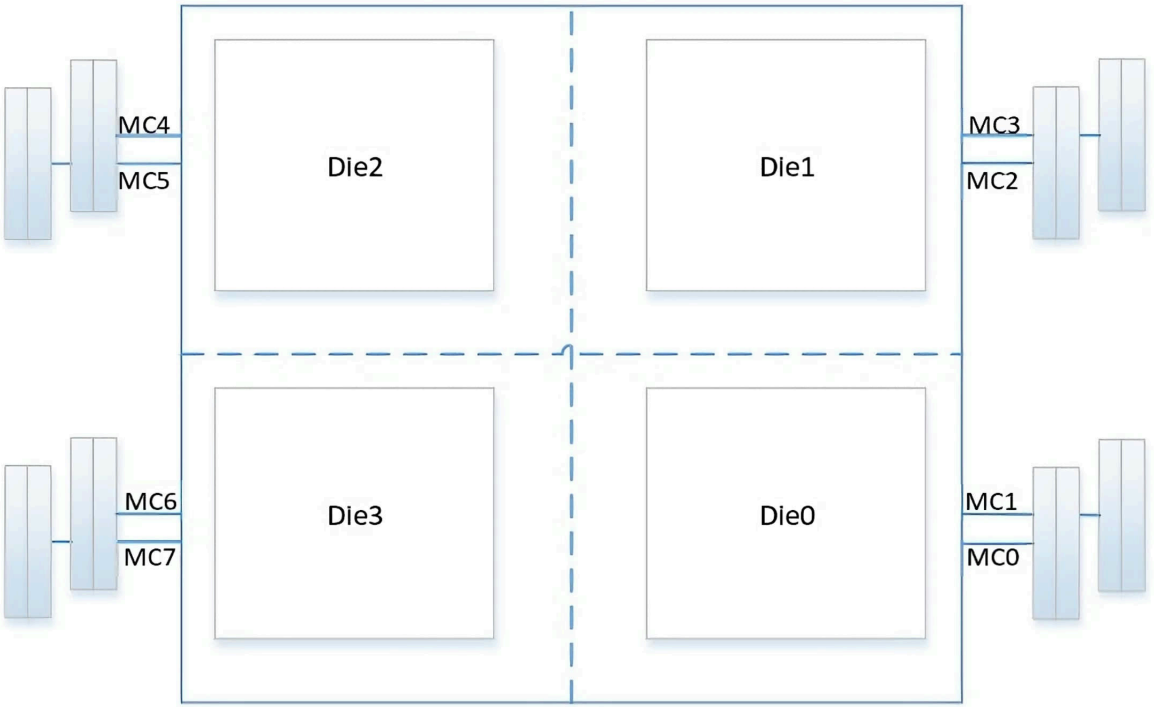

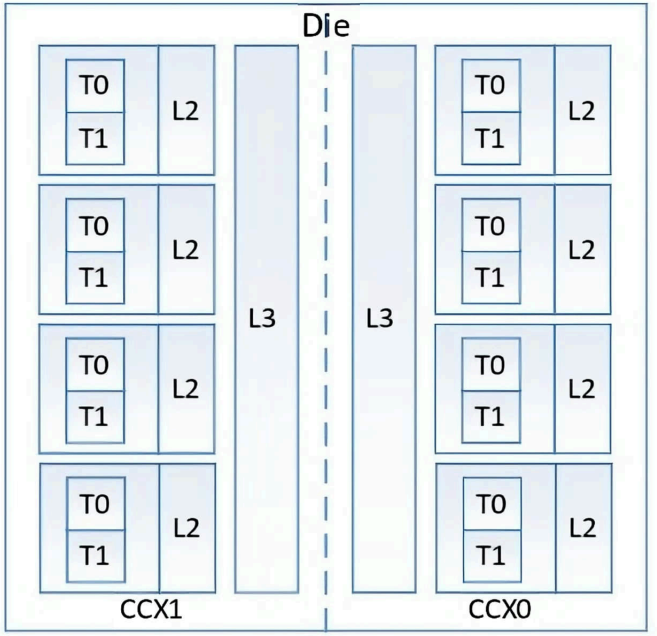

Legend

Tx = thread Lx = L2/3 cache MCx = Memory Controller CCXx = Core Complex

[4] https://www.amd.com/content/dam/amd/en/documents/epyc-technical-docs/specifications/56308-numa-topology-for-epyc-naples-family-processors.pdf

AMD Naples (1st Gen. EPYC)

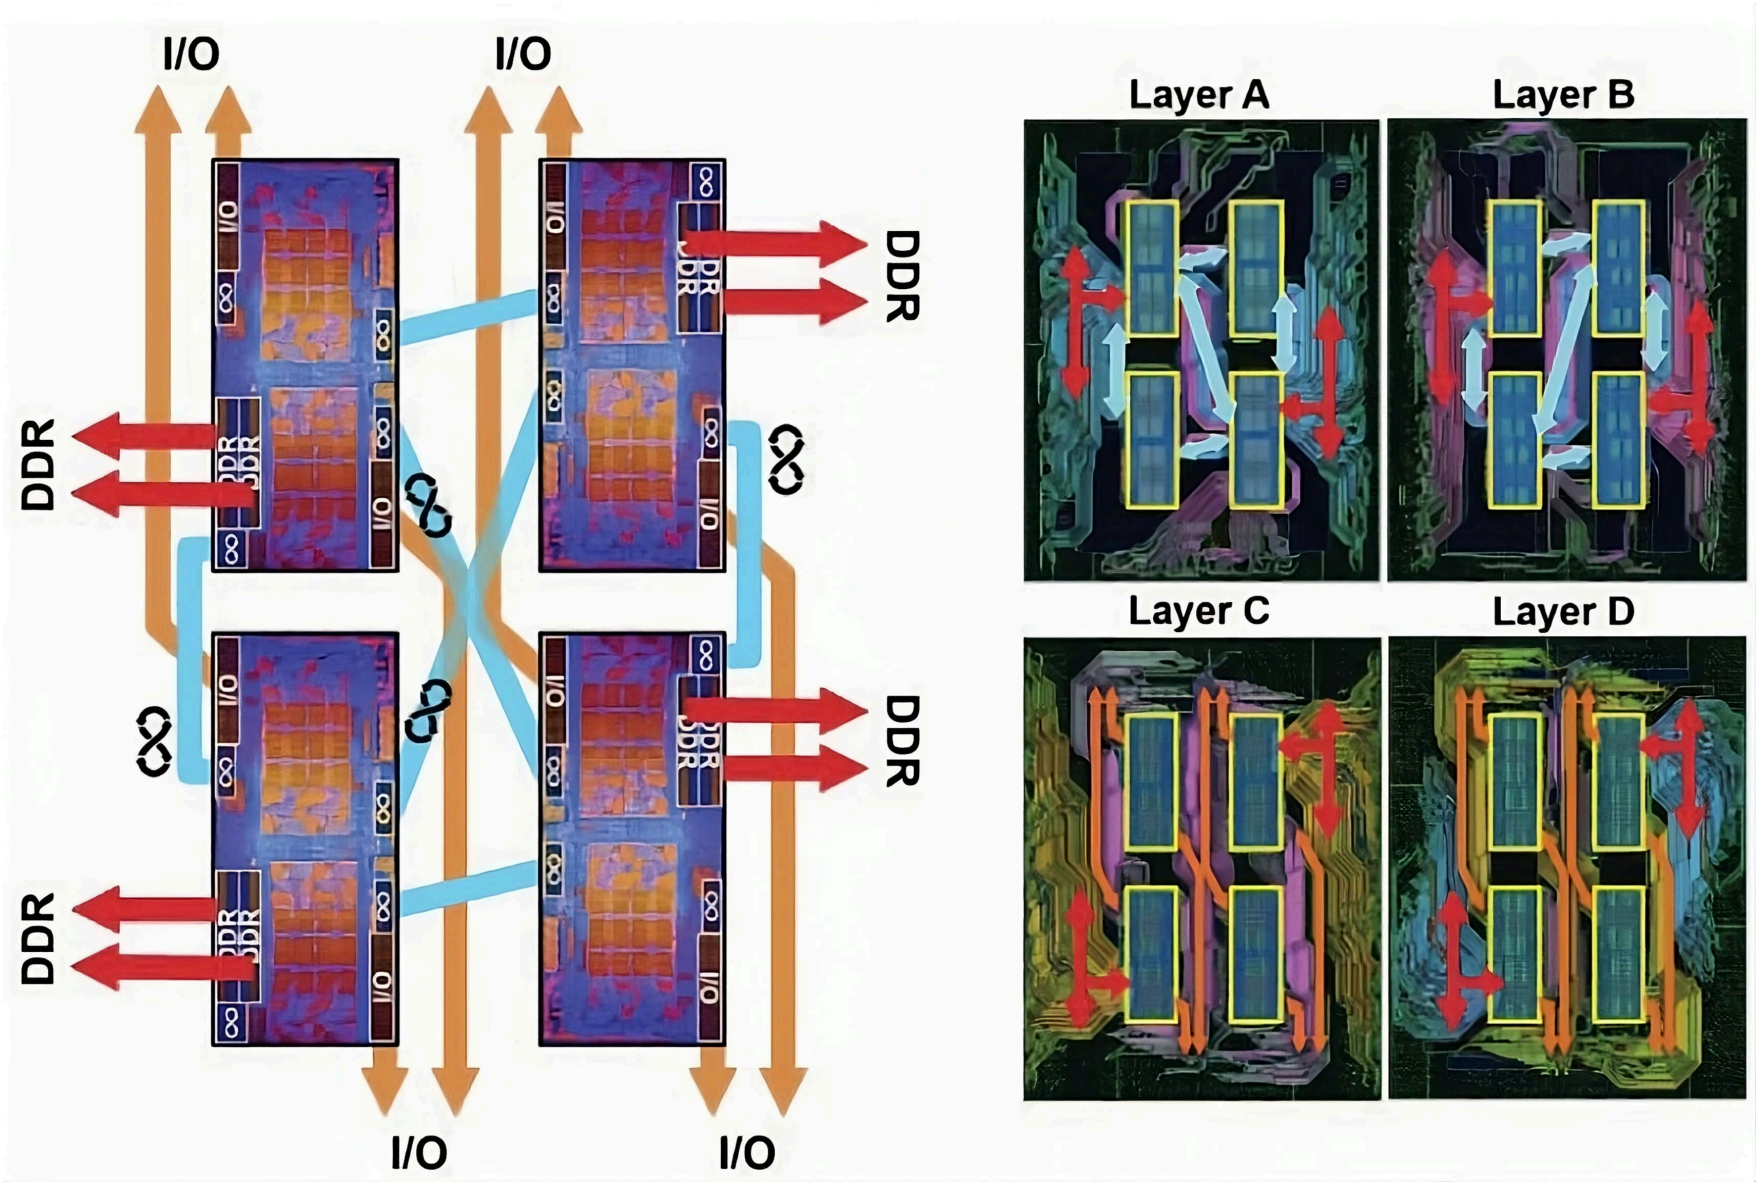

Figure 4: Multi-layer package routing, DDR (red), IO (orange), infinity-fabric (blue)[5]

[5] [1] S. Naffziger *et al.*, “Pioneering Chiplet Technology and Design for the AMD EPYC™ and Ryzen™ Processor Families : Industrial Product”, IEEE, Jun. 2021. doi: 10.1109/isca52012.2021.00014.

AMD Rome (2nd Gen. EPYC)[6]

[6] https://www.amd.com/content/dam/amd/en/documents/epyc-technical-docs/tuning-guides/amd-epyc-7002-tg-hpc-56827.pdf

Memory Access Latencies for Naples and Rome, Naffziger et al.[7]

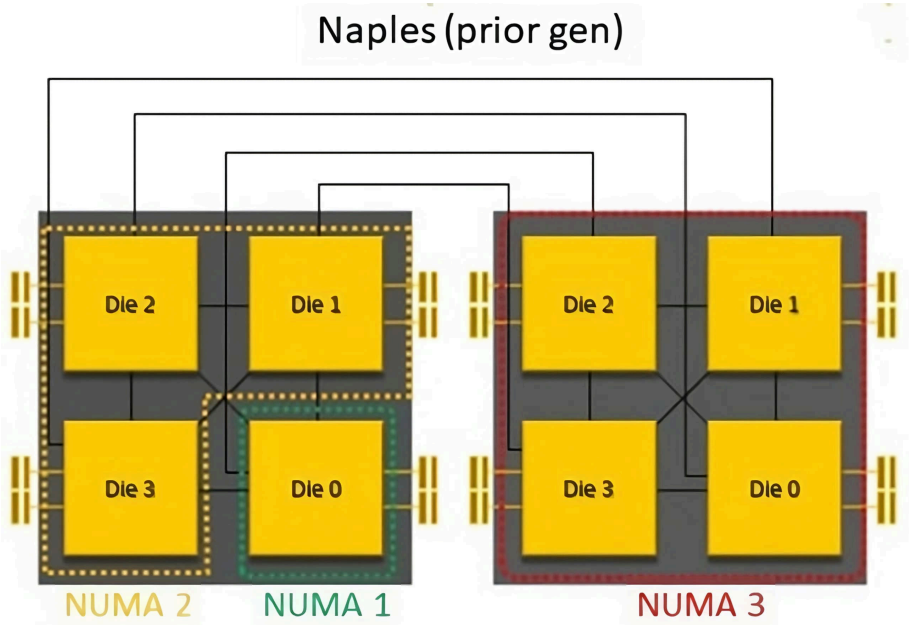

Domain	Latency (ns)
NUMA1	90
NUMA2	141
NUMA3	234
Avg. Local	128

3 NUMA Distances
8 NUMA Domains

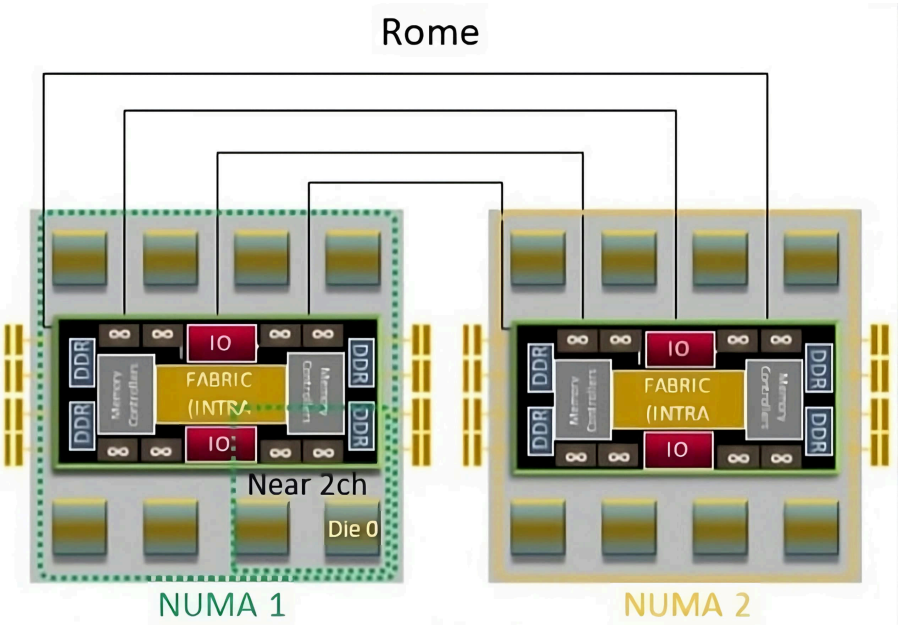

Domain	Latency (ns)
Near 2ch	94
NUMA1	104
NUMA2	201
Local NUMA1 avg. latency vs. Naples	.81X

2 NUMA Distances
2 NUMA Domains

[7] [2] S. Naffziger, K. Lepak, M. Paraschou, and M. Subramony, "2.2 AMD Chiplet Architecture for High-Performance Server and Desktop Products", IEEE, Feb. 2020. doi: 10.1109/isscc19947.2020.9063103.

Impact of NUMA on Applications

User Space Network Drivers

Paul Emmerich
Technical University of Munich
emmericp@net.in.tum.de

Maximilian Pudelko
Technical University of Munich
Open Networking Foundation
max@opennetworking.org

Simon Bauer
Technical University of Munich
bauersi@net.in.tum.de

Stefan Huber
Technical University of Munich
hubestef@net.in.tum.de

Thomas Zwickl
Technical University of Munich
zwickl@net.in.tum.de

Georg Carle
Technical University of Munich
carle@net.in.tum.de

Impact of NUMA on Applications

Emmerich et al.[8] – User Space Networking Drivers

Ingress NIC	Egress NIC	CPU	Memory	Throughput
Node 0	Node 0	Node 0	Node 0	10.8M pps
Node 0	Node 0	Node 0	Node 1	10.8M pps
Node 0	Node 0	Node 1	Node 0	7.6M pps
Node 0	Node 0	Node 1	Node 1	6.6M pps
Node 0	Node 1	Node 0	Node 0	7.9M pps
Node 0	Node 1	Node 0	Node 1	10.0M pps
Node 0	Node 1	Node 1	Node 0	8.6M pps
Node 0	Node 1	Node 1	Node 1	8.1M pps

Figure 7: Forwarding performance in packets per second, columns indicates pinning of each resource, based on Emmerich et al. [3]

[8] [3] P. Emmerich, M. Pudelko, S. Bauer, and G. Carle, “User Space Network Drivers”, in ANRW ‘18. Montreal, QC, Canada: Association for Computing Machinery, 2018, p. 91. doi: 10.1145/3232755.3232767.

Impact of NUMA on Applications

Tales of the Tail:
Hardware, OS, and Application-level Sources of Tail Latency

Jialin Li, Naveen Kr. Sharma, Dan R. K. Ports, and Steven D. Gribble

Department of Computer Science & Engineering, University of Washington

{lijl,naveenks,drkp,gribble}@cs.washington.edu

Impact of NUMA on Applications

Li et al.[9] – Memcached

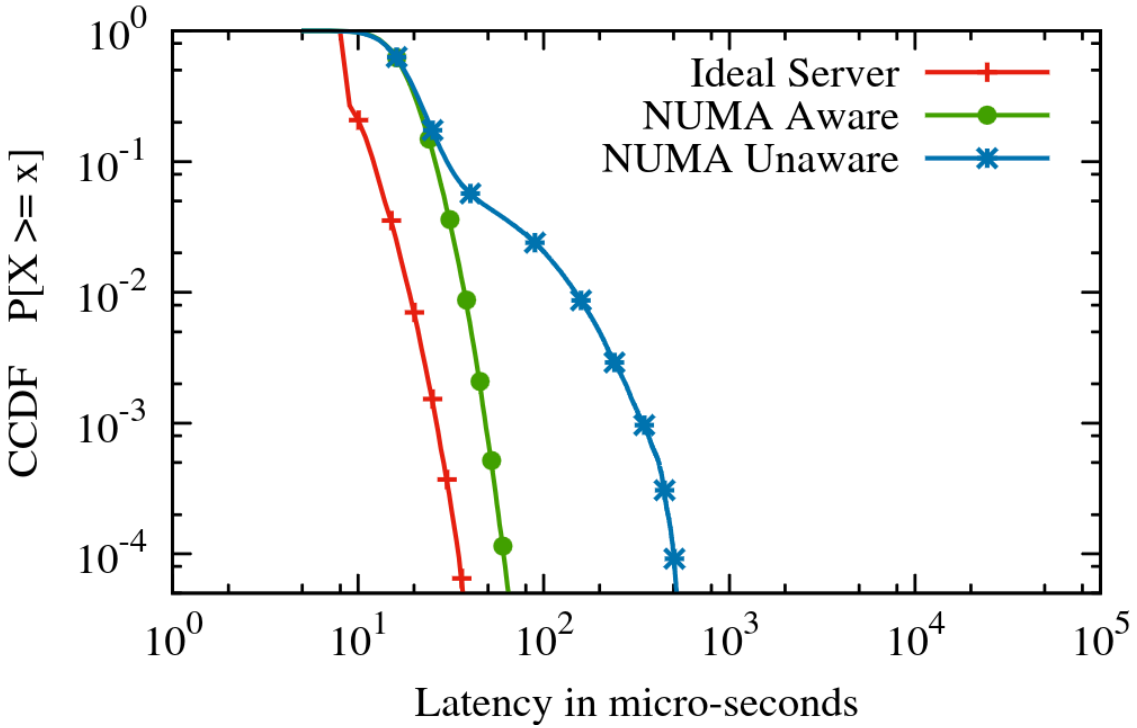

Figure 9: Memcached tail latency; 2 sockets; two instances (green), one instance (blue), based on Li et al. [4]

[9] [4] J. Li, N. K. Sharma, D. R. K. Ports, and S. D. Gribble, “Tales of the Tail”, ACM, Nov. 2014. doi: 10.1145/2670979.2670988.

Conclusion

- Chiplet technology is a fundamental part of future CPU architectures
- Inconsistent memory access latencies are a challenge for applications
- CPU architecture matters

Bibliography

[1] S. Naffziger *et al.*, “Pioneering Chiplet Technology and Design for the AMD EPYC™ and Ryzen™ Processor Families : Industrial Product”, IEEE, Jun. 2021. doi: 10.1109/isca52012.2021.00014.

[2] S. Naffziger, K. Lepak, M. Paraschou, and M. Subramony, “2.2 AMD Chiplet Architecture for High-Performance Server and Desktop Products”, IEEE, Feb. 2020. doi: 10.1109/isscc19947.2020.9063103.

[3] P. Emmerich, M. Pudelko, S. Bauer, and G. Carle, “User Space Network Drivers”, in ANRW '18. Montreal, QC, Canada: Association for Computing Machinery, 2018, p. 91. doi: 10.1145/3232755.3232767.

[4] J. Li, N. K. Sharma, D. R. K. Ports, and S. D. Gribble, “Tales of the Tail”, ACM, Nov. 2014. doi: 10.1145/2670979.2670988.